

# FROM STORAGE TO EXPERIENCE: A SURVEY ON THE EVOLUTION OF LLM AGENT MEMORY MECHANISMS

**Anonymous authors**

Paper under double-blind review

## 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.

## 1 INTRODUCTION

In recent years, the rapid advancement of Large Language Models (LLMs) has fundamentally reshaped the landscape of artificial intelligence (Touvron et al., 2023; Hurst et al., 2024; Yang et al., 2025a). 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 (Yao et al., 2022; Qin et al., 2024; Luo et al., 2025c). 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., 2023a; 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. (2025a).

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 (Packer et al., 2023; Hu et al., 2024; Kang et al., 2025), while the other draws inspiration from cognitive science and psychology to simulate mechanisms for the formation, consolidation, and retrieval of human memory (Zhong et al., 2023; Hou et al., 2024; Xu et al., 2025b). 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., 2025b; Yang et al., 2025b; Zhang et al., 2025j). Existing surveys have not sufficiently isolated these key technical drivers from general methodologies Wu et al. (2025c); Du et al. (2025c); Wu & Shu (2025); Cao 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.

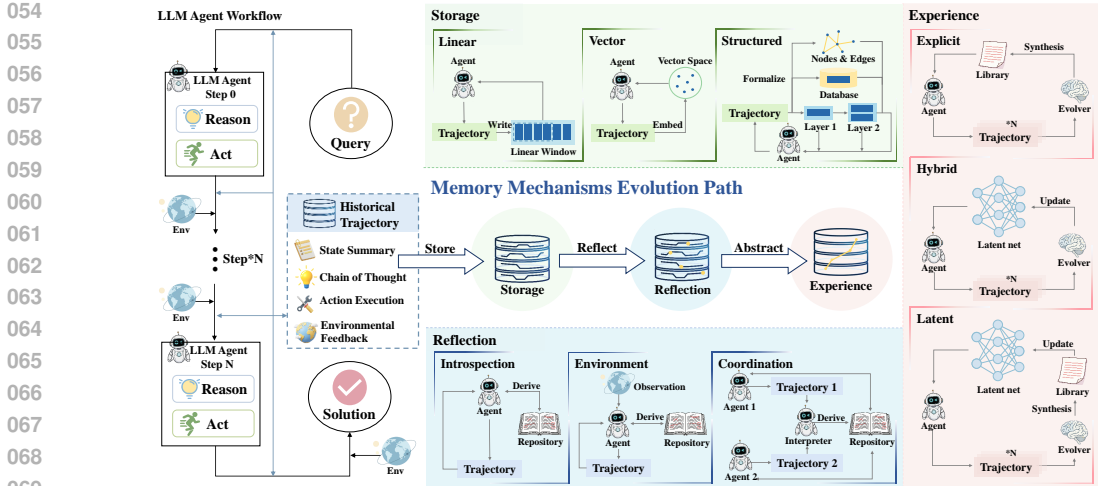


Figure 1: Overview of the LLM agent Memory Mechanisms System.

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. (2025b) 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 D, respectively.

## 2 BACKGROUND

### 2.1 THE LLM AGENT FRAMEWORK

We formalize an LLM-based agent as a decision-making entity parameterized by  $\theta$ , interacting with a dynamic environment  $\mathcal{E}$ . The agent’s operation is governed by a policy  $\pi_\theta$ , 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  $\theta$  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  $\tau$  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  $\tau_i$  to generate a refined memory unit  $m'_i$ , which encapsulates critiques or corrective insights:

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

where  $\phi$  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  $\tau_i$  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  $\mathcal{K}$ :

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

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

Formally,  $\mathcal{K}$  serves as a policy prior that elevates  $\pi_\theta$  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.

### 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., 2023b; Summers et al., 2023; Packer et al., 2023). First, these modules maintain internal states for reasoning to ensure the coherence of thought (Yao et al., 2023; Sun et al., 2025c); 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., 2025b); finally, they internalize interactions into persistent traits of the persona to ensure uniformity in behavior (Park et al., 2023; Westh ufer et al., 2025; Liang 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 (Huang et al., 2024; Everitt et al., 2025). Memory mechanisms mitigate this drift by providing persistent and explicit goals at a high level (Hu et al., 2024; Li et al., 2025e). 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., 2025c).

#### 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.

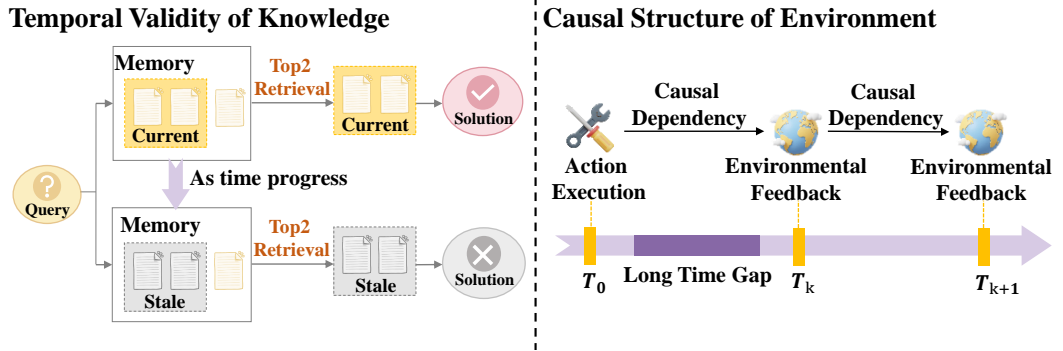


Figure 2: The Drivers in Dynamic Environments.

**The Temporal Validity of Knowledge.** Knowledge within environments of a dynamic nature is typically conditional rather than eternally valid (Lazaridou et al., 2021; Jang et al., 2022; Ko et al., 2024). 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 (Luu et al., 2022; Kalai & Vempala, 2023; Kasai et al., 2024); 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 (Zhong et al., 2023; Siyue et al., 2024; Salama et al., 2025; Du et al., 2025a; Houichime et al., 2025).

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

### 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 (Xiong et al., 2025; Srivastava & He, 2025). This necessitates the exploration of more strategic policies for the addition and deletion of information within memory mechanisms Du et al. (2025c); Liu et al. (2025e).

**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 (Tang et al., 2025; Cai et al., 2025; Xia et al., 2025b; Alakuijala et al., 2025; Guo 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  $\tau_i$  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 (Ratner et al., 2022; Xiao et al., 2023; Jin et al., 2024), as well as the achievement of information sparsification through the mechanical reduction of noise (Zhang et al., 2023b; Jiang et al., 2023; Xiao et al., 2024).

**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 (Melz, 2023; Liu et al., 2024; Das et al., 2024) as well as weighted retrieval that incorporates temporal decay and scores for importance (Zhong et al., 2023; Park 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; Xue et al., 2023; Lee & Ko, 2025), partition memory into distinct hierarchies to address the trade-off between storage capacity and speed of retrieval (Packer et al., 2023; Lu et al., 2023), and directly model the history of interaction as a topological network of entities and relations (Modarressi et al., 2024; Li 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 (Liu et al., 2023; Zhang et al., 2025h; Bohnet et al., 2025), the maintenance of the lifecycle of memory (Li et al., 2025a; Kang et al., 2025; Chhikara et al., 2025), and the compression and distillation of long trajectories Huang et al. (2025b); Han et al. (2025); Yang et al. (2025b); 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., 2025b;a) and calibrate internal models of the world (Sun et al., 2024; Xiao et al., 2025; Sun et al., 2025b).

**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 Bo et al. (2024); Balestri & Pescatore (2025); Wang et al. (2025c); Ozer 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 edible 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; Zhang et al., 2025b; Hassell et al., 2025; Wan et al., 2025) or directly abstracting them into executable entities (Wang et al., 2025e; Zhang et al., 2025a; Shi et al., 2025a).

**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; Zhai et al., 2025; Zhang et al., 2025f; Tandon et al., 2025; Yu et al., 2025b). 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., 2025d;c).

**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 (Wu et al., 2025a; Anonymous, 2025; Ouyang 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 (Zheng et al., 2024; Pan et al., 2025; Sun et al., 2025a); drivers based on curricula facilitate exploration tasks of increasing difficulty through the dynamic generation and adjustment of sequences for tasks (Wei et al., 2025b; Ahn et al., 2022); and drivers based on reuse enable exploration of high efficiency through the abstraction and reuse of trajectories from history (Wang et al., 2025d; Cai 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 (Qi et al., 2025; Zhai et al., 2025; Cheng et al., 2025a); 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 (Xia et al., 2025b; Liu et al., 2025a); 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 Shi et al. (2025b); Bidochko & Vyklyuk (2026).

## 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.

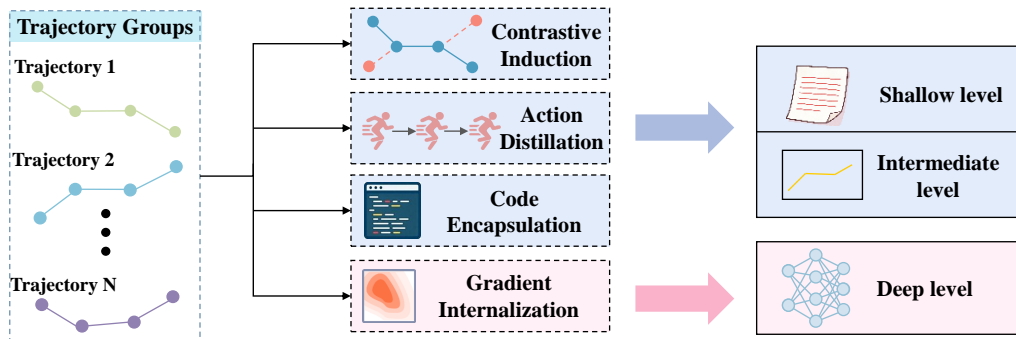


Figure 3: Overview of Cross-Trajectory Abstraction.

**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., 2024b); 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., 2025b; Latimer et al., 2025); the encapsulation of recurring patterns of behavior into program functions that are reusable by leveraging the compositionality of code (Yang et al., 2025c; Zhang et al., 2025e; Wang et al., 2025e); and the internalization of groups of trajectories into the parameters of the model through techniques for fine-tuning (Ding et al., 2025; Chen et al., 2025b).

**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., 2025c; Wei et al., 2025a; Hayashi et al., 2025); abstraction at the intermediate level completely removes redundancies of natural language, extracting only modular skeletons for execution as experience (Wang et al., 2024c; Liu et al., 2025g; Yu et al., 2025a); 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., 2025b; Luo et al., 2025b; Wang et al., 2025a).

### 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 (Xiong et al., 2025; Gharat et al., 2025; Srivastava & He, 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 (Wang et al., 2025b; Rasmussen et al., 2025). More importantly, the persistent retrieval of irrelevant or obsolete memories can disrupt the coherence in reasoning of the LLM agent Xu et al. (2025b); Tan 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., 2025a). 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 (Sun et al., 2025c; Zhang et al., 2025i; Nan et al., 2025).

**Benchmark for Experience.** Existing datasets primarily evaluate the capacity for retrieval and denoising 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 D). 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., 2025a). 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. (2025b). 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 (Tran et al., 2025; Liao 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 (Yuen et al., 2025; Rezaadeh 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 (Liu et al., 2025d; Zhou et al., 2025a; He 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., 2025b; 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.

## REFERENCES

- 486  
487  
488 Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea  
489 Finn, Keerthana Gopalakrishnan, Karol Hausman, Alexander Herzog, Daniel Ho, Jasmine Hsu,  
490 Julian Ibarz, Brian Ichter, Alex Irpan, Eric Jang, Rosario M Jauregui Ruano, Kyle Jeffrey, Sally Jes-  
491 month, Nikhil Jayant Joshi, Ryan C. Julian, Dmitry Kalashnikov, Yuheng Kuang, Kuang-Huei Lee,  
492 Sergey Levine, Yao Lu, Linda Luu, Carolina Parada, Peter Pastor, Jornell Quiambao, Kanishka Rao,  
493 Jarek Rettinghouse, Diego M Reyes, Pierre Sermanet, Nicolas Sievers, Clayton Tan, Alexander  
494 Toshev, Vincent Vanhoucke, F. Xia, Ted Xiao, Peng Xu, Sichun Xu, and Mengyuan Yan. Do as i  
495 can, not as i say: Grounding language in robotic affordances. In *Conference on Robot Learning*,  
496 2022. URL <https://api.semanticscholar.org/CorpusID:247939706>.
- 497 Qingyao Ai, Yichen Tang, Changyue Wang, Jianming Long, Weihang Su, and Yiqun Liu. Memory-  
498 bench: A benchmark for memory and continual learning in llm systems. *ArXiv*, abs/2510.17281,  
499 2025. URL <https://api.semanticscholar.org/CorpusID:282210434>.
- 500 Minttu Alakuijala, Ya Gao, Georgy Ananov, Samuel Kaski, Pekka Marttinen, Alexander Ilin, and  
501 Harri Valpola. Memento no more: Coaching ai agents to master multiple tasks via hints inter-  
502 nalization. *ArXiv*, abs/2502.01562, 2025. URL [https://api.semanticscholar.org/](https://api.semanticscholar.org/CorpusID:276107931)  
503 [CorpusID:276107931](https://api.semanticscholar.org/CorpusID:276107931).
- 504 Petr Anokhin, Nikita Semenov, Artyom Y. Sorokin, Dmitry Evseev, Mikhail Burtsev, and Evgeny  
505 Burnaev. Arigraph: Learning knowledge graph world models with episodic memory for llm  
506 agents. In *International Joint Conference on Artificial Intelligence*, 2024. URL [https://api.](https://api.semanticscholar.org/CorpusID:271039035)  
507 [semanticscholar.org/CorpusID:271039035](https://api.semanticscholar.org/CorpusID:271039035).
- 508 Anonymous. Exploratory memory-augmented LLM agent via hybrid on- and off-policy optimization.  
509 In *Submitted to The Fourteenth International Conference on Learning Representations*, 2025. URL  
510 <https://openreview.net/forum?id=UOzxviKVFO>. under review.
- 511  
512 Yushi Bai, Xin Lv, Jiajie Zhang, Hong Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao  
513 Liu, Aohan Zeng, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. Longbench: A bilingual,  
514 multitask benchmark for long context understanding. *ArXiv*, abs/2308.14508, 2023. URL [https://](https://api.semanticscholar.org/CorpusID:261245264)  
515 [api.semanticscholar.org/CorpusID:261245264](https://api.semanticscholar.org/CorpusID:261245264).
- 516 Yushi Bai, Shangqing Tu, Jiajie Zhang, Hao Peng, Xiaozhi Wang, Xin Lv, Shulin Cao, Jiazheng Xu,  
517 Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. Longbench v2: Towards deeper understanding  
518 and reasoning on realistic long-context multitasks. *ArXiv*, abs/2412.15204, 2024. URL [https://](https://api.semanticscholar.org/CorpusID:274859535)  
519 [api.semanticscholar.org/CorpusID:274859535](https://api.semanticscholar.org/CorpusID:274859535).
- 520 Roberto Balestri and Guglielmo Pescatore. Narrative memory in machines: Multi-agent arc extraction  
521 in serialized tv. *ArXiv*, abs/2508.07010, 2025. URL [https://api.semanticscholar.](https://api.semanticscholar.org/CorpusID:280565831)  
522 [org/CorpusID:280565831](https://api.semanticscholar.org/CorpusID:280565831).
- 523  
524 Ali Behrouz, Peilin Zhong, and Vahab S. Mirrokni. Titans: Learning to memorize at test time.  
525 *ArXiv*, abs/2501.00663, 2024. URL [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:275212078)  
526 [275212078](https://api.semanticscholar.org/CorpusID:275212078).
- 527 Andrii Bidochko and Yaroslav Vyklyuk. Thought management system for long-horizon, goal-  
528 driven llm agents. *Journal of Computational Science*, 93:102740, 2026. ISSN 1877-7503. doi:  
529 <https://doi.org/10.1016/j.jocs.2025.102740>. URL [https://www.sciencedirect.com/](https://www.sciencedirect.com/science/article/pii/S1877750325002170)  
530 [science/article/pii/S1877750325002170](https://www.sciencedirect.com/science/article/pii/S1877750325002170).
- 531  
532 Xiaohe Bo, Zeyu Zhang, Quanyu Dai, Xueyang Feng, Lei Wang, Rui Li, Xu Chen, and Ji-Rong  
533 Wen. Reflective multi-agent collaboration based on large language models. *Advances in Neural*  
534 *Information Processing Systems 37*, 2024. URL [https://api.semanticscholar.org/](https://api.semanticscholar.org/CorpusID:276318441)  
535 [CorpusID:276318441](https://api.semanticscholar.org/CorpusID:276318441).
- 536 Bernd Bohnet, P. Kamienny, Hanie Sedghi, Dilan Gorur, Pranjal Awasthi, Aaron T Parisi, Kevin  
537 Swersky, Rosanne Liu, Azade Nova, and Noah Fiedel. Enhancing llm planning capabilities through  
538 intrinsic self-critique. 2025. URL [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:284350956)  
539 [284350956](https://api.semanticscholar.org/CorpusID:284350956).

- 540 Zhicheng Cai, Xinyuan Guo, Yu Pei, Jiangtao Feng, Jiangjie Chen, Ya-Qin Zhang, Wei-Ying Ma,  
541 Mingxuan Wang, and Hao Zhou. Flex: Continuous agent evolution via forward learning from expe-  
542 rience. 2025. URL <https://api.semanticscholar.org/CorpusID:282912514>.  
543
- 544 Zouying Cao, Jiayi Deng, Li Yu, Weikang Zhou, Zhaoyang Liu, Bolin Ding, and Hai Zhao. Remember  
545 me, refine me: A dynamic procedural memory framework for experience-driven agent evolution.  
546 2025. URL <https://api.semanticscholar.org/CorpusID:283737683>.  
547
- 548 Ding Chen, Simin Niu, Kehang Li, Peng Liu, Xiangping Zheng, Bo Tang, Xinchu Li, Feiyu  
549 Xiong, and Zhiyu Li. Halumem: Evaluating hallucinations in memory systems of agents.  
550 *ArXiv*, abs/2511.03506, 2025a. URL <https://api.semanticscholar.org/CorpusID:282758378>.  
551
- 552 Shiqi Chen, Tongyao Zhu, Zian Wang, Jinghan Zhang, Kangrui Wang, Siyang Gao, Teng Xiao,  
553 Yee Whye Teh, Junxian He, and Manling Li. Internalizing world models via self-play finetuning  
554 for agentic rl. *ArXiv*, abs/2510.15047, 2025b. URL <https://api.semanticscholar.org/CorpusID:282203038>.  
555
- 556 Silin Chen, Shaoxin Lin, Xiaodong Gu, Yuling Shi, Heng Lian, Longfei Yun, Dong Chen, Weiguo  
557 Sun, Lin Cao, and Qianxiang Wang. Swe-exp: Experience-driven software issue resolution.  
558 *ArXiv*, abs/2507.23361, 2025c. URL <https://api.semanticscholar.org/CorpusID:280401697>.  
559
- 560 Jiali Cheng, Anjishnu Kumar, Roshan Lal, Rishi Rajasekaran, Hani Ramezani, Omar Zia Khan,  
561 Oleg Rokhlenko, Sunny Chiu-Webster, Gang Hua, and Hadi Amiri. Webatlas: An llm  
562 agent with experience-driven memory and action simulation. 2025a. URL <https://api.semanticscholar.org/CorpusID:282389238>.  
563
- 564 Mingyue Cheng, Ouyang Jie, Shuo Yu, Ruiran Yan, Yucong Luo, Zirui Liu, Daoyu Wang, Qi Liu,  
565 and Enhong Chen. Agent-rl: Training powerful llm agents with end-to-end reinforcement learning.  
566 2025b. URL <https://api.semanticscholar.org/CorpusID:283081049>.  
567
- 568 Prateek Chhikara, Dev Khant, Saket Aryan, Taranjeet Singh, and Deshraj Yadav. Mem0: Building  
569 production-ready ai agents with scalable long-term memory. *ArXiv*, abs/2504.19413, 2025. URL  
570 <https://api.semanticscholar.org/CorpusID:278165315>.  
571
- 572 Shaobo Cui, Luca Mouchel, and Boi Faltings. Uncertainty in causality: A new frontier. In *An-  
573 nual Meeting of the Association for Computational Linguistics*, 2025. URL <https://api.semanticscholar.org/CorpusID:280035098>.  
574
- 575 Payel Das, Subhajit Chaudhury, Elliot Nelson, Igor Melnyk, Sarath Swaminathan, Sihui Dai, Au-  
576 r lie C. Lozano, Georgios Kollias, Vijil Chenthamarakshan, Jir  Navr til, Soham Dan, and Pin-Yu  
577 Chen. Larimar: Large language models with episodic memory control. *ArXiv*, abs/2403.11901,  
578 2024. URL <https://api.semanticscholar.org/CorpusID:268532114>.  
579
- 580 Yang Deng, Xuan Zhang, Wenxuan Zhang, Yifei Yuan, See-Kiong Ng, and Tat-Seng Chua. On  
581 the multi-turn instruction following for conversational web agents. In *Annual Meeting of the  
582 Association for Computational Linguistics*, 2024. URL <https://api.semanticscholar.org/CorpusID:267897510>.  
583
- 584 Bowen Ding, Yuhang Chen, Jiayang Lv, Jiyao Yuan, Qi Zhu, Shuangshuang Tian, Dantong Zhu, Futing  
585 Wang, Heyuan Deng, Fei Mi, Lifeng Shang, and Tao Lin. Rethinking expert trajectory utilization  
586 in llm post-training. 2025. URL <https://api.semanticscholar.org/CorpusID:283883758>.  
587
- 588 Xingbo Du, Loka Li, Duzhen Zhang, and Le Song. Memr<sup>3</sup>: Memory retrieval via reflective rea-  
589 soning for llm agents. 2025a. URL <https://api.semanticscholar.org/CorpusID:284133206>.  
590
- 591 Xingbo Du, Loka Li, Duzhen Zhang, and Le Song. Memr<sup>3</sup>: Memory retrieval via reflective reasoning  
592 for llm agents, 2025b. URL <https://arxiv.org/abs/2512.20237>.  
593

- 594 Yiming Du, Hongru Wang, Zhengyi Zhao, Bin Liang, Baojun Wang, Wanjun Zhong, Zezhong Wang,  
595 and Kam-Fai Wong. *Perltqa: A personal long-term memory dataset for memory classification,*  
596 *retrieval, and synthesis in question answering.* *ArXiv*, abs/2402.16288, 2024. URL <https://api.semanticscholar.org/CorpusID:267938190>.  
597
- 598 Yiming Du, Wenyu Huang, Danna Zheng, Zhaowei Wang, Sébastien Montella, Mirella Lapata,  
599 Kam-Fai Wong, and Jeff Z. Pan. Rethinking memory in ai: Taxonomy, operations, topics, and  
600 future directions. *ArXiv*, abs/2505.00675, 2025c. URL <https://api.semanticscholar.org/CorpusID:278237720>.  
601  
602
- 603 Yiming Du, Baojun Wang, Yifan Xiang, Zhaowei Wang, Wenyu Huang, Boyang Xue, Bin Liang,  
604 Xingshan Zeng, Fei Mi, Haoli Bai, Lifeng Shang, Jeff Z. Pan, Yuxin Jiang, and Kam-Fai Wong.  
605 *Memory-t1: Reinforcement learning for temporal reasoning in multi-session agents.* 2025d. URL  
606 <https://api.semanticscholar.org/CorpusID:284132930>.  
607
- 608 Tom Everitt, Cristina Garbacea, Alexis Bellot, Jonathan Richens, Henry Papadatos, Siméon Campos,  
609 and Rohin Shah. Evaluating the goal-directedness of large language models. *ArXiv*, abs/2504.11844,  
610 2025. URL <https://api.semanticscholar.org/CorpusID:277824195>.
- 611 Jizhan Fang, Xinle Deng, Haoming Xu, Ziyang Jiang, Yuqi Tang, Ziwen Xu, Shumin Deng, Yunzhi  
612 Yao, Mengru Wang, Shuofei Qiao, Huajun Chen, and Ningyu Zhang. *Lightmem: Lightweight*  
613 *and efficient memory-augmented generation.* *ArXiv*, abs/2510.18866, 2025a. URL <https://api.semanticscholar.org/CorpusID:282245881>.  
614
- 615 Runnan Fang, Yuan Liang, Xiaobin Wang, Jialong Wu, Shuofei Qiao, Pengjun Xie, Fei Huang, Huajun  
616 Chen, and Ningyu Zhang. *Memp: Exploring agent procedural memory.* *ArXiv*, abs/2508.06433,  
617 2025b. URL <https://api.semanticscholar.org/CorpusID:280561810>.  
618
- 619 Lang Feng, Zhenghai Xue, Tingcong Liu, and Bo An. Group-in-group policy optimization for llm  
620 agent training. *ArXiv*, abs/2505.10978, 2025a. URL <https://api.semanticscholar.org/CorpusID:278715074>.  
621  
622
- 623 Tongtong Feng, Xin Wang, Yu-Gang Jiang, and Wenwu Zhu. Embodied ai: From llms to world  
624 models. *ArXiv*, abs/2509.20021, 2025b. URL <https://api.semanticscholar.org/CorpusID:281505303>.  
625
- 626 Saman Forouzandeh, Wei Peng, Parham Moradi, Xinghuo Yu, and Mahdi Jalili. Learning hierarchical  
627 procedural memory for llm agents through bayesian selection and contrastive refinement. 2025.  
628 URL <https://api.semanticscholar.org/CorpusID:284077809>.  
629
- 630 Dayuan Fu, Keqing He, Yejie Wang, Wentao Hong, Zhuoma Gongque, Weihao Zeng, Wei Wang,  
631 Jingang Wang, Xunliang Cai, and Weiran Xu. *Agentrefine: Enhancing agent generalization through*  
632 *refinement tuning.* *ArXiv*, abs/2501.01702, 2025. URL <https://api.semanticscholar.org/CorpusID:275323978>.  
633
- 634 Dawei Gao, Zitao Li, Weirui Kuang, Xuchen Pan, Daoyuan Chen, Zhijian Ma, Bingchen Qian,  
635 Liuyi Yao, Lin Zhu, Chen Cheng, Hongzhu Shi, Yaliang Li, Bolin Ding, and Jingren Zhou.  
636 *Agentscope: A flexible yet robust multi-agent platform.* *ArXiv*, abs/2402.14034, 2024. URL  
637 <https://api.semanticscholar.org/CorpusID:267782737>.  
638
- 639 Himanshu Gharat, Himanshi Agrawal, and Gourab K. Patro. From personalization to prejudice:  
640 Bias and discrimination in memory-enhanced ai agents for recruitment. 2025. URL <https://api.semanticscholar.org/CorpusID:283933650>.  
641
- 642 Amirhosein Ghasemabadi and Di Niu. Can llms predict their own failures? self-awareness via internal  
643 circuits. 2025. URL <https://api.semanticscholar.org/CorpusID:284133177>.  
644
- 645 Jiacheng Guo, Ling Yang, Peter Chen, Qixin Xiao, Yinjie Wang, Xinzhe Juan, Jiahao Qiu, Ke Shen,  
646 and Mengdi Wang. *Genenv: Difficulty-aligned co-evolution between llm agents and envi-*  
647 *ronment simulators.* 2025. URL <https://api.semanticscholar.org/CorpusID:284077413>.

- 648 Dongge Han, Camille Couturier, Daniel Madrigal Diaz, Xuchao Zhang, Victor Rühle, and Saravan  
649 Rajmohan. Legomem: Modular procedural memory for multi-agent llm systems for workflow  
650 automation. *ArXiv*, abs/2510.04851, 2025. URL [https://api.semanticscholar.org/  
651 CorpusID:281843936](https://api.semanticscholar.org/CorpusID:281843936).
- 652 Jackson Hassell, Dan Zhang, Han Jun Kim, Tom Mitchell, and Estevam Hruschka. Learning  
653 from supervision with semantic and episodic memory: A reflective approach to agent adaptation.  
654 *ArXiv*, abs/2510.19897, 2025. URL [https://api.semanticscholar.org/CorpusID:  
655 282304825](https://api.semanticscholar.org/CorpusID:282304825).
- 656 Hiroaki Hayashi, Bo Pang, Wenting Zhao, Ye Liu, Akash Gokul, Srijan Bansal, Caiming Xiong,  
657 Semih Yavuz, and Yingbo Zhou. Self-abstraction from grounded experience for plan-guided  
658 policy refinement. 2025. URL [https://api.semanticscholar.org/CorpusID:  
659 282911910](https://api.semanticscholar.org/CorpusID:282911910).
- 660 Junqing He, Liang Zhu, Rui Wang, Xi Wang, Gholamreza Haffari, and Jiaying Zhang. Madial-  
661 bench: Towards real-world evaluation of memory-augmented dialogue generation. In *North  
662 American Chapter of the Association for Computational Linguistics*, 2024a. URL [https://  
663 api.semanticscholar.org/CorpusID:272827841](https://api.semanticscholar.org/CorpusID:272827841).
- 664 Kaiyu He, Mian Zhang, Shuo Yan, Peilin Wu, and Zhiyu Chen. Idea: Enhancing the rule  
665 learning ability of large language model agent through induction, deduction, and abduction.  
666 In *Annual Meeting of the Association for Computational Linguistics*, 2024b. URL [https://  
667 api.semanticscholar.org/CorpusID:273025550](https://api.semanticscholar.org/CorpusID:273025550).
- 668 Xingqi He, Yujie Zhang, Shuyong Gao, Wenjie Li, Lingyi Hong, Mingxi Chen, Kaixun Jiang,  
669 Jiyuan Fu, and Wenqiang Zhang. Rsagent: Learning to reason and act for text-guided segmen-  
670 tation via multi-turn tool invocations. 2025. URL [https://api.semanticscholar.org/  
671 CorpusID:284350430](https://api.semanticscholar.org/CorpusID:284350430).
- 672 Matthew Ho, Chen Si, Zhaoxiang Feng, Fangxu Yu, Yichi Yang, Zhijian Liu, Zhiting Hu, and Lianhui  
673 Qin. Arcmemo: Abstract reasoning composition with lifelong llm memory. *ArXiv*, abs/2509.04439,  
674 2025. URL <https://api.semanticscholar.org/CorpusID:281103677>.
- 675 Chuanyang Hong and Qingyun He. Enhancing memory retrieval in generative agents through  
676 llm-trained cross attention networks. *Frontiers in Psychology*, 16, 2025. URL [https://api.  
677 semanticscholar.org/CorpusID:278420582](https://api.semanticscholar.org/CorpusID:278420582).
- 678 Yuki Hou, Haruki Tamoto, and Homei Miyashita. "my agent understands me better": Integrating  
679 dynamic human-like memory recall and consolidation in llm-based agents. *Extended Abstracts  
680 of the CHI Conference on Human Factors in Computing Systems*, 2024. URL [https://api.  
681 semanticscholar.org/CorpusID:268819055](https://api.semanticscholar.org/CorpusID:268819055).
- 682 Tarik Houichime, Abdelghani Souhar, and Younès El Amrani. Memory as resonance: A biomimetic  
683 architecture for infinite context memory on ergodic phonetic manifolds. 2025. URL [https://  
684 api.semanticscholar.org/CorpusID:284132830](https://api.semanticscholar.org/CorpusID:284132830).
- 685 Cheng-Ping Hsieh, Simeng Sun, Samuel Kriman, Shantanu Acharya, Dima Rekesh, Fei Jia, and  
686 Boris Ginsburg. Ruler: What's the real context size of your long-context language models?  
687 *ArXiv*, abs/2404.06654, 2024. URL [https://api.semanticscholar.org/CorpusID:  
688 269032933](https://api.semanticscholar.org/CorpusID:269032933).
- 689 Chenxu Hu, Jie Fu, Chenzhuang Du, Simian Luo, Junbo Jake Zhao, and Hang Zhao. Chatdb:  
690 Augmenting llms with databases as their symbolic memory. *ArXiv*, abs/2306.03901, 2023. URL  
691 <https://api.semanticscholar.org/CorpusID:259088875>.
- 692 Mengkang Hu, Tianxing Chen, Qiguang Chen, Yao Mu, Wenqi Shao, and Ping Luo. Hiagent:  
693 Hierarchical working memory management for solving long-horizon agent tasks with large lan-  
694 guage model. *ArXiv*, abs/2408.09559, 2024. URL [https://api.semanticscholar.org/  
695 CorpusID:271903203](https://api.semanticscholar.org/CorpusID:271903203).
- 696 Yuanzhe Hu, Yu Wang, and Julian McAuley. Evaluating memory in llm agents via incremental multi-  
697 turn interactions. *ArXiv*, abs/2507.05257, 2025a. URL [https://api.semanticscholar.  
698 org/CorpusID:280136659](https://api.semanticscholar.org/CorpusID:280136659).
- 700  
701

- 702 Yuyang Hu, Shichun Liu, Yanwei Yue, Guibin Zhang, Boyang Liu, Fangyi Zhu, Jiahang Lin, Honglin  
703 Guo, Shihan Dou, Zhiheng Xi, Senjie Jin, Jiejun Tan, Yanbin Yin, Jiongnan Liu, Zeyu Zhang,  
704 Zhongxiang Sun, Yutao Zhu, Hao Sun, Boci Peng, Zhenrong Cheng, Xuanbo Fan, Jiabin Guo,  
705 Xinlei Yu, Zhenhong Zhou, Zewen Hu, Jiahao Huo, Junhao Wang, Yuwei Niu, Yu Wang, Zhenfei  
706 Yin, Xiaobin Hu, Yue Liao, Qiankun Li, Kun Wang, Wangchunshu Zhou, Yixin Liu, Dawei Cheng,  
707 Qi Zhang, Tao Gui, Shirui Pan, Yan Zhang, Philip Torr, Zhicheng Dou, Ji-Rong Wen, Xuanjing  
708 Huang, Yu-Gang Jiang, and Shuicheng Yan. Memory in the age of ai agents. 2025b. URL  
709 <https://api.semanticscholar.org/CorpusID:283897233>.
- 710 Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song,  
711 and Denny Zhou. Large language models cannot self-correct reasoning yet. *ArXiv*, abs/2310.01798,  
712 2023a. URL <https://api.semanticscholar.org/CorpusID:263609132>.
- 713  
714 Xu Huang, Weiwen Liu, Xiaolong Chen, Xingmei Wang, Hao Wang, Defu Lian, Yasheng Wang,  
715 Ruiming Tang, and Enhong Chen. Understanding the planning of llm agents: A survey.  
716 *ArXiv*, abs/2402.02716, 2024. URL <https://api.semanticscholar.org/CorpusID:267411892>.
- 717  
718 Xu Huang, Junwu Chen, Yuxing Fei, Zhuohan Li, Philippe Schwaller, and Gerbrand Ceder. Cascade:  
719 Cumulative agentic skill creation through autonomous development and evolution. 2025a. URL  
720 <https://api.semanticscholar.org/CorpusID:284350716>.
- 721  
722 Yizhe Huang, Yang Liu, Ruiyu Zhao, Xiaolong Zhong, Xingming Yue, and Ling Jiang. Mem-  
723 orb: A plug-and-play verbal-reinforcement memory layer for e-commerce customer service.  
724 *ArXiv*, abs/2509.18713, 2025b. URL <https://api.semanticscholar.org/CorpusID:281496282>.
- 725  
726 Yunpeng Huang, Jingwei Xu, Zixu Jiang, Junyu Lai, Zenan Li, Yuan Yao, Taolue Chen, Lijuan  
727 Yang, Zhou Xin, and Xiaoxing Ma. Advancing transformer architecture in long-context large  
728 language models: A comprehensive survey. *ArXiv*, abs/2311.12351, 2023b. URL <https://api.semanticscholar.org/CorpusID:265308945>.
- 729  
730 Zhaopei Huang, Qifeng Dai, Guozheng Wu, Xiaopeng Wu, Kehan Chen, Chuan Yu, Xubin  
731 Li, Tiezheng Ge, Wenxuan Wang, and Qin Jin. Mem-pal: Towards memory-based personal-  
732 ized dialogue assistants for long-term user-agent interaction. 2025c. URL <https://api.semanticscholar.org/CorpusID:283073534>.
- 733  
734  
735  
736 Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Os-  
737 trow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Madry, Alex Baker-Whitcomb, Alex  
738 Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, Alex Paino, Alex  
739 Renzin, Alex Tachard Passos, Alexander Kirillov, Alexi Christakis, Alexis Conneau, Ali Kamali,  
740 Allan Jabri, Allison Moyer, Allison Tam, Amadou Crookes, Amin Tootoonchian, Ananya Kumar,  
741 Andrea Vallone, Andrej Karpathy, Andrew Braunstein, Andrew Cann, Andrew Codisoti, Andrew  
742 Galu, Andrew Kondrich, Andrew Tulloch, Andrey Mishchenko, Angela Baek, Angela Jiang, An-  
743 toine Pelisse, Antonia Woodford, Anuj Gosalia, Arka Dhar, Ashley Pantuliano, Avi Nayak, Avital  
744 Oliver, Barret Zoph, Behrooz Ghorbani, Ben Leimberger, Ben Rossen, Ben Sokolowsky, Ben  
745 Wang, Benjamin Zweig, Beth Hoover, Blake Samic, Bob McGrew, Bobby Spero, Bogo Gertler,  
746 Bowen Cheng, Brad Lightcap, Brandon Walkin, Brendan Quinn, Brian Guarraci, Brian Hsu, Bright  
747 Kellogg, Brydon Eastman, Camillo Lugaresi, Carroll L. Wainwright, Cary Bassin, Cary Hudson,  
748 Casey Chu, Chad Nelson, Chak Li, Chan Jun Shern, Channing Conger, Charlotte Barette, Chelsea  
749 Voss, Chen Ding, Cheng Lu, Chong Zhang, Chris Beaumont, Chris Hallacy, Chris Koch, Christian  
750 Gibson, Christina Kim, Christine Choi, Christine McLeavey, Christopher Hesse, Claudia Fischer,  
751 Clemens Winter, Coley Czarnecki, Colin Jarvis, Colin Wei, Constantin Koumouzelis, and Dane  
752 Sherburn. Gpt-4o system card. *CoRR*, abs/2410.21276, 2024. doi: 10.48550/ARXIV.2410.21276.  
753 URL <https://doi.org/10.48550/arXiv.2410.21276>.
- 754  
755 Joel Jang, Seonghyeon Ye, Changho Lee, Sohee Yang, Joongbo Shin, Janghoon Han, Gyeonghun  
Kim, and Minjoon Seo. Temporalwiki: A lifelong benchmark for training and evaluating ever-  
evolving language models. In *Conference on Empirical Methods in Natural Language Processing*,  
2022. URL <https://api.semanticscholar.org/CorpusID:248476156>.

- 756 Zixi Jia, Qinghua Liu, Hexiao Li, Yuyan Chen, and Jiqiang Liu. Evaluating the long-term  
757 memory of large language models. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova,  
758 and Mohammad Taher Pilehvar (eds.), *Findings of the Association for Computational Lin-*  
759 *guistics: ACL 2025*, pp. 19759–19777, Vienna, Austria, July 2025. Association for Compu-  
760 tational Linguistics. ISBN 979-8-89176-256-5. doi: 10.18653/v1/2025.findings-acl.1014. URL  
761 <https://aclanthology.org/2025.findings-acl.1014/>.
- 762 Bowen Jiang, Yuan Yuan, Maohao Shen, Zhuoqun Hao, Zhangchen Xu, Zichen Chen, Ziyi  
763 Liu, Anvesh Rao Vijjini, Jiashu He, Hanchao Yu, Radha Poovendran, Greg Wornell, Lyle  
764 Ungar, Dan Roth, Sihao Chen, and Camillo Jose Taylor. Personamem-v2: Towards person-  
765 alized intelligence via learning implicit user personas and agentic memory. 2025a. URL  
766 <https://api.semanticscholar.org/CorpusID:283693901>.
- 767 Dongfu Jiang, Yi Lu, Zhuofeng Li, Zhiheng Lyu, Ping Nie, Haozhe Wang, Alex Su, Hui Chen,  
768 Kai Zou, Chao Du, Tianyu Pang, and Wenhui Chen. Verltool: Towards holistic agentic re-  
769 inforcement learning with tool use. *ArXiv*, abs/2509.01055, 2025b. URL [https://api.](https://api.semanticscholar.org/CorpusID:281080546)  
770 [semanticscholar.org/CorpusID:281080546](https://api.semanticscholar.org/CorpusID:281080546).
- 771 Huiqiang Jiang, Qianhui Wu, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. LlmLingua: Compressing  
772 prompts for accelerated inference of large language models. In *Conference on Empirical Methods*  
773 *in Natural Language Processing*, 2023. URL [https://api.semanticscholar.org/](https://api.semanticscholar.org/CorpusID:263830701)  
774 [CorpusID:263830701](https://api.semanticscholar.org/CorpusID:263830701).
- 775 Hongye Jin, Xiaotian Han, Jingfeng Yang, Zhimeng Jiang, Zirui Liu, Chia yuan Chang, Huiyuan  
776 Chen, and Xia Hu. Llm maybe longlm: Self-extend llm context window without tuning.  
777 *ArXiv*, abs/2401.01325, 2024. URL [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:266725385)  
778 [266725385](https://api.semanticscholar.org/CorpusID:266725385).
- 779 Abhinav Joshi, Areeb Ahmad, and Ashutosh Modi. Cold: Causal reasoning in closed daily activities.  
780 *ArXiv*, abs/2411.19500, 2024. URL [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:274422558)  
781 [274422558](https://api.semanticscholar.org/CorpusID:274422558).
- 782 Adam Tauman Kalai and Santosh S. Vempala. Calibrated language models must hallucinate.  
783 *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, 2023. URL  
784 <https://api.semanticscholar.org/CorpusID:265445593>.
- 785 Jiazheng Kang, Mingming Ji, Zhe Zhao, and Ting Bai. Memory os of ai agent. *ArXiv*, abs/2506.06326,  
786 2025. URL <https://api.semanticscholar.org/CorpusID:279250574>.
- 787 Jungo Kasai, Keisuke Sakaguchi, Yoichi Takahashi, Ronan Le Bras, Akari Asai, Xinyan Yu, Dragomir  
788 Radev, Noah A. Smith, Yejin Choi, and Kentaro Inui. Realtime qa: What’s the answer right now?,  
789 2024. URL <https://arxiv.org/abs/2207.13332>.
- 790 Eunwon Kim, Chanho Park, and Buru Chang. Share: Shared memory-aware open-domain long-  
791 term dialogue dataset constructed from movie script. *ArXiv*, abs/2410.20682, 2024a. URL  
792 <https://api.semanticscholar.org/CorpusID:273654255>.
- 793 Jiho Kim, Woosog Chay, Hyeonji Hwang, Daeun Kyung, Hyunseung Chung, Eunbyeol Cho,  
794 Yohan Jo, and Edward Choi. Dialsim: A real-time simulator for evaluating long-term dia-  
795 logue understanding of conversational agents. *ArXiv*, abs/2406.13144, 2024b. URL [https:](https://api.semanticscholar.org/CorpusID:273234634)  
796 [//api.semanticscholar.org/CorpusID:273234634](https://api.semanticscholar.org/CorpusID:273234634).
- 797 Minsoo Kim and Seung won Hwang. Coex - co-evolving world-model and exploration.  
798 *ArXiv*, abs/2507.22281, 2025. URL [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:280391153)  
799 [280391153](https://api.semanticscholar.org/CorpusID:280391153).
- 800 Dayoon Ko, Jinyoung Kim, Hahyeon Choi, and Gunhee Kim. Growover: How can llms adapt to grow-  
801 ing real-world knowledge? In *Annual Meeting of the Association for Computational Linguistics*,  
802 2024. URL <https://api.semanticscholar.org/CorpusID:270371589>.
- 803 Yuri Kuratov, Aydar Bulatov, Petr Anokhin, Ivan Rodkin, Dmitry Sorokin, Artyom Y. Sorokin, and  
804 Mikhail Burtsev. Babilong: Testing the limits of llms with long context reasoning-in-a-haystack.  
805 *ArXiv*, abs/2406.10149, 2024. URL [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:270521583)  
806 [270521583](https://api.semanticscholar.org/CorpusID:270521583).

- 810 Chris Latimer, Nicol'o Boschi, Andrew Neeser, Chris Bartholomew, Gaurav Srivastava, Xuan Wang,  
811 and Naren Ramakrishnan. Hindsight is 20/20: Building agent memory that retains, recalls, and  
812 reflects. 2025. URL <https://api.semanticscholar.org/CorpusID:283897719>.  
813
- 814 Angeliki Lazaridou, Adhiguna Kuncoro, Elena Gribovskaya, Devang Agrawal, Adam Liska, Tayfun  
815 Terzi, Mai Giménez, Cyprien de Masson d'Autume, Tomás Kociský, Sebastian Ruder, Dani  
816 Yogatama, Kris Cao, Susannah Young, and Phil Blunsom. Mind the gap: Assessing temporal  
817 generalization in neural language models. In *Neural Information Processing Systems*, 2021. URL  
818 <https://api.semanticscholar.org/CorpusID:239886013>.
- 819 Seokhan Lee and Hanseok Ko. Training a team of language models as options to build an sql-  
820 based memory. *Applied Sciences*, 2025. URL [https://api.semanticscholar.org/](https://api.semanticscholar.org/CorpusID:282448398)  
821 [CorpusID:282448398](https://api.semanticscholar.org/CorpusID:282448398).  
822
- 823 Rui Li, Zeyu Zhang, Xiaohe Bo, Zihang Tian, Xu Chen, Quanyu Dai, Zhenhua Dong, and Ruiming  
824 Tang. Cam: A constructivist view of agentic memory for llm-based reading comprehension.  
825 *ArXiv*, abs/2510.05520, 2025a. URL [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:281886660)  
826 [281886660](https://api.semanticscholar.org/CorpusID:281886660).
- 827 Shilong Li, Yancheng He, Hangyu Guo, Xingyuan Bu, Ge Bai, Jie Liu, Jiaheng Liu, Xingwei Qu,  
828 Yangguang Li, Wanli Ouyang, Wenbo Su, and Bo Zheng. Graphreader: Building graph-based  
829 agent to enhance long-context abilities of large language models. In *Conference on Empirical*  
830 *Methods in Natural Language Processing*, 2024. URL [https://api.semanticscholar.](https://api.semanticscholar.org/CorpusID:270620354)  
831 [org/CorpusID:270620354](https://api.semanticscholar.org/CorpusID:270620354).
- 832 Xiaoxi Li, Wenxiang Jiao, Jiarui Jin, Guanting Dong, Jiajie Jin, Yinuo Wang, Hao Wang, Yutao Zhu,  
833 Ji-Rong Wen, Yuan Lu, and Zhicheng Dou. Deepagent: A general reasoning agent with scalable  
834 toolsets. *ArXiv*, abs/2510.21618, 2025b. URL [https://api.semanticscholar.org/](https://api.semanticscholar.org/CorpusID:282384893)  
835 [CorpusID:282384893](https://api.semanticscholar.org/CorpusID:282384893).  
836
- 837 Xintong Li, Jalend Bantupalli, Ria Dharmani, Yuwei Zhang, and Jingbo Shang. Toward multi-session  
838 personalized conversation: A large-scale dataset and hierarchical tree framework for implicit  
839 reasoning. *ArXiv*, abs/2503.07018, 2025c. URL [https://api.semanticscholar.org/](https://api.semanticscholar.org/CorpusID:276903780)  
840 [CorpusID:276903780](https://api.semanticscholar.org/CorpusID:276903780).
- 841 Zhiyu Li, Shichao Song, Chenyang Xi, Hanyu Wang, Chen Tang, Simin Niu, Ding Chen, Jiawei Yang,  
842 Chunyu Li, Qingchen Yu, Jihao Zhao, Yezhaohui Wang, Peng Liu, Zehao Lin, Pengyuan Wang,  
843 Jiahao Huo, Tianyi Chen, Kai Chen, Ke-Rong Li, Zhenzhen Tao, Junpeng Ren, Huayi Lai, Hao Wu,  
844 Bo Tang, Zhenren Wang, Zhaoxin Fan, Ningyu Zhang, Linfeng Zhang, Junchi Yan, Ming-Zhou  
845 Yang, Tong Xu, Wei Xu, Huajun Chen, Haofeng Wang, Hongkang Yang, Wentao Zhang, Zhikun  
846 Xu, Siheng Chen, and Feiyu Xiong. Memos: A memory os for ai system. *ArXiv*, abs/2507.03724,  
847 2025d. URL <https://api.semanticscholar.org/CorpusID:280093879>.  
848
- 849 Ziyue Li, Yuan Chang, Gaihong Yu, and Xiaoqiu Le. Hiplan: Hierarchical planning for llm-  
850 based agents with adaptive global-local guidance. *ArXiv*, abs/2508.19076, 2025e. URL [https://](https://api.semanticscholar.org/CorpusID:280870607)  
851 [api.semanticscholar.org/CorpusID:280870607](https://api.semanticscholar.org/CorpusID:280870607).
- 852 Jiafeng Liang, Hao Li, Chang Li, Jiaqi Zhou, Shixin Jiang, Zekun Wang, Changkai Ji, Zhihao Zhu,  
853 Runxuan Liu, Taolin Ren, Jinlan Fu, See-Kiong Ng, Xia Liang, Ming Liu, and Bing Qin. Ai  
854 meets brain: Memory systems from cognitive neuroscience to autonomous agents. 2025. URL  
855 <https://api.semanticscholar.org/CorpusID:284311192>.  
856
- 857 Callie C. Liao, Duoduo Liao, and Sai Surya Gadiraju. Agentmaster: A multi-agent conversa-  
858 tional framework using a2a and mcp protocols for multimodal information retrieval and analysis.  
859 *ArXiv*, abs/2507.21105, 2025. URL [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:280337378)  
860 [280337378](https://api.semanticscholar.org/CorpusID:280337378).
- 861 Bill Yuchen Lin, Yicheng Fu, Karina Yang, Prithviraj Ammanabrolu, Faeze Brahma, Shiyu Huang,  
862 Chandra Bhagavatula, Yejin Choi, and Xiang Ren. Swiftsage: A generative agent with fast  
863 and slow thinking for complex interactive tasks. *ArXiv*, abs/2305.17390, 2023. URL [https://](https://api.semanticscholar.org/CorpusID:258960143)  
[api.semanticscholar.org/CorpusID:258960143](https://api.semanticscholar.org/CorpusID:258960143).

- 864 Jiaqi Liu, Kaiwen Xiong, Peng Xia, Yiyang Zhou, Haonian Ji, Lu Feng, Siwei Han, Mingyu  
865 Ding, and Huaxiu Yao. Agent0-vl: Exploring self-evolving agent for tool-integrated vision-  
866 language reasoning. 2025a. URL [https://api.semanticscholar.org/CorpusID:  
867 283250484](https://api.semanticscholar.org/CorpusID:283250484).
- 868 Jiateng Liu, Zhenhailong Wang, Xiaojiang Huang, Yingjie Li, Xing Fan, Xiang Li, Chenlei Guo,  
869 Ruhi Sarikaya, and Heng Ji. Analyzing and internalizing complex policy documents for llm agents.  
870 *ArXiv*, abs/2510.11588, 2025b. URL [https://api.semanticscholar.org/CorpusID:  
871 282057538](https://api.semanticscholar.org/CorpusID:282057538).
- 872 Jun Liu, Zhenglun Kong, Changdi Yang, Fan Yang, Tianqi Li, Peiyan Dong, Joannah Nanjekye, Hao  
873 Tang, Geng Yuan, Wei Niu, Wenbin Zhang, Pu Zhao, Xue Lin, Dong-Xu Huang, and Yanzhi Wang.  
874 Rcr-router: Efficient role-aware context routing for multi-agent llm systems with structured memory.  
875 *ArXiv*, abs/2508.04903, 2025c. URL [https://api.semanticscholar.org/CorpusID:  
876 280546606](https://api.semanticscholar.org/CorpusID:280546606).
- 877 Junming Liu, Yifei Sun, Weihua Cheng, Haodong Lei, Yirong Chen, Licheng Wen, Xuemeng  
878 Yang, Daocheng Fu, Pinlong Cai, Nianchen Deng, Yi Yu, Shuyue Hu, Botian Shi, and Ding  
879 Wang. Memverse: Multimodal memory for lifelong learning agents. 2025d. URL [https:  
880 //api.semanticscholar.org/CorpusID:283466545](https://api.semanticscholar.org/CorpusID:283466545).
- 881 Lei Liu, Xiaoyan Yang, Yue Shen, Binbin Hu, Zhiqiang Zhang, Jinjie Gu, and Guannan  
882 Zhang. Think-in-memory: Recalling and post-thinking enable llms with long-term memory.  
883 *ArXiv*, abs/2311.08719, 2023. URL [https://api.semanticscholar.org/CorpusID:  
884 265212826](https://api.semanticscholar.org/CorpusID:265212826).
- 885 Shukai Liu, Jian Yang, Bo Jiang, Yizhi Li, Jinyang Guo, Xianglong Liu, and Bryan Dai. Context  
886 as a tool: Context management for long-horizon swe-agents. 2025e. URL [https://api.  
887 semanticscholar.org/CorpusID:284275525](https://api.semanticscholar.org/CorpusID:284275525).
- 888 Weijie Liu, Zecheng Tang, Juntao Li, Kehai Chen, and Min Zhang. Memlong: Memory-  
889 augmented retrieval for long text modeling. *ArXiv*, abs/2408.16967, 2024. URL [https:  
890 //api.semanticscholar.org/CorpusID:272310589](https://api.semanticscholar.org/CorpusID:272310589).
- 891 Xiaoyu Liu, Paiheng Xu, Junda Wu, Jiabin Yuan, Yifan Yang, Yuhang Zhou, Fuxiao Liu, Tianrui  
892 Guan, Haoliang Wang, Tong Yu, Julian J. McAuley, Wei Ai, and Furong Huang. Large language  
893 models and causal inference in collaboration: A comprehensive survey. *ArXiv*, abs/2403.09606,  
894 2025f. URL <https://api.semanticscholar.org/CorpusID:277245089>.
- 895 Yitao Liu, Chenglei Si, Karthik R. Narasimhan, and Shunyu Yao. Contextual experience replay  
896 for self-improvement of language agents. In *Annual Meeting of the Association for Computa-  
897 tional Linguistics*, 2025g. URL [https://api.semanticscholar.org/CorpusID:  
898 279250810](https://api.semanticscholar.org/CorpusID:279250810).
- 899 Xiaoxiao Long, Qingrui Zhao, Kaiwen Zhang, Zihao Zhang, Dingrui Wang, Yumeng Liu, Zhengjie  
900 Shu, Yi Lu, Shouzheng Wang, Xinzhe Wei, Wei Li, Wei Yin, Yao Yao, Jiangtian Pan, Qiu Shen,  
901 Ruigang Yang, Xun Cao, and Qionghai Dai. A survey: Learning embodied intelligence from  
902 physical simulators and world models. *ArXiv*, abs/2507.00917, 2025. URL [https://api.  
903 semanticscholar.org/CorpusID:280137292](https://api.semanticscholar.org/CorpusID:280137292).
- 904 Junru Lu, Siyu An, Mingbao Lin, Gabriele Pergola, Yulan He, Di Yin, Xing Sun, and Yunsheng  
905 Wu. Memochat: Tuning llms to use memos for consistent long-range open-domain conversation.  
906 *ArXiv*, abs/2308.08239, 2023. URL [https://api.semanticscholar.org/CorpusID:  
907 260926502](https://api.semanticscholar.org/CorpusID:260926502).
- 908 Hongyin Luo, Nathaniel Morgan, Tina Li, Derek Zhao, Ai Vy Ngo, Philip Schroeder, Lijie  
909 Yang, Assaf Ben-Kish, Jack O'Brien, and James R. Glass. Beyond context limits: Subcon-  
910 scious threads for long-horizon reasoning. *ArXiv*, abs/2507.16784, 2025a. URL [https:  
911 //api.semanticscholar.org/CorpusID:280675636](https://api.semanticscholar.org/CorpusID:280675636).
- 912 Xufang Luo, Yuge Zhang, Zhiyuan He, Zilong Wang, Siyun Zhao, Dongsheng Li, Luna K.  
913 Qiu, and Yuqing Yang. Agent lightning: Train any ai agents with reinforcement learning.  
914 *ArXiv*, abs/2508.03680, 2025b. URL [https://api.semanticscholar.org/CorpusID:  
915 280526917](https://api.semanticscholar.org/CorpusID:280526917).

- 918 Ziyang Luo, Zhiqi Shen, Wenzhuo Yang, Zirui Zhao, Prathyusha Jwalapuram, Amrita Saha, Doyen  
919 Sahoo, Silvio Savarese, Caiming Xiong, and Junnan Li. Mcp-universe: Benchmarking large  
920 language models with real-world model context protocol servers. *CoRR*, abs/2508.14704, 2025c.  
921 doi: 10.48550/ARXIV.2508.14704. URL [https://doi.org/10.48550/arXiv.2508.](https://doi.org/10.48550/arXiv.2508.14704)  
922 14704.
- 923 Kelvin Luu, Daniel Khashabi, Suchin Gururangan, Karishma Mandyam, and Noah A. Smith. Time  
924 waits for no one! analysis and challenges of temporal misalignment, 2022. URL <https://arxiv.org/abs/2111.07408>.
- 925  
926 Yuanjie Lyu, Chengyu Wang, Jun Huang, and Tong Xu. From correction to mastery: Reinforced  
927 distillation of large language model agents. *ArXiv*, abs/2509.14257, 2025. URL [https://api.](https://api.semanticscholar.org/CorpusID:281393943)  
928 [semanticscholar.org/CorpusID:281393943](https://api.semanticscholar.org/CorpusID:281393943).
- 929  
930 Adyasha Maharana, Dong-Ho Lee, S. Tulyakov, Mohit Bansal, Francesco Barbieri, and Yuwei Fang.  
931 Evaluating very long-term conversational memory of llm agents. *ArXiv*, abs/2402.17753, 2024.  
932 URL <https://api.semanticscholar.org/CorpusID:268041615>.
- 933  
934 Bodhisattwa Prasad Majumder, Bhavana Dalvi, Peter Alexander Jansen, Oyvind Tafjord, Niket  
935 Tandon, Li Zhang, Chris Callison-Burch, and Peter Clark. Clin: A continually learning language  
936 agent for rapid task adaptation and generalization. *ArXiv*, abs/2310.10134, 2023. URL [https://api.](https://api.semanticscholar.org/CorpusID:264146262)  
937 [semanticscholar.org/CorpusID:264146262](https://api.semanticscholar.org/CorpusID:264146262).
- 938 Eric Melz. Enhancing llm intelligence with arm-rag: Auxiliary rationale memory for retrieval aug-  
939 mented generation. *ArXiv*, abs/2311.04177, 2023. URL [https://api.semanticscholar.](https://api.semanticscholar.org/CorpusID:265043634)  
940 [org/CorpusID:265043634](https://api.semanticscholar.org/CorpusID:265043634).
- 941  
942 Atsuyuki Miyai, Zaiying Zhao, Kazuki Egashira, Atsuki Sato, Tatsumi Sunada, Shota Onohara,  
943 Hiromasa Yamanishi, Mashiro Toyooka, Kunato Nishina, Ryoma Maeda, Kiyoharu Aizawa, and  
944 T. Yamasaki. Webchorearena: Evaluating web browsing agents on realistic tedious web tasks.  
945 *ArXiv*, abs/2506.01952, 2025. URL [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:279119756)  
946 [279119756](https://api.semanticscholar.org/CorpusID:279119756).
- 947  
948 Ali Modarressi, Abdullatif Köksal, Ayyoob Imani, Mohsen Fayyaz, and Hinrich Schütze. Memllm:  
949 Finetuning llms to use an explicit read-write memory. *ArXiv*, abs/2404.11672, 2024. URL  
<https://api.semanticscholar.org/CorpusID:269214524>.
- 950  
951 Jiayan Nan, Wenquan Ma, Wenlong Wu, and Yize Chen. Nemori: Self-organizing agent mem-  
952 ory inspired by cognitive science. *ArXiv*, abs/2508.03341, 2025. URL [https://api.](https://api.semanticscholar.org/CorpusID:280526452)  
[semanticscholar.org/CorpusID:280526452](https://api.semanticscholar.org/CorpusID:280526452).
- 953  
954 Siru Ouyang, Jun Yan, I-Hung Hsu, Yanfei Chen, Ke Jiang, Zifeng Wang, Rujun Han, Long T. Le,  
955 Samira Daruki, Xiangru Tang, Vishy Tirumalashetty, George Lee, Mahsan Rofouei, Hangfei Lin,  
956 Jiawei Han, Chen-Yu Lee, and Tomas Pfister. Reasoningbank: Scaling agent self-evolving with  
957 reasoning memory. *ArXiv*, abs/2509.25140, 2025. URL [https://api.semanticscholar.](https://api.semanticscholar.org/CorpusID:281674540)  
[org/CorpusID:281674540](https://api.semanticscholar.org/CorpusID:281674540).
- 958  
959 Onat Ozer, Grace Wu, Yuchen Wang, Daniel Dosti, Honghao Zhang, and Vivi De La Rue.  
960 Mar:multi-agent reflexion improves reasoning abilities in llms. 2025. URL [https://api.](https://api.semanticscholar.org/CorpusID:284154082)  
961 [semanticscholar.org/CorpusID:284154082](https://api.semanticscholar.org/CorpusID:284154082).
- 962  
963 Charles Packer, Vivian Fang, Shishir G. Patil, Kevin Lin, Sarah Wooders, and Joseph Gonzalez.  
964 Memgpt: Towards llms as operating systems. *ArXiv*, abs/2310.08560, 2023. URL [https://api.](https://api.semanticscholar.org/CorpusID:263909014)  
965 [semanticscholar.org/CorpusID:263909014](https://api.semanticscholar.org/CorpusID:263909014).
- 966  
967 Yiyuan Pan, Zhe Liu, and Hesheng Wang. Wonder wins ways: Curiosity-driven exploration  
968 through multi-agent contextual calibration. *ArXiv*, abs/2509.20648, 2025. URL [https://api.](https://api.semanticscholar.org/CorpusID:281525435)  
[semanticscholar.org/CorpusID:281525435](https://api.semanticscholar.org/CorpusID:281525435).
- 969  
970 Joon Sung Park, Joseph C. O’Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and  
971 Michael S. Bernstein. Generative agents: Interactive simulacra of human behavior. *Proceedings*  
*of the 36th Annual ACM Symposium on User Interface Software and Technology*, 2023. URL  
<https://api.semanticscholar.org/CorpusID:258040990>.

- 972 Zehan Qi, Xiao Liu, Iat Long Iong, Hanyu Lai, Xueqiao Sun, Wenyi Zhao, Yu Yang, Xinyue  
973 Yang, Jiadai Sun, Shuntian Yao, Tianjie Zhang, Wei Xu, Jie Tang, and Yuxiao Dong. Webrl:  
974 Training llm web agents via self-evolving online curriculum reinforcement learning, 2025. URL  
975 <https://arxiv.org/abs/2411.02337>.
- 976 Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong,  
977 Xiangru Tang, Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian, Ruobing Xie, Jie Zhou,  
978 Mark Gerstein, Dahai Li, Zhiyuan Liu, and Maosong Sun. Toolllm: Facilitating large language  
979 models to master 16000+ real-world apis. In *The Twelfth International Conference on Learning  
980 Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL  
981 <https://openreview.net/forum?id=dHng200Jjr>.
- 982 Vishal Raman, R VijaiAravindh, and Abhijith Ragav. Remi: A novel causal schema memory  
983 architecture for personalized lifestyle recommendation agents. *ArXiv*, abs/2509.06269, 2025. URL  
984 <https://api.semanticscholar.org/CorpusID:280922051>.
- 985 Preston Rasmussen, Pavlo Paliychuk, Travis Beauvais, Jack Ryan, and Daniel Chalef. Zep: A  
986 temporal knowledge graph architecture for agent memory. *ArXiv*, abs/2501.13956, 2025. URL  
987 <https://api.semanticscholar.org/CorpusID:275907122>.
- 988 Nir Ratner, Yoav Levine, Yonatan Belinkov, Ori Ram, Inbal Magar, Omri Abend, Ehud Karpas,  
989 Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. Parallel context windows for large  
990 language models. In *Annual Meeting of the Association for Computational Linguistics*, 2022. URL  
991 <https://api.semanticscholar.org/CorpusID:258686160>.
- 992 Matthew Renze and Erhan Guven. Self-reflection in large language model agents: Effects on  
993 problem-solving performance. *2024 2nd International Conference on Foundation and Large  
994 Language Models (FLLM)*, pp. 516–525, 2024. URL <https://api.semanticscholar.org/CorpusID:275956778>.
- 995 Alireza Rezazadeh, Zichao Li, Ange Lou, Yuying Zhao, Wei Wei, and Yujia Bao. Collaborative mem-  
996 ory: Multi-user memory sharing in llm agents with dynamic access control. *ArXiv*, abs/2505.18279,  
997 2025. URL <https://api.semanticscholar.org/CorpusID:278904585>.
- 998 Rana Salama, Jason Cai, Michelle Yuan, Anna Currey, Monica Sunkara, Yi Zhang, and Yassine  
999 Benajiba. Meminsight: Autonomous memory augmentation for llm agents. *ArXiv*, abs/2503.21760,  
1000 2025. URL <https://api.semanticscholar.org/CorpusID:277349587>.
- 1001 Yuchen Shi, Yuzheng Cai, Siqi Cai, Zihan Xu, Lichao Chen, Yulei Qin, Zhijian Zhou, Xiang Fei,  
1002 Chaofan Qiu, Xiaoyu Tan, Gang Li, Zongyi Li, Haojia Lin, Guocan Cai, Yong Mao, Yunsheng  
1003 Wu, Ke Li, and Xing Sun. Youtu-agent: Scaling agent productivity with automated generation  
1004 and hybrid policy optimization. 2025a. URL <https://api.semanticscholar.org/CorpusID:284350437>.
- 1005 Zijing Shi, Meng Fang, and Ling Chen. Monte carlo planning with large language model for text-  
1006 based game agents. *ArXiv*, abs/2504.16855, 2025b. URL <https://api.semanticscholar.org/CorpusID:277999605>.
- 1007 Noah Shinn, Federico Cassano, Beck Labash, Ashwin Gopinath, Karthik Narasimhan, and Shunyu  
1008 Yao. Reflexion: language agents with verbal reinforcement learning. In *Neural Information  
1009 Processing Systems*, 2023. URL <https://api.semanticscholar.org/CorpusID:258833055>.
- 1010 Zhang Siyue, Yuxiang Xue, Yiming Zhang, Xiaobao Wu, Anh Tuan Luu, and Zhao Chen. Mrag: A  
1011 modular retrieval framework for time-sensitive question answering. *ArXiv*, abs/2412.15540, 2024.  
1012 URL <https://api.semanticscholar.org/CorpusID:274965713>.
- 1013 Saksham Sahai Srivastava and Haoyu He. Memorygraft: Persistent compromise of llm agents via  
1014 poisoned experience retrieval, 2025. URL <https://arxiv.org/abs/2512.16962>.
- 1015 Theodore R. Sumers, Shunyu Yao, Karthik Narasimhan, and Thomas L. Griffiths. Cognitive  
1016 architectures for language agents. *Trans. Mach. Learn. Res.*, 2024, 2023. URL <https://api.semanticscholar.org/CorpusID:261556862>.
- 1017  
1018  
1019  
1020  
1021  
1022  
1023  
1024  
1025

- 1026 Haoran Sun, Yekun Chai, Shuohuan Wang, Yu Sun, Hua Wu, and Haifeng Wang. Curiosity-driven  
1027 reinforcement learning from human feedback. In *Annual Meeting of the Association for Com-*  
1028 *putational Linguistics*, 2025a. URL [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:275758441)  
1029 275758441.
- 1030 Haoran Sun, Zekun Zhang, and Shaoning Zeng. Preference-aware memory update for long-term  
1031 llm agents. *ArXiv*, abs/2510.09720, 2025b. URL [https://api.semanticscholar.org/](https://api.semanticscholar.org/CorpusID:282058600)  
1032 [CorpusID:282058600](https://api.semanticscholar.org/CorpusID:282058600).
- 1033 Weiwei Sun, Miao Lu, Zhan Ling, Kang Liu, Xuesong Yao, Yiming Yang, and Jiecao Chen. Scaling  
1034 long-horizon llm agent via context-folding. *ArXiv*, abs/2510.11967, 2025c. URL [https://api.](https://api.semanticscholar.org/CorpusID:282064490)  
1035 [semanticscholar.org/CorpusID:282064490](https://api.semanticscholar.org/CorpusID:282064490).
- 1036 Zhiyuan Sun, Haochen Shi, Marc-Alexandre Côté, Glen Berseth, Xingdi Yuan, and Bang Liu.  
1037 Enhancing agent learning through world dynamics modeling. *ArXiv*, abs/2407.17695, 2024. URL  
1038 <https://api.semanticscholar.org/CorpusID:271431868>.
- 1039 Mirac Suzgun, Mert Yükeşgönül, Federico Bianchi, Daniel Jurafsky, and James Zou. Dynamic  
1040 cheatsheet: Test-time learning with adaptive memory. *ArXiv*, abs/2504.07952, 2025. URL  
1041 <https://api.semanticscholar.org/CorpusID:277667675>.
- 1042 Zhen Tan, Jun Yan, I-Hung Hsu, Rujun Han, Zifeng Wang, Long T. Le, Yiwen Song, Yanfei Chen,  
1043 Hamid Palangi, George Lee, Anand Iyer, Tianlong Chen, Huan Liu, Chen-Yu Lee, and Tomas  
1044 Pfister. In prospect and retrospect: Reflective memory management for long-term personalized  
1045 dialogue agents. *ArXiv*, abs/2503.08026, 2025. URL [https://api.semanticscholar.](https://api.semanticscholar.org/CorpusID:276928772)  
1046 [org/CorpusID:276928772](https://api.semanticscholar.org/CorpusID:276928772).
- 1047 Arnav Tandon, Karan Dalal, Xinhao Li, Daniel Kocaja, Marcel Rod, Sam Buchanan, Xiaolong  
1048 Wang, Jure Leskovec, Sanmi Koyejo, Tatsunori Hashimoto, Carlos Guestrin, Jed McCaleb, Yejin  
1049 Choi, and Yu Sun. End-to-end test-time training for long context. 2025. URL [https://api.](https://api.semanticscholar.org/CorpusID:284313495)  
1050 [semanticscholar.org/CorpusID:284313495](https://api.semanticscholar.org/CorpusID:284313495).
- 1051 Hao Tang, Darren Key, and Kevin Ellis. Worldcoder, a model-based llm agent: Building world  
1052 models by writing code and interacting with the environment. *ArXiv*, abs/2402.12275, 2024a. URL  
1053 <https://api.semanticscholar.org/CorpusID:267751283>.
- 1054 Jiaming Tang, Yilong Zhao, Kan Zhu, Guangxuan Xiao, Baris Kasikci, and Song Han. Quest:  
1055 Query-aware sparsity for efficient long-context llm inference. *ArXiv*, abs/2406.10774, 2024b. URL  
1056 <https://api.semanticscholar.org/CorpusID:270559146>.
- 1057 Xiangru Tang, Tianrui Qin, Tianhao Peng, Ziyang Zhou, Yanjun Shao, Tingting Du, Xinming Wei,  
1058 Peng Xia, Fang Wu, He Zhu, Ge Zhang, Jiaheng Liu, Xingyao Wang, Sirui Hong, Chenglin  
1059 Wu, Hao Cheng, Chi Wang, and Wangchunshu Zhou. Agent kb: Leveraging cross-domain  
1060 experience for agentic problem solving. *ArXiv*, abs/2507.06229, 2025. URL [https://api.](https://api.semanticscholar.org/CorpusID:280047833)  
1061 [semanticscholar.org/CorpusID:280047833](https://api.semanticscholar.org/CorpusID:280047833).
- 1062 Mohammad Tavakoli, Alireza Salemi, Carrie Ye, Mohamed Abdalla, Hamed Zamani, and J. Ross  
1063 Mitchell. Beyond a million tokens: Benchmarking and enhancing long-term memory in llms.  
1064 *ArXiv*, abs/2510.27246, 2025. URL [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:282719168)  
1065 [282719168](https://api.semanticscholar.org/CorpusID:282719168).
- 1066 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée  
1067 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand  
1068 Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language  
1069 models. *CoRR*, abs/2302.13971, 2023. doi: 10.48550/ARXIV.2302.13971. URL [https:](https://doi.org/10.48550/arXiv.2302.13971)  
1070 [//doi.org/10.48550/arXiv.2302.13971](https://doi.org/10.48550/arXiv.2302.13971).
- 1071 Khanh-Tung Tran, Dung Dao, Minh-Duong Nguyen, Quoc-Viet Pham, Barry O’Sullivan,  
1072 and Hoang D. Nguyen. Multi-agent collaboration mechanisms: A survey of llms.  
1073 *ArXiv*, abs/2501.06322, 2025. URL [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:275471465)  
1074 [275471465](https://api.semanticscholar.org/CorpusID:275471465).

- 1080 Ioannis C. Tsaknakis, Bingqing Song, Shuyu Gan, Dongyeop Kang, Alfredo García, Gaowen Liu,  
1081 Charles Fleming, and Mingyi Hong. Do llms recognize your latent preferences? a benchmark  
1082 for latent information discovery in personalized interaction. *ArXiv*, abs/2510.17132, 2025. URL  
1083 <https://api.semanticscholar.org/CorpusID:282209649>.  
1084
- 1085 Chunhui Wan, Xunan Dai, Zhuo Wang, Minglei Li, Yanpeng Wang, Yinan Mao, Yu Lan, and Zhiwen  
1086 Xiao. Loongflow: Directed evolutionary search via a cognitive plan-execute-summarize paradigm.  
1087 2025. URL <https://api.semanticscholar.org/CorpusID:284351394>.  
1088
- 1089 Luanbo Wan and Weizhi Ma. Storybench: A dynamic benchmark for evaluating long-term memory  
1090 with multi turns. *ArXiv*, abs/2506.13356, 2025. URL <https://api.semanticscholar.org/CorpusID:279402292>.  
1091
- 1092 Hengyi Wang, Haizhou Shi, Shiwei Tan, Weiyi Qin, Wenyuan Wang, Tunyu Zhang, Akshay Uttama  
1093 Nambi, Tanuja Ganu, and Hao Wang. Multimodal needle in a haystack: Benchmarking long-  
1094 context capability of multimodal large language models. In *North American Chapter of the*  
1095 *Association for Computational Linguistics*, 2024a. URL <https://api.semanticscholar.org/CorpusID:270559255>.  
1096
- 1097 Jiongxiao Wang, Qiaojing Yan, Yawei Wang, Yijun Tian, Soumya Smruti Mishra, Zhichao Xu,  
1098 Megha Gandhi, Panpan Xu, and Lin Lee Cheong. Reinforcement learning for self-improving  
1099 agent with skill library. 2025a. URL <https://api.semanticscholar.org/CorpusID:284058483>.  
1100  
1101
- 1102 Lei Wang, Chengbang Ma, Xueyang Feng, Zeyu Zhang, Hao ran Yang, Jingsen Zhang, Zhi-Yang  
1103 Chen, Jiakai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Ji rong Wen. A  
1104 survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18,  
1105 2023. URL <https://api.semanticscholar.org/CorpusID:261064713>.  
1106
- 1107 Piaohong Wang, Motong Tian, Jiaxian Li, Yuan Liang, Yuqing Wang, Qianben Chen, Tiannan  
1108 Wang, Zhicong Lu, Jiawei Ma, Yuchen Eleanor Jiang, and Wangchunshu Zhou. O-mem: Omni  
1109 memory system for personalized, long horizon, self-evolving agents. 2025b. URL <https://api.semanticscholar.org/CorpusID:283073241>.  
1110
- 1111 Rongzheng Wang, Shuang Liang, Qizhi Chen, Yihong Huang, Muquan Li, Yizhuo Ma, Dongyang  
1112 Zhang, Ke Qin, and Man-Fai Leung. Graphcogent: Mitigating llms’ working memory constraints  
1113 via multi-agent collaboration in complex graph understanding. 2025c. URL <https://api.semanticscholar.org/CorpusID:281682226>.  
1114
- 1115 Siyuan Wang, Zhongyu Wei, Yejin Choi, and Xiang Ren. Symbolic working memory enhances  
1116 language models for complex rule application. *ArXiv*, abs/2408.13654, 2024b. URL <https://api.semanticscholar.org/CorpusID:271957585>.  
1117  
1118
- 1119 Yu Wang and Xi Chen. Mirix: Multi-agent memory system for llm-based agents. *ArXiv*,  
1120 abs/2507.07957, 2025. URL <https://api.semanticscholar.org/CorpusID:280277519>.  
1121  
1122
- 1123 Zihan Wang, Kangrui Wang, Qineng Wang, Pingyue Zhang, Linjie Li, Zhengyuan Yang, Kefan  
1124 Yu, Minh Nhat Nguyen, Licheng Liu, Eli Gottlieb, Monica Lam, Yiping Lu, Kyunghyun Cho,  
1125 Jiajun Wu, Fei-Fei Li, Lijuan Wang, Yejin Choi, and Manling Li. Ragen: Understanding self-  
1126 evolution in llm agents via multi-turn reinforcement learning. *ArXiv*, abs/2504.20073, 2025d. URL  
1127 <https://api.semanticscholar.org/CorpusID:278170861>.  
1128
- 1129 Zora Zhiruo Wang, Jiayuan Mao, Daniel Fried, and Graham Neubig. Agent workflow memory.  
1130 *ArXiv*, abs/2409.07429, 2024c. URL <https://api.semanticscholar.org/CorpusID:272592995>.  
1131
- 1132 Zora Zhiruo Wang, Apurva Gandhi, Graham Neubig, and Daniel Fried. Inducing programmatic skills  
1133 for agentic tasks. *ArXiv*, abs/2504.06821, 2025e. URL <https://api.semanticscholar.org/CorpusID:277634286>.

- 1134 Tianxin Wei, Noveen Sachdeva, Benjamin Coleman, Zhankui He, Yuanchen Bei, Xuying Ning,  
1135 Mengting Ai, Yunzhe Li, Jingrui He, Ed Huai hsin Chi, Chi Wang, Shuo Chen, Fernando Pereira,  
1136 Wang-Cheng Kang, and Derek Zhiyuan Cheng. Evo-memory: Benchmarking llm agent test-time  
1137 learning with self-evolving memory. 2025a. URL [https://api.semanticscholar.org/  
1138 CorpusID:283261706](https://api.semanticscholar.org/CorpusID:283261706).
- 1139 Zhepei Wei, Wenlin Yao, Yao Liu, Weizhi Zhang, Qin Lu, Liang Qiu, Changlong Yu, Puyang Xu,  
1140 Chao Zhang, Bing Yin, Hyokun Yun, and Lihong Li. Webagent-rl: Training web agents via  
1141 end-to-end multi-turn reinforcement learning. *ArXiv*, abs/2505.16421, 2025b. URL [https:  
1142 //api.semanticscholar.org/CorpusID:278788476](https://api.semanticscholar.org/CorpusID:278788476).
- 1144 Rebecca Westh auer, Wolfgang Minker, and Sebastian Zepf. Enabling personalized long-term inter-  
1145 actions in llm-based agents through persistent memory and user profiles. *ArXiv*, abs/2510.07925,  
1146 2025. URL <https://api.semanticscholar.org/CorpusID:281951598>.
- 1147 Cheng-Kuang Wu, Zhi Rui Tam, Chieh-Yen Lin, Yun-Nung Chen, and Hung yi Lee. Streambench:  
1148 Towards benchmarking continuous improvement of language agents. *ArXiv*, abs/2406.08747, 2024.  
1149 URL <https://api.semanticscholar.org/CorpusID:270440494>.
- 1151 Rong Wu, Xiaoman Wang, Jianbiao Mei, Pinlong Cai, Daocheng Fu, Cheng Yang, Licheng Wen,  
1152 Xuemeng Yang, Yufan Shen, Yuxin Wang, and Botian Shi. Evolver: Self-evolving llm agents  
1153 through an experience-driven lifecycle. *ArXiv*, abs/2510.16079, 2025a. URL [https://api.  
1154 semanticscholar.org/CorpusID:282210254](https://api.semanticscholar.org/CorpusID:282210254).
- 1155 Shanglin Wu and Kai Shu. Memory in llm-based multi-agent systems: Mechanisms, challenges, and  
1156 collective intelligence. December 2025. doi: 10.36227/techrxiv.176539617.79044553/v1. URL  
1157 <http://dx.doi.org/10.36227/techrxiv.176539617.79044553/v1>.
- 1159 Shanglin Wu et al. Memory in llm-based multi-agent systems: Mechanisms, challenges, and col-  
1160 lective intelligence. *TechRxiv preprint*, 2025b. URL [https://github.com/ShanglinWu/  
1161 LLM-MAS\\_Memory\\_Survey](https://github.com/ShanglinWu/LLM-MAS_Memory_Survey).
- 1162 Yaxiong Wu, Sheng Liang, Chen Zhang, Yichao Wang, Yongyue Zhang, Huifeng Guo, Ruiming  
1163 Tang, and Yong Liu. From human memory to ai memory: A survey on memory mechanisms in the  
1164 era of llms. *ArXiv*, abs/2504.15965, 2025c. URL [https://api.semanticscholar.org/  
1165 CorpusID:277993681](https://api.semanticscholar.org/CorpusID:277993681).
- 1167 Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe  
1168 Wang, Senjie Jin, Enyu Zhou, Rui Zheng, Xiaoran Fan, Xiao Wang, Limao Xiong, Qin Liu,  
1169 Yuhao Zhou, Weiran Wang, Changhao Jiang, Yicheng Zou, Xiangyang Liu, Zhangyue Yin, Shihan  
1170 Dou, Rongxiang Weng, Wensen Cheng, Qi Zhang, Wenjuan Qin, Yongyan Zheng, Xipeng Qiu,  
1171 Xuanjing Huan, and Tao Gui. The rise and potential of large language model based agents: A survey.  
1172 *ArXiv*, abs/2309.07864, 2023. URL [https://api.semanticscholar.org/CorpusID:  
1173 261817592](https://api.semanticscholar.org/CorpusID:261817592).
- 1174 Menglin Xia, Victor Ruele, Saravan Rajmohan, and Reza Shokri. Minerva: A programmable  
1175 memory test benchmark for language models. *ArXiv*, abs/2502.03358, 2025a. URL [https:  
1176 //api.semanticscholar.org/CorpusID:276116638](https://api.semanticscholar.org/CorpusID:276116638).
- 1178 Peng Xia, Peng Xia, Kaide Zeng, Jiaqi Liu, Can Qin, Fang Wu, Yiyang Zhou, Caiming Xiong, and  
1179 Huaxiu Yao. Agent0: Unleashing self-evolving agents from zero data via tool-integrated reasoning.  
1180 2025b. URL <https://api.semanticscholar.org/CorpusID:283110047>.
- 1181 Chaojun Xiao, Penge Zhang, Xu Han, Guangxuan Xiao, Yankai Lin, Zhengyan Zhang, Zhiyuan Liu,  
1182 Song Han, and Maosong Sun. Inllm: Training-free long-context extrapolation for llms with an  
1183 efficient context memory. *Advances in Neural Information Processing Systems 37*, 2024. URL  
1184 <https://api.semanticscholar.org/CorpusID:267523068>.
- 1185 Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. Efficient streaming  
1186 language models with attention sinks. *ArXiv*, abs/2309.17453, 2023. URL [https://api.  
1187 semanticscholar.org/CorpusID:263310483](https://api.semanticscholar.org/CorpusID:263310483).

- 1188 Yunzhong Xiao, Yangmin Li, Hwei Wang, Yunlong Tang, and Zora Zhiruo Wang. Toolmem:  
1189 Enhancing multimodal agents with learnable tool capability memory. *ArXiv*, abs/2510.06664,  
1190 2025. URL <https://api.semanticscholar.org/CorpusID:281891854>.  
1191
- 1192 Zidi Xiong, Yuping Lin, Wenya Xie, Pengfei He, Jiliang Tang, Himabindu Lakkaraju, and Zhen  
1193 Xiang. How memory management impacts llm agents: An empirical study of experience-following  
1194 behavior. *ArXiv*, abs/2505.16067, 2025. URL <https://api.semanticscholar.org/CorpusID:278788765>.  
1195
- 1196 Haoran Xu, Jiacong Hu, Ke Zhang, Lei Yu, Yuxin Tang, Xinyuan Song, Yiqun Duan, Lynn Ai, and  
1197 Bill Shi. Sedm: Scalable self-evolving distributed memory for agents. *ArXiv*, abs/2509.09498,  
1198 2025a. URL <https://api.semanticscholar.org/CorpusID:281252587>.  
1199
- 1200 Wujiang Xu, Zujie Liang, Kai Mei, Hang Gao, Juntao Tan, and Yongfeng Zhang. A-mem:  
1201 Agentic memory for llm agents. *ArXiv*, abs/2502.12110, 2025b. URL <https://api.semanticscholar.org/CorpusID:276421617>.  
1202
- 1203 Siqiao Xue, Caigao Jiang, Wenhui Shi, Fangyin Cheng, Keting Chen, Hongjun Yang, Zhiping  
1204 Zhang, Jianshan He, Hongyang Zhang, Ganglin Wei, Wang Zhao, Fan Zhou, Danrui Qi, Hong Yi,  
1205 Shadong Liu, and Faqiang Chen. Db-gpt: Empowering database interactions with private large  
1206 language models. *ArXiv*, abs/2312.17449, 2023. URL <https://api.semanticscholar.org/CorpusID:266690744>.  
1207
- 1208 Sikuan Yan, Xiufeng Yang, Zuchao Huang, Ercong Nie, Zifeng Ding, Zonggen Li, Xiaowen Ma,  
1209 Hinrich Schutze, Volker Tresp, and Yunpu Ma. Memory-r1: Enhancing large language model  
1210 agents to manage and utilize memories via reinforcement learning. *ArXiv*, abs/2508.19828, 2025a.  
1211 URL <https://api.semanticscholar.org/CorpusID:280918480>.  
1212
- 1213 Xue Yan, Zijing Ou, Mengyue Yang, Yan Song, Haifeng Zhang, Yingzhen Li, and Jun Wang. Memory-  
1214 driven self-improvement for decision making with large language models. *ArXiv*, abs/2509.26340,  
1215 2025b. URL <https://api.semanticscholar.org/CorpusID:281682979>.  
1216
- 1217 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang  
1218 Gao, Chengen Huang, Chenxu Lv, Chuji Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng  
1219 Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang,  
1220 Jian Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu,  
1221 Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men,  
1222 Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren,  
1223 Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang,  
1224 Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu.  
1225 Qwen3 technical report. *CoRR*, abs/2505.09388, 2025a. doi: 10.48550/ARXIV.2505.09388. URL  
1226 <https://doi.org/10.48550/arXiv.2505.09388>.  
1227
- 1228 Ruihan Yang, Jiangjie Chen, Yikai Zhang, Siyu Yuan, Aili Chen, Kyle Richardson, Yanghua Xiao,  
1229 and Deqing Yang. Selfgoal: Your language agents already know how to achieve high-level goals.  
1230 *ArXiv*, abs/2406.04784, 2024. URL <https://api.semanticscholar.org/CorpusID:270357425>.  
1231
- 1232 Wei Yang, Jinwei Xiao, Hongming Zhang, Qingyang Zhang, Yanna Wang, and Bo Xu. Coarse-to-  
1233 fine grounded memory for llm agent planning. *ArXiv*, abs/2508.15305, 2025b. URL <https://api.semanticscholar.org/CorpusID:280699719>.  
1234
- 1235 Yongjin Yang, Sinjae Kang, Juyong Lee, Dongjun Lee, Se young Yun, and Kimin Lee.  
1236 Automated skill discovery for language agents through exploration and iterative feedback.  
1237 *ArXiv*, abs/2506.04287, 2025c. URL <https://api.semanticscholar.org/CorpusID:279243862>.  
1238
- 1239 Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov,  
1240 and Christopher D. Manning. Hotpotqa: A dataset for diverse, explainable multi-hop question  
1241 answering. In *Conference on Empirical Methods in Natural Language Processing*, 2018. URL  
<https://api.semanticscholar.org/CorpusID:52822214>.

- 1242 Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao.  
1243 React: Synergizing reasoning and acting in language models. *ArXiv*, abs/2210.03629, 2022. URL  
1244 <https://api.semanticscholar.org/CorpusID:252762395>.  
1245
- 1246 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik  
1247 Narasimhan. Tree of thoughts: Deliberate problem solving with large language models.  
1248 *ArXiv*, abs/2305.10601, 2023. URL [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:258762525)  
1249 258762525.
- 1250 Rui Ye, Zhongwang Zhang, Kuan Li, Huifeng Yin, Zhengwei Tao, Yida Zhao, Liangcai Su, Li-  
1251 wen Zhang, Zile Qiao, Xinyu Wang, Pengjun Xie, Fei Huang, Siheng Chen, Jingren Zhou,  
1252 and Yong Jiang. Agentfold: Long-horizon web agents with proactive context management.  
1253 *ArXiv*, abs/2510.24699, 2025. URL [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:282400798)  
1254 282400798.
- 1255
- 1256 Howard Yen, Tianyu Gao, Minmin Hou, Ke Ding, Daniel Fleischer, Peter Izsak, Moshe Wasserblat,  
1257 and Danqi Chen. Helmet: How to evaluate long-context language models effectively and thoroughly.  
1258 *ArXiv*, abs/2410.02694, 2024. URL [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:273098808)  
1259 273098808.
- 1260 Zhangyue Yin, Qiushi Sun, Qipeng Guo, Zhiyuan Zeng, Qinyuan Cheng, Xipeng Qiu, and Xuanjing  
1261 Huang. Explicit memory learning with expectation maximization. In *Conference on Empirical*  
1262 *Methods in Natural Language Processing*, 2024. URL [https://api.semanticscholar.](https://api.semanticscholar.org/CorpusID:273901309)  
1263 [org/CorpusID:273901309](https://api.semanticscholar.org/CorpusID:273901309).  
1264
- 1265 Simon Yu, Gang Li, Weiyan Shi, and Pengyuan Qi. Polyskill: Learning generalizable skills  
1266 through polymorphic abstraction. *ArXiv*, abs/2510.15863, 2025a. URL [https://api.](https://api.semanticscholar.org/CorpusID:282203716)  
1267 [semanticscholar.org/CorpusID:282203716](https://api.semanticscholar.org/CorpusID:282203716).
- 1268
- 1269 Wenhao Yu, Zhenwen Liang, Chengsong Huang, Kishan Panaganti, Tianqing Fang, Haitao Mi, and  
1270 Dong Yu. Guided self-evolving llms with minimal human supervision. 2025b. URL [https:](https://api.semanticscholar.org/CorpusID:283458198)  
1271 [//api.semanticscholar.org/CorpusID:283458198](https://api.semanticscholar.org/CorpusID:283458198).
- 1272 Ruifeng Yuan, Shichao Sun, Zili Wang, Ziqiang Cao, and Wenjie Li. Personalized large lan-  
1273 guage model assistant with evolving conditional memory. In *International Conference on Com-*  
1274 *putational Linguistics*, 2023. URL [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:266690798)  
1275 266690798.
- 1276
- 1277 Sizhe Yuen, Francisco Gomez Medina, Ting Su, Yali Du, and Adam J. Sobey. Intrinsic mem-  
1278 ory agents: Heterogeneous multi-agent llm systems through structured contextual memory.  
1279 *ArXiv*, abs/2508.08997, 2025. URL [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:280635601)  
1280 280635601.
- 1281 Yunpeng Zhai, Shuchang Tao, Cheng Chen, Anni Zou, Ziqian Chen, Qingxu Fu, Shinji Mai, Li Yu,  
1282 Jiaji Deng, Zouying Cao, Zhaoyang Liu, Bolin Ding, and Jingren Zhou. Agentevolver: Towards  
1283 efficient self-evolving agent system. 2025. URL [https://api.semanticscholar.org/](https://api.semanticscholar.org/CorpusID:282992095)  
1284 [CorpusID:282992095](https://api.semanticscholar.org/CorpusID:282992095).
- 1285
- 1286 Genghan Zhang, Shaowei Zhu, Anjiang Wei, Zhenyu Song, Allen Nie, Zhen Jia, Nandita Vijaykumar,  
1287 Yida Wang, and Kunle Olukotun. Accelopt: A self-improving llm agentic system for ai accelerator  
1288 kernel optimization. 2025a. URL [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:283109792)  
1289 283109792.
- 1290
- 1291 Gui-Min Zhang, Muxin Fu, Guancheng Wan, Miao Yu, Kun Wang, and Shuicheng Yan. G-memory:  
1292 Tracing hierarchical memory for multi-agent systems. *ArXiv*, abs/2506.07398, 2025b. URL  
1293 <https://api.semanticscholar.org/CorpusID:279250852>.
- 1294
- 1295 Gui-Min Zhang, Muxin Fu, and Shuicheng Yan. Memgen: Weaving generative latent memory for self-  
evolving agents. *ArXiv*, abs/2509.24704, 2025c. URL [https://api.semanticscholar.](https://api.semanticscholar.org/CorpusID:281676243)  
[org/CorpusID:281676243](https://api.semanticscholar.org/CorpusID:281676243).

- 1296 Gui-Min Zhang, Fanci Meng, Guancheng Wan, Zherui Li, Kun Wang, Zhenfei Yin, Lei  
1297 Bai, and Shuicheng Yan. Latentevolve: Self-evolving test-time scaling in latent space.  
1298 *ArXiv*, abs/2509.24771, 2025d. URL [https://api.semanticscholar.org/CorpusID:  
1299 281675514](https://api.semanticscholar.org/CorpusID:281675514).
- 1300 Guibin Zhang, Haotian Ren, Chong Zhan, Zhenhong Zhou, Junhao Wang, He Zhu, Wangchunshu  
1301 Zhou, and Shuicheng Yan. Memevolve: Meta-evolution of agent memory systems. 2025e. URL  
1302 <https://api.semanticscholar.org/CorpusID:284078455>.
- 1303 Kai Zhang, Xiangchao Chen, Bo Liu, Tianci Xue, Zeyi Liao, Zhihan Liu, Xiyao Wang, Yuting  
1304 Ning, Zhaorun Chen, Xiaohan Fu, Jian Xie, Yuxuan Sun, Boyu Gou, Qi Qi, Zihang Meng,  
1305 Jianwei Yang, Ning Zhang, Xian Li, Ashish Shah, Dat Huynh, Hengduo Li, Zi Xian Yang,  
1306 Sara Cao, Lawrence Jang, Shuyan Zhou, Jiacheng Zhu, Huan Sun, Jason Weston, Yu Su, and  
1307 Yifan Wu. Agent learning via early experience. *ArXiv*, abs/2510.08558, 2025f. URL [https:  
1308 //api.semanticscholar.org/CorpusID:281951121](https://api.semanticscholar.org/CorpusID:281951121).
- 1309 Kongcheng Zhang, Qi Yao, Shunyu Liu, Wenjian Zhang, Mingcan Cen, Yang Zhou, Wenkai Fang,  
1310 Yiru Zhao, Baisheng Lai, and Mingli Song. Replay failures as successes: Sample-efficient reinforce-  
1311 ment learning for instruction following. 2025g. URL [https://api.semanticscholar.  
1312 org/CorpusID:284311803](https://api.semanticscholar.org/CorpusID:284311803).
- 1313 Muru Zhang, Ofir Press, William Merrill, Alisa Liu, and Noah A. Smith. How language  
1314 model hallucinations can snowball. *ArXiv*, abs/2305.13534, 2023a. URL [https://api.  
1315 semanticscholar.org/CorpusID:258841857](https://api.semanticscholar.org/CorpusID:258841857).
- 1316 Puzhen Zhang, Xuyang Chen, Yu Feng, Yuhan Jiang, and Liqiu Meng. Constructing coherent  
1317 spatial memory in llm agents through graph rectification. *ArXiv*, abs/2510.04195, 2025h. URL  
1318 <https://api.semanticscholar.org/CorpusID:281843644>.
- 1319 Yuxiang Zhang, Jiangming Shu, Ye Ma, Xueyuan Lin, Shangxi Wu, and Jitao Sang. Memory  
1320 as action: Autonomous context curation for long-horizon agentic tasks, 2025i. URL [https:  
1321 //arxiv.org/abs/2510.12635](https://arxiv.org/abs/2510.12635).
- 1322 Zeyu Zhang, Quanyu Dai, Xiaohe Bo, Chen Ma, Rui Li, Xu Chen, Jieming Zhu, Zhenhua  
1323 Dong, and Ji-Rong Wen. A survey on the memory mechanism of large language model-  
1324 based agents. *ACM Transactions on Information Systems*, 43:1 – 47, 2024. URL [https:  
1325 //api.semanticscholar.org/CorpusID:269293320](https://api.semanticscholar.org/CorpusID:269293320).
- 1326 Zeyu Zhang, Quanyu Dai, Xu Chen, Rui Li, Zhongyang Li, and Zhenhua Dong. Memengine:  
1327 A unified and modular library for developing advanced memory of llm-based agents. *Com-  
1328 panion Proceedings of the ACM on Web Conference 2025*, 2025j. URL [https://api.  
1329 semanticscholar.org/CorpusID:278327095](https://api.semanticscholar.org/CorpusID:278327095).
- 1330 Zeyu Zhang, Yang Zhang, Haoran Tan, Rui Li, and Xu Chen. Explicit v.s. implicit memory: Exploring  
1331 multi-hop complex reasoning over personalized information. *ArXiv*, abs/2508.13250, 2025k. URL  
1332 <https://api.semanticscholar.org/CorpusID:280686168>.
- 1333 Zhenyu (Allen) Zhang, Ying Sheng, Tianyi Zhou, Tianlong Chen, Lianmin Zheng, Ruisi Cai,  
1334 Zhao Song, Yuandong Tian, Christopher Ré, Clark W. Barrett, Zhangyang Wang, and Beidi  
1335 Chen. H2o: Heavy-hitter oracle for efficient generative inference of large language models.  
1336 *ArXiv*, abs/2306.14048, 2023b. URL [https://api.semanticscholar.org/CorpusID:  
1337 259263947](https://api.semanticscholar.org/CorpusID:259263947).
- 1338 Siyan Zhao, Mingyi Hong, Yang Liu, Devamanyu Hazarika, and Kaixiang Lin. Do llms recognize  
1339 your preferences? evaluating personalized preference following in llms. *ArXiv*, abs/2502.09597,  
1340 2025. URL <https://api.semanticscholar.org/CorpusID:276317480>.
- 1341 Junhao Zheng, Xidi Cai, Qiuke Li, Duzhen Zhang, Zhongzhi Li, Yingying Zhang, Le Song, and  
1342 Qianli Ma. Lifelongagentbench: Evaluating llm agents as lifelong learners. *ArXiv*, abs/2505.11942,  
1343 2025. URL <https://api.semanticscholar.org/CorpusID:278739762>.
- 1344 Qinqing Zheng, Mikael Henaff, Amy Zhang, Aditya Grover, and Brandon Amos. Online intrinsic  
1345 rewards for decision making agents from large language model feedback. *ArXiv*, abs/2410.23022,  
1346 2024. URL <https://api.semanticscholar.org/CorpusID:273695775>.
- 1347  
1348  
1349

- 1350 Wanjun Zhong, Lianghong Guo, Qi-Fei Gao, He Ye, and Yanlin Wang. Memorybank: Enhancing  
1351 large language models with long-term memory. *ArXiv*, abs/2305.10250, 2023. URL <https://api.semanticscholar.org/CorpusID:258741194>.  
1352  
1353
- 1354 Wangchunshu Zhou, Yuchen Eleanor Jiang, Peng Cui, Tiannan Wang, Zhenxin Xiao, Yifan Hou, Ryan  
1355 Cotterell, and Mrinmaya Sachan. Recurrentgpt: Interactive generation of (arbitrarily) long text.  
1356 *ArXiv*, abs/2305.13304, 2023. URL <https://api.semanticscholar.org/CorpusID:258832617>.  
1357
- 1358 Yanfang Zhou, Xiaodong Li, Yuntao Liu, Yongqiang Zhao, Xintong Wang, Zhenyu Li, Jinlong  
1359 Tian, and Xinhai Xu. M2pa: A multi-memory planning agent for open worlds inspired by  
1360 cognitive theory. In *Annual Meeting of the Association for Computational Linguistics, 2025a*. URL  
1361 <https://api.semanticscholar.org/CorpusID:280322030>.
- 1362 Zijian Zhou, Ao Qu, Zhaoxuan Wu, Sunghwan Kim, Sunghwan Kim, Alok Prakash, Daniela  
1363 Rus, Jinhua Zhao, Bryan Kian Hsiang Low, and Paul Liang. Mem1: Learning to synergize  
1364 memory and reasoning for efficient long-horizon agents. *ArXiv*, abs/2506.15841, 2025b. URL  
1365 <https://api.semanticscholar.org/CorpusID:279465470>.  
1366
- 1367 Kunlun Zhu, Zijia Liu, Bingxuan Li, Muxin Tian, Yingxuan Yang, Jiaxun Zhang, Pengrui Han,  
1368 Qipeng Xie, Fuyang Cui, Weijia Zhang, Xiaoteng Ma, Xiaodong Yu, Gowtham Ramesh, Jialian  
1369 Wu, Zicheng Liu, Pan Lu, James Zou, and Jiaxuan You. Where llm agents fail and how they can  
1370 learn from failures. *ArXiv*, abs/2509.25370, 2025. URL <https://api.semanticscholar.org/CorpusID:281681143>.  
1371
- 1372 Jiaru Zou, Xiyuan Yang, Ruizhong Qiu, Gaotang Li, Katherine Tieu, Pan Lu, Ke Shen, Hanghang  
1373 Tong, Yejin Choi, Jingrui He, James Zou, Mengdi Wang, and Ling Yang. Latent collaboration  
1374 in multi-agent systems. 2025. URL <https://api.semanticscholar.org/CorpusID:283251109>.  
1375  
1376

## 1377 A OVERVIEW

1378  
1379 We first provide a formal definition for the operational framework of LLM agents and the evolutionary  
1380 paradigms of memory mechanisms. Furthermore, we categorize the primary drivers for this evolution  
1381 into three dimensions, emphasizing that these driving forces constitute the fundamental factors  
1382 supporting transformations in memory mechanisms and the inherent capabilities of LLM agents.  
1383 Based on the depth to which historical trajectories are utilized, this survey proposes an evolutionary  
1384 framework consisting of three distinct stages:

- 1385 • **Storage:** As the foundational layer of evolution, this stage focuses on the faithful preservation  
1386 of trajectories from interactions over a long duration to address constraints regarding the  
1387 memory capacity of LLM agents.
- 1388 • **Reflection:** Through the introduction of loops for dynamic evaluation, the memory mechanism  
1389 transitions from a recorder of information to an evaluator, thereby mitigating issues related to  
1390 hallucinations and logical errors within the memory of LLM agents.
- 1391 • **Experience:** Representing the highest level of cognition, this stage employs abstraction across  
1392 multiple trajectories to extract behavioral patterns of a higher order. This process compresses  
1393 redundant memory into heuristic strategies that are transferable and reusable.  
1394

1395 Furthermore, we provide an in-depth discussion on two pivotal technological shifts required for  
1396 memory mechanisms to advance toward the stage of experience: active exploration and abstraction  
1397 across trajectories. These advancements enable LLM agents to transition from passive recipients of  
1398 information to collectors of experience driven by specific goals, thereby enhancing the capacity for  
1399 proactive generalization in tasks that are unknown. Finally, this survey discusses several valuable  
1400 directions for the future of memory mechanisms.

1401 **Summary of Contribution.** As a survey, our primary objective is the synthesis and analysis  
1402 of existing research, while providing novel insights and perspectives for researchers who aim to  
1403 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.

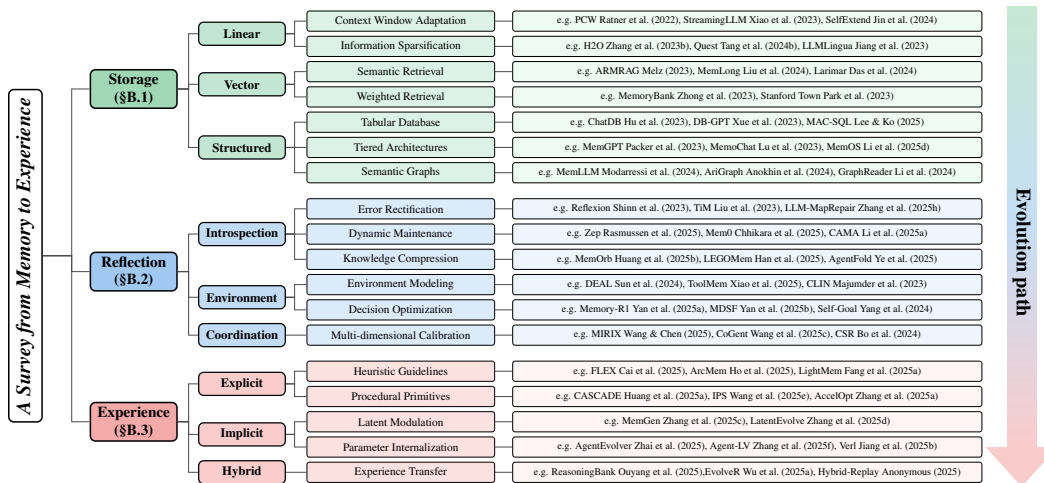


Figure 4: Taxonomy of LLM agent Memory Mechanisms

## 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. (2023b) evicts tokens based on cumulative attention scores, while Tang et al. (2024b) 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 & 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 Zhou et al. (2023); Lu 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. (2024b). 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 Zhou et al. (2023); Lu 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 Zhang et al. (2023a); Ghasemabadi & Niu (2025); Zhang et al. (2025g). 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. (2025h) 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. (2025a) 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 Packer et al. (2023); Zhou et al. (2025b); Kang et al. (2025).
- **Knowledge Compression:** Knowledge compression distills high-dimensional trajectories into compact and reusable representations. Huang et al. (2025b) 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. (2025b), and context-folding techniques preserve working-context efficiency during reasoning Sun et al. (2025c); Ye et al. (2025); Li et al. (2025b).

**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. (2025b), 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. (2025a) learns discrete actions from outcome-based rewards, while Yan et al. (2025b) 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 & Chen (2025) coordinates core, episodic, and semantic memory modules to process multimodal long contexts. Wang et al. (2025c) 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 & 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. (Shinn et al., 2023; Renze & Guven, 2024) 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. (Hong & He, 2025; Zhu et al., 2025; Fu 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. (2025a) 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. (2025b) 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. (2025e) 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. (2025a) 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. (2025a) 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. (2025c) introduces the MemGen framework, employing a "Memory Weaver" to dynamically generate and inject latent token sequences conditioned on current reasoning state. Zhang et al. (2025d) 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. (2025b) 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. (2025f) 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. (2025a) proposes group-based policy optimization for fine-grained experience internalization across multi-turn interactions. Jiang et al. (2025b) 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. (2025a) 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 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.

## D 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. (2024a). 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

1728 the ability of the model for the integration of facts across documents and reasoning in multiple  
1729 steps within the context of personal trivia (Zhang et al., 2025k; Yang et al., 2018). Furthermore,  
1730 certain frameworks focus on emotional support and interactive scenarios to evaluate the model’s  
1731 capacity for memory recall across proactive and passive paradigms He et al. (2024a).

1732  
1733 **Reflection Stage.** The core of the Reflection stage is the evaluation of how agents transform raw  
1734 trajectories into memory of high quality, with an emphasis on the denoising and fidelity of memory,  
1735 the deep alignment with characteristics of users, and the support for perception within complex  
1736 environments.

- 1737 • **Error Correction:** Error correction primarily evaluates whether errors or hallucinations  
1738 emerge during the lifecycle of the memory system. For instance, it involves testing operations  
1739 for the search, editing, and matching of memory Xia et al. (2025a), examining the presence of  
1740 hallucinations during the stages of extraction or update Chen et al. (2025a).
- 1741 • **Personalization:** Personalization focuses on the capacity for the extraction of deep personal-  
1742 ization from the history of the agent, which includes the mining of latent information through  
1743 reflection to identify implicit preferences (Jiang et al., 2025a; Huang et al., 2025c), traits of  
1744 users (Du et al., 2024; Zhao et al., 2025), key information (Yuan et al., 2023; Li et al., 2025c),  
1745 and shared components (Tsaknakis et al., 2025; Kim et al., 2024a).
- 1746 • **Dynamic Reasoning:** Dynamic reasoning emphasizes the critical role of memory in reasoning  
1747 across multiple steps and the perception of environments with high complexity. This involves  
1748 the selective forgetting of memory Hu et al. (2025a), backtracking on decisions Wan & Ma  
1749 (2025), scenarios in the real world (Deng et al., 2024; Miyai et al., 2025), and the mechanisms  
1750 for summarization and transition Maharana et al. (2024).

1751  
1752 **Experience Stage.** The Experience Stage represents the pinnacle of the evolutionary path of memory  
1753 mechanisms; at this phase, the focus shifts toward how LLM agents abstract general experience  
1754 from fragmented trajectories within dynamic environments to facilitate continuous evolution through  
1755 practical application. While benchmarks for this stage are relatively scarce, they possess a strong  
1756 empirical character: Wu et al. (2024) and Ai et al. (2025) simulate environments for authentic  
1757 deployment to evaluate the capacity of LLM agents for the extraction and internalization of experience  
1758 within cycles of input and feedback; conversely, Wei et al. (2025a) and Zheng et al. (2025) emphasize  
1759 the capacity for the transfer of experience, measuring levels of abstraction and generalization by  
1760 assessing the transfer of acquired experience to a diverse range of other tasks.

1761  
1762  
1763  
1764  
1765  
1766  
1767  
1768  
1769  
1770  
1771  
1772  
1773  
1774  
1775  
1776  
1777  
1778  
1779  
1780  
1781

1782  
 1783  
 1784  
 1785  
 1786  
 1787  
 1788  
 1789  
 1790  
 1791  
 1792  
 1793  
 1794  
 1795  
 1796  
 1797  
 1798  
 1799  
 1800  
 1801  
 1802  
 1803  
 1804  
 1805  
 1806  
 1807  
 1808  
 1809  
 1810  
 1811  
 1812  
 1813  
 1814  
 1815  
 1816  
 1817  
 1818  
 1819  
 1820  
 1821  
 1822  
 1823  
 1824  
 1825  
 1826  
 1827  
 1828  
 1829  
 1830  
 1831  
 1832  
 1833  
 1834  
 1835

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. (2024a)	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	Kuratov 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. (2024b)	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. (2025k)	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. (2024a)	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.

Stage	Dataset	Reference	Size	Description
Reflection Stage Benchmark	Minerva	Xia et al. (2025a)	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. (2025a)	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. (2025a)	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. (2025a)	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 & 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. (2025c)	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. (2024a)	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. (2025c)	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.
	Experience Stage Benchmark	LIDB	Tsaknakis et al. (2025)	210
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. (2025a)	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.	

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