

# AMA: Adaptive Memory via Multi-Agent Collaboration

Weiquan Huang<sup>1\*</sup>, Zixuan Wang<sup>1\*</sup>, Hehai Lin<sup>1</sup>, Sudong Wang<sup>1</sup>, Bo Xu<sup>1</sup>,  
Qian Li<sup>2</sup>, Beier Zhu<sup>3</sup>, Linyi Yang<sup>4</sup>, Chengwei Qin<sup>1†</sup>

<sup>1</sup>The Hong Kong University of Science and Technology (Guangzhou)

<sup>2</sup>Shandong University <sup>3</sup>Nanyang Technological University

<sup>4</sup>Southern University of Science and Technology

## Abstract

The rapid evolution of Large Language Model (LLM) agents has necessitated robust memory systems to support cohesive long-term interaction and complex reasoning. Benefiting from the strong capabilities of LLMs, recent research focus has shifted from simple context extension to the development of dedicated agentic memory systems. However, existing approaches typically rely on rigid retrieval granularity, accumulation-heavy maintenance strategies, and coarse-grained update mechanisms. These design choices create a persistent mismatch between stored information and task-specific reasoning demands, while leading to the unchecked accumulation of logical inconsistencies over time. To address these challenges, we propose **Adaptive Memory via Multi-Agent Collaboration (AMA)**, a novel framework that leverages coordinated agents to manage memory across multiple granularities. AMA employs a hierarchical memory design that dynamically aligns retrieval granularity with task complexity. Specifically, the **Constructor** and **Retriever** jointly enable multi-granularity memory construction and adaptive query routing. The **Judge** verifies the relevance and consistency of retrieved content, triggering iterative retrieval when evidence is insufficient or invoking the **Refresher** upon detecting logical conflicts. The **Refresher** then enforces memory consistency by performing targeted updates or removing outdated entries. Extensive experiments on challenging long-context benchmarks show that AMA significantly outperforms state-of-the-art baselines while reducing token consumption by approximately 80% compared to full-context methods, demonstrating its effectiveness in maintaining retrieval precision and long-term memory consistency.

\*Equal contribution.

†Corresponding author.

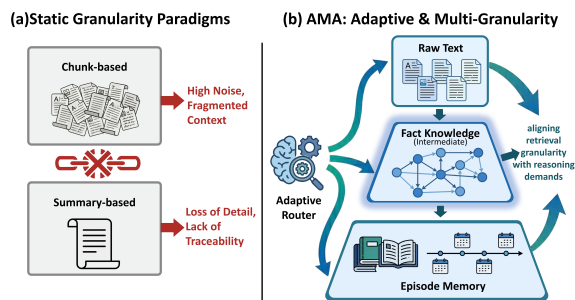


Figure 1: **Comparison of static paradigms and the AMA framework.** (a) Static methods suffer from the dilemma of fixed granularity, leading to either noise or information loss. (b) AMA dynamically determines the memory granularity to use, aligning retrieval precision with reasoning demands.

## 1 Introduction

Large Language Model (LLM) agents have demonstrated strong capabilities in complex reasoning, tool use, and multi-turn interaction scenarios (Deng et al., 2023; Liang and Tong, 2025; Comanici et al., 2025; Mei et al., 2024). Supporting such behaviors requires long-term memory to preserve contextual coherence and consistency (Liu et al., 2023; Sumers et al., 2023). Existing approaches to long-term memory can be broadly categorized into internal and external memory paradigms (Zhang et al., 2025). Internal memory implicitly absorbs historical information into model parameters, but is constrained by limited capacity (Mallen et al., 2023) and incurs substantial costs for continual updates (Wang et al., 2024b; Thede et al., 2025). In contrast, external memory relies on explicit storage and retrieval, providing superior scalability and editability (Wang and Chen, 2025; Qian et al., 2025; Rezazadeh et al., 2024). As a result, it has become the dominant approach, making the design of efficient and reliable external memory systems a critical foundation for sustained agent evolution.

Building on the growing adoption of external memory, many systems support dynamic memory

management through explicit Create-Read-Update-Delete operations, enabling agents to incrementally maintain memory over time (Zhong et al., 2024; Wang et al., 2024c; Yan et al., 2025; Rasmussen et al., 2025). Despite these advantages, they exhibit a fundamental limitation: a mismatch between the granularity at which memories are stored and the granularity required for effective retrieval and reasoning. As illustrated in Figure 1, these approaches typically rely on static text chunking with fixed lengths or coarse-grained summaries (Zhang et al., 2025; Wu et al., 2025). Such static strategies often disrupt the inherent semantic coherence of stored information, which in turn leads to suboptimal retrieval behavior: overly coarse retrieval introduces substantial irrelevant noise, while excessively fine-grained or isolated chunks fragment essential logical dependencies, ultimately leading to reasoning failures in complex tasks (Hu et al., 2025; Lee et al., 2025; Wang et al., 2024a). These limitations highlight the necessity of an adaptive memory paradigm capable of dynamically aligning memory granularity with task-specific requirements.

To address these challenges, recent work has shifted toward agentic memory mechanisms (Xu et al., 2025; Wang et al., 2025; Wang and Chen, 2025), leveraging the generative capabilities of LLMs to mitigate the rigidity of static storage granularity. Typically, these frameworks employ LLMs to synthesize interaction history into flexible representations like summaries or vector entries, extending the effective context window. While these designs improve representation flexibility, they leave two fundamental challenges largely unaddressed (Packer et al., 2023b; Chhikara et al., 2025). First, the absence of an explicit adaptive routing mechanism prevents agents from selecting the appropriate memory granularity at inference time, leading to persistent mismatches with task demands. Second, reliance on accumulation-heavy strategies and coarse-grained update mechanisms fails to support precise modifications, resulting in the unchecked accumulation of redundancy and errors (Wu et al., 2024; Hu et al., 2025).

To overcome the coupled challenges of adaptive retrieval control and long-term memory evolution, we propose **Adaptive Memory via Multi-Agent Collaboration (AMA)**, as illustrated in Figures 1 and 2. Unlike prior agentic memory systems that mainly rely on a monolithic controller, AMA adopts a multi-agent design that decomposes the memory lifecycle into four functionally distinct yet

interdependent roles: the **Constructor**, **Retriever**, **Judge**, and **Refresher**. Specifically, the **Constructor** transforms unstructured dialogue streams into hierarchical granularities, including Raw Text, Fact Knowledge, and Episode Memory, to accommodate diverse storage requirements. The **Retriever** acts as an adaptive gateway, dynamically routing queries to the most appropriate memory form based on current reasoning demands. To ensure consistency, the **Judge** serves as a logic auditor, verifying relevance to trigger feedback loops and detecting conflicts to activate the **Refresher** for updates. This separation of responsibilities enables fine-grained control over retrieval, verification, and memory evolution, which would be difficult to achieve within a single-agent design without entangling conflicting objectives. Extensive experiments across multiple long-term memory benchmarks demonstrate that AMA consistently outperforms strong memory baselines. By adaptively controlling retrieval granularity and explicitly maintaining memory consistency over time, AMA achieves state-of-the-art performance while reducing token consumption by up to 80% compared to using full context. Moreover, our analysis highlights the importance of the logic-driven Refresher, which plays a critical role in dynamic knowledge maintenance and enables AMA to achieve nearly 90% accuracy in knowledge update scenarios.

In summary, our main contributions are three-fold: (1) We introduce a comprehensive memory paradigm featuring multi-granularity storage and adaptive routing, which incorporates logic-driven conflict detection to maintain long-term consistency and reasoning fidelity. (2) We design a unified multi-agent framework to orchestrate storage, retrieval, and maintenance, facilitating robust memory evolution in long-context applications. (3) Through extensive experiments and analysis, we demonstrate that AMA significantly outperforms state-of-the-art baselines, verifying its effectiveness and robustness in complex long-context tasks.

## 2 Related Work

### 2.1 Memory for LLM Agents

Prior research on memory for LLM agents has investigated a wide range of approaches, ranging from full interaction storage to system-level frameworks (Zhong et al., 2024; Wang et al., 2023; Mei et al., 2024; Liu et al., 2024). These methods typically evolve from context extension to structured

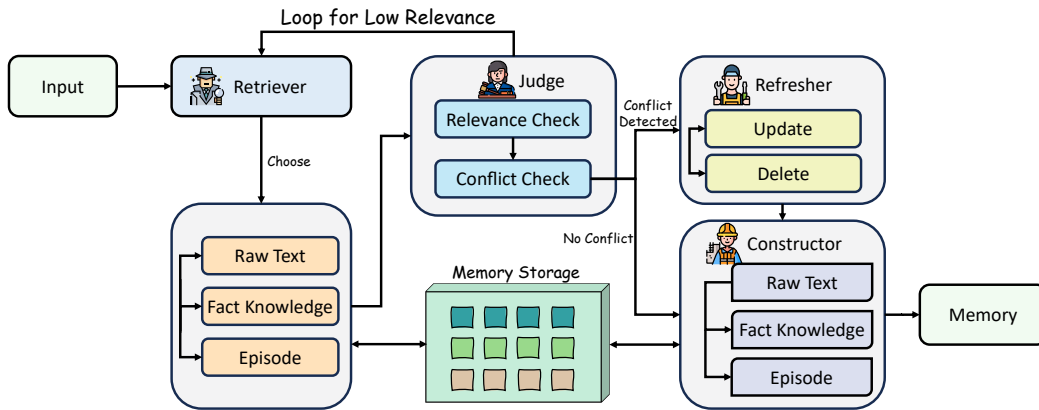


Figure 2: Overview of the AMA framework. The system orchestrates four agents to enable adaptive memory evolution. The **Retriever** routes inputs to optimal granularities based on intent. The **Judge** audits content relevance to trigger feedback loops and detects conflicts. The **Refresher** executes updates or deletions to rectify these inconsistencies. Finally, the **Constructor** synthesizes the validated context into structured memory entries.

organization. Specifically, MemGPT (Packer et al., 2023a) focuses on context management, adopting a cache-like organization to prioritize salient information. Moving towards modularity, Mem0 (Chhikara et al., 2025) abstracts memory as an independent layer dedicated to long-term management. To further enhance retrieval precision, Nemori (Nan et al., 2025) and Zep (Rasmussen et al., 2025) introduce semantic structures, leveraging self-organizing events and temporal knowledge graphs, respectively. Despite their progress, these methods rely on static retrieval strategies, which limits their ability to adaptively coordinate information across different abstraction levels and task stages. Therefore, designing an adaptive memory system that can robustly support long-term interactions remains a critical challenge.

## 2.2 Multi-Agent System

Multi-agent systems have demonstrated clear advantages in tackling complex tasks by enabling role-based collaboration and interactive decision making (Lin et al., 2025; Haji et al., 2024; Abbasnejad et al., 2025; Huot et al., 2024). In software engineering, multi-agent approaches improve system reliability through explicit role specialization and structured workflows (Qian et al., 2024; Hong et al., 2023). In mathematical reasoning, multi-agent frameworks enhance solution accuracy via collaborative interaction and process-level verification (Zhang and Xiong, 2025; Wu et al., 2023). In parallel, a growing body of work on agentic memory focuses on improving long-term information modeling for LLM agents (Xu et al., 2025; Yan et al., 2025). While this line of research pro-

vides valuable insights into memory abstraction and maintenance, most existing approaches are built around a monolithic controller and do not explicitly leverage multi-agent collaboration. A notable recent exception is MIRIX<sup>1</sup> (Wang and Chen, 2025), which explores assigning specialized agents for memory organization, but lacks dedicated mechanisms for long-term memory consistency. Building on these complementary lines of research, our work integrates multi-agent collaboration with agentic memory design to support long-term memory for LLM agents.

## 3 Method

We introduce Adaptive Memory via Multi-Agent Collaboration (AMA) to address the critical challenge of aligning retrieval granularity with diverse task requirements, as well as the unchecked accumulation of logical inconsistencies. As illustrated in Figure 2, the framework operates through a coordinated multi-agent pipeline. The process begins with the **Retriever**, which accesses memory across multiple granularities based on the input intent. The **Judge** then evaluates the relevance of the retrieved content and identifies potential conflicts, triggering feedback retrieval or activating the **Refresher** to perform targeted memory updates when necessary. Finally, the **Constructor** consolidates the validated information and organizes it into memory representations at different granularities, supporting continual memory evolution. In the following sections, we present the detailed design

<sup>1</sup>We did not include MIRIX as a baseline in this work because its official implementation was not publicly available during our experimental phase.

of the Constructor, Retriever, Judge, and Refresher.

### 3.1 Constructor

To clearly delineate the functional roles of different memory granularities within the overall pipeline, we begin by introducing the Constructor. Its primary responsibility is to construct multi-granular memory by generating structured semantic components from the current input  $u_t$ , context window  $W_t$ , and conflict-free memory history  $\mathcal{H}_t^*$ , conditioned on a carefully designed prompt  $P_{con}$ . Drawing inspiration from prior work (Tan et al., 2025) and established linguistic theory (Huddleston and Pullum, 2005), the Constructor decomposes natural language into stable and parsable fact templates. Specifically, it leverages five fundamental sentence patterns defined by combinations of Subject (S), Verb (V), Object (O), and Complement (C): S-V, S-V-O, S-V-C, S-V-O-O, and S-V-O-C. Through this decomposition, the Constructor simultaneously extracts a set of facts and the indices of conversation turns relevant to the current input:  $K_t, R_t \leftarrow \text{Constructor}(u_t, W_t \parallel \mathcal{H}_t^* \parallel P_{con})$ .

The set  $K_t = \{k_{t,1}, k_{t,2}, \dots\}$  represents the structured fact knowledge parsed from the current content. We index dialogue contents by a unique identifier  $D_{s:j}$ , which denotes the  $j$ -th turn in the  $s$ -th session. Based on this indexing scheme, the Constructor automatically selects a subset of relevant historical turns  $R_t \subseteq \{D_{s:j}\}$ . In parallel, the Constructor constructs unified meta-information  $\Omega_t = \{\tau_t, d_t, \text{speaker}_t\}$  for the current turn  $t$ , where  $d_t = D_{s:t}$ . The timestamp  $\tau_t$  encodes precise temporal information, and  $\text{speaker}_t \in \{\text{user}, \text{assistant}\}$ . If the input contains an explicit temporal expression (e.g., dates or event times), it is directly extracted as  $\tau_t$ ; otherwise, the current system time is assigned. This design ensures chronological consistency across multi-turn memories and facilitates time-sensitive conflict detection. Given the tuple  $(u_t, K_t, R_t, \Omega_t)$ , the Constructor then generates memory entries at varying granularities (Figure 3).

**Raw Text Memory.** This component records the content of the current turn in its original form  $u_t$ , together with the reference information ( $R_t$  and  $\Omega_t$ ) generated by the Constructor. Formally, we define  $m_t^{\text{raw}} = \{u_t, R_t, \Omega_t\}$ . This granularity preserves the fundamental conversational trajectory, ensuring both data traceability and retrieval flexibility.

**Fact Knowledge Memory.** Each extracted fact is treated as an independent memory unit. Accordingly, we define Fact Knowledge Memory as

$m_{t,i}^{\text{fact}} = \{k_{t,i}, R_t, \Omega_t\}$  with  $k_{t,i} \in K_t$ . By transforming unstructured text into structured knowledge units, Fact Knowledge Memory enables associative retrieval, facilitates conflict detection, and supports the long-term accumulation and refinement of knowledge within the AMA framework.

**Episode Memory.** It is designed to capture high-level abstractions across turns. Following a gate-keeping mechanism inspired by prior work (Park et al., 2023; Nan et al., 2025), we introduce a trigger function with prompt  $P_{tri}$  to determine a binary activation state  $T_t \in \{0, 1\}$ . The trigger is activated under three conditions: detection of a topic shift, an explicit user request, or saturation of the context window threshold. This generation process is formalized as  $T_t = \text{Constructor}(u_t, W_t \parallel P_{tri})$ .

When activated ( $T_t = 1$ ), the Constructor employs a dedicated prompt  $P_{epi}$  to synthesize an abstract summary  $E_t$ , which directly constitutes the episodic memory entry:  $m_t^{\text{epi}} = E_t = \text{Constructor}(u_t, W_t \parallel P_{epi})$ .

**Memory Encoding.** To support efficient retrieval across memory granularities, we compute a dense vector representation for each memory entry based on its primary semantic content. For a memory entry  $m_i$ , we extract its core text  $c_i$  (i.e., the raw utterance  $u_t$ , the fact  $k_{t,i}$ , or the episode summary  $E_t$ ) and encode it into a high-dimensional embedding using a text encoder (Reimers and Gurevych, 2019):  $e_i = f_{\text{enc}}(c_i)$ . These embeddings serve as keys for the granularity-specific retrieval mechanisms, which are detailed in Section 3.2.

### 3.2 Retriever

The Retriever functions as the memory access gateway within the AMA framework. Its primary role is to dynamically route queries to the most appropriate memory granularity. To address referential ambiguity and missing context commonly observed in raw dialogue, the Retriever first rewrites the query into a self-contained form and then performs adaptive retrieval based on multi-dimensional intent analysis.

**Query Rewriting and Intent Routing.** Given the current input  $u_t$  and the context window  $W_t$ , the Retriever employs a dedicated prompt  $P_{ret}$  to guide the LLM in simultaneously generating three outputs: a context-independent rewritten query  $u'_t$ , a four-dimensional binary intent vector  $\mathbf{B}$ , and a dynamic retrieval count  $K_{dyn}$ :  $u'_t, \mathbf{B}, K_{dyn} \leftarrow \text{Retriever}(u_t, W_t \parallel P_{ret})$ .

The rewritten query  $u'_t$  resolves ambiguous refer-

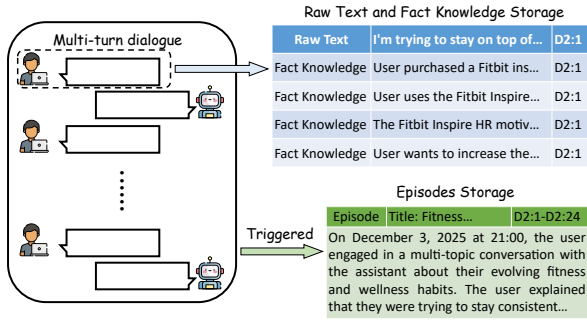


Figure 3: Memory Construction Stage. In this stage Constructor generates raw text and fact knowledge memories from utterances, while conditionally synthesizing abstract episodes upon trigger activation.

ences and omissions, which is suitable for retrieval. The intent vector  $\mathbf{B} = [b_{\text{fine}}, b_{\text{abs}}, b_{\text{event}}, b_{\text{atomic}}]$  encodes the activation of four query dimensions: *fine-grained details*, *abstract summaries*, *cross-temporal events*, and *atomic facts*. Based on  $\mathbf{B}$ , a routing function  $f_M$  dynamically selects the appropriate retrieval operator  $O$ . The mapping  $O = f_M(\mathbf{B})$  is determined by priority:  $O = O_{\text{raw}}$  (if  $b_{\text{fine}} = 1$ );  $O_{\text{epi}}$  (if  $b_{\text{abs}} \vee b_{\text{event}} = 1$ ); else  $O_{\text{fact}}$ .

This routing strategy explicitly prioritizes specialized retrieval intents. When fine-grained detail is required ( $b_{\text{fine}} = 1$ ), the Retriever chooses Raw Text Memory for precise phrasing. When abstract or event-level understanding is needed ( $b_{\text{abs}} = 1$  or  $b_{\text{event}} = 1$ ), Episodic Memory is queried to obtain high-level semantic representations. In all other cases, the Retriever defaults to Fact Knowledge Memory to access structured information.

**Similarity-based Retrieval.** Once the target memory repository  $M$  is determined, we compute the cosine similarity between the embedding of the rewritten query  $u'_t$  and the pre-computed embedding  $e_i$  of each memory entry  $m_i$  within the selected repository  $M$ , defined as  $s_i = \cos(f_{\text{enc}}(u'_t), e_i) = \frac{f_{\text{enc}}(u'_t) \cdot e_i}{\|f_{\text{enc}}(u'_t)\| \|e_i\|}$ . Memory entries are ranked by their similarity scores, with the Top- $K$  entries forming the final retrieval set. To prevent the predicted  $K_{\text{dyn}}$  from being too small to capture sufficient information, we enforce a minimum threshold  $K_m$  and set the effective cutoff as  $K = \max(K_{\text{dyn}}, K_m)$ , where  $K_{\text{dyn}}$  is dynamically predicted by the Retriever. Accordingly, the candidate memory set  $\mathcal{H}_t$  is obtained as  $\mathcal{H}_t = \text{Top}_K(\{m_i\}_{i=1}^{|M|}, \text{key} = s_i)$ . The resulting set  $\mathcal{H}_t$  is then passed to the Judge for verification.

### 3.3 Judge

While the Retriever recalls a candidate memory set  $\mathcal{H}_t$  based on vector similarity, directly injecting unverified memories may introduce noise or amplify hallucinations. To ensure robustness and reliability, the Judge acts as a dynamic filter, performing a sequential dual-verification to refine  $\mathcal{H}_t$  into a validated set  $\mathcal{H}_t^*$ , guided by the prompt  $P_{\text{jud}}$ .

**Relevance Assessment.** The Judge first evaluates the pragmatic utility of the retrieved content with respect to the current input  $u_t$ . To optimize the utilization of  $\mathcal{H}_t$ , we incorporate a relevance-based rejection mechanism. If the density of valid information in  $\mathcal{H}_t$  falls below a predefined threshold, the system triggers a **Retry** action. This issues a feedback signal to the Retriever, prompting it to traverse remaining memory granularities or perform retrieval expansion using relation indices  $R_t$  to broaden the scope. This feedback loop is bounded by the retrieval round limit  $K_r$ . Upon successfully passing this check, the retained relevant memory set is denoted as  $\mathcal{H}_r$ , which serves as the input for the conflict detection phase.

**Conflict Detection.** Subsequently, the Judge conducts logical consistency checks to identify contradictions between the current input  $u_t$  and the filtered memory  $\mathcal{H}_r$ . Typical conflicts include outdated facts that contradict updated user status. When detecting such inconsistencies, the Judge isolates a conflict set  $C_{\text{err}}$ , which comprises the specific memory entries identified as contradictory, and triggers a **Refresh** action to activate the Refresher for targeted updates. In the absence of conflicts, the filtered memory  $\mathcal{H}_r$  is directly instantiated as the validated memory set  $\mathcal{H}_t^*$ , ready for downstream utilization.

The overall verification process is formalized as:  $\mathcal{H}_t^*, C_{\text{err}}, \text{Action} \leftarrow \text{Judge}(u_t, \mathcal{H}_t, W_t \parallel P_{\text{jud}})$ , where  $\text{Action} \in \{\text{Pass}, \text{Retry}, \text{Refresh}\}$  dictates the system flow based on the verification outcome: (1) Action = Retry is triggered when retrieval relevance is insufficient; (2) Action = Refresh is triggered when a non-empty conflict set  $C_{\text{err}}$  is detected, detailed in the section 3.4; (3) Action = Pass occurs when memories are both relevant and consistent, forwarding the validated set  $\mathcal{H}_t^*$  to the Constructor for memory synthesis and to the downstream agent for generating the final response.

Method	Single-Hop			Multi-Hop			Temporal Reasoning			Open Domain			Overall			
	LLM Score	F1	BLEU-1	LLM Score	F1	BLEU-1	LLM Score	F1	BLEU-1	LLM Score	F1	BLEU-1	LLM Score	F1	BLEU-1	
GPT-4o-mini	FullContext	0.823	0.526	0.443	0.663	0.351	0.259	0.557	0.437	0.358	<b>0.482</b>	0.243	0.171	0.717	0.458	0.375
	RAG	0.318	0.220	0.185	0.311	0.184	0.116	0.235	0.194	0.156	0.323	0.189	0.134	0.300	0.206	0.163
	LangMem	0.616	0.389	0.332	0.526	0.336	0.240	0.250	0.321	0.263	0.477	<b>0.295</b>	<b>0.236</b>	0.515	0.359	0.295
	MemGPT	0.479	0.249	0.186	0.229	0.112	0.111	0.388	0.326	0.261	0.500	0.308	0.255	0.415	0.244	0.192
	Zep	0.627	0.394	0.334	0.501	0.273	0.192	0.584	0.444	0.378	0.393	0.227	0.156	0.580	0.372	0.307
	A-Mem	0.518	0.270	0.201	0.248	0.121	0.120	0.546	0.459	0.367	0.541	0.333	0.276	0.476	0.286	0.225
	Mem0	0.676	0.441	0.340	0.598	0.340	0.250	0.500	0.440	0.373	0.403	0.269	0.192	0.608	0.412	0.339
	Nemori	0.817	0.541	0.430	0.650	0.363	0.255	0.707	0.564	<b>0.463</b>	0.446	0.207	0.150	0.740	0.492	0.383
AMA	<b>0.849</b>	<b>0.622</b>	<b>0.548</b>	<b>0.681</b>	<b>0.423</b>	<b>0.303</b>	<b>0.748</b>	<b>0.589</b>	0.461	0.479	0.283	0.228	<b>0.774</b>	<b>0.558</b>	<b>0.465</b>	
GPT-4.1-mini	FullContext	0.848	0.599	0.521	0.753	0.431	0.329	0.723	0.463	0.390	0.552	0.277	0.217	0.786	0.519	0.439
	RAG	0.351	0.252	0.215	0.309	0.196	0.125	0.267	0.217	0.186	0.281	0.174	0.136	0.321	0.229	0.187
	LangMem	0.824	0.498	0.425	0.693	0.405	0.317	0.496	0.473	0.399	<b>0.575</b>	0.320	0.257	0.717	0.464	0.390
	MemGPT	0.490	0.255	0.190	0.234	0.115	0.114	0.397	0.334	0.267	0.512	0.315	0.261	0.425	0.250	0.196
	Zep	0.652	0.444	0.390	0.523	0.297	0.199	0.587	0.233	0.195	0.427	0.236	0.188	0.601	0.360	0.301
	A-Mem	0.530	0.276	0.206	0.254	0.124	0.123	0.559	0.470	0.375	0.553	0.341	0.282	0.487	0.293	0.230
	Mem0	0.697	0.474	0.410	0.666	0.391	0.295	0.555	0.382	0.323	0.468	0.232	0.173	0.647	0.424	0.356
	Nemori	0.828	0.573	0.502	0.732	0.407	0.311	0.757	0.562	<b>0.490</b>	0.498	0.252	0.188	0.774	0.521	0.445
AMA	<b>0.888</b>	<b>0.636</b>	<b>0.570</b>	<b>0.716</b>	<b>0.453</b>	<b>0.324</b>	<b>0.764</b>	<b>0.608</b>	0.484	0.514	<b>0.333</b>	<b>0.251</b>	<b>0.805</b>	<b>0.580</b>	<b>0.492</b>	
Qwen3-30B-Instruct	FullContext	0.845	0.528	0.453	0.665	0.369	0.273	0.555	0.443	0.369	<b>0.551</b>	0.273	0.185	0.733	0.466	0.390
	RAG	0.327	0.221	0.188	0.312	0.193	0.122	0.234	0.197	0.160	0.368	0.219	0.146	0.307	0.209	0.169
	LangMem	0.633	0.390	0.337	0.529	0.353	0.254	0.253	0.325	0.270	0.545	<b>0.341</b>	<b>0.249</b>	0.525	0.365	0.309
	A-Mem	0.518	0.270	0.201	0.248	0.121	0.120	0.546	0.459	0.367	0.541	0.333	0.276	0.476	0.286	0.225
	Nemori	0.839	0.543	0.439	0.652	0.381	0.269	0.704	0.571	<b>0.477</b>	0.510	0.239	0.163	0.756	0.500	0.399
	AMA	<b>0.872</b>	<b>0.625</b>	<b>0.560</b>	<b>0.703</b>	<b>0.445</b>	<b>0.318</b>	<b>0.751</b>	<b>0.597</b>	0.475	0.505	0.327	0.247	<b>0.791</b>	<b>0.570</b>	<b>0.483</b>
Qwen3-8B-Instruct	FullContext	<b>0.803</b>	0.502	0.430	<b>0.632</b>	0.351	0.259	0.527	0.421	0.351	0.523	0.259	0.176	0.696	0.443	0.371
	RAG	0.317	0.214	0.182	0.303	0.187	0.118	0.227	0.191	0.155	0.357	0.212	0.142	0.298	0.203	0.164
	LangMem	0.617	0.380	0.328	0.515	0.344	0.247	0.246	0.317	0.263	<b>0.531</b>	<b>0.332</b>	0.243	0.512	0.356	0.301
	A-Mem	0.504	0.263	0.196	0.241	0.118	0.117	0.532	0.447	0.357	0.527	0.324	<b>0.269</b>	0.464	0.279	0.219
	Nemori	0.762	0.493	0.399	0.592	0.346	0.244	0.639	0.519	<b>0.433</b>	0.463	0.217	0.148	0.686	0.454	0.362
	AMA	0.780	<b>0.559</b>	<b>0.500</b>	0.628	<b>0.398</b>	<b>0.284</b>	<b>0.671</b>	<b>0.533</b>	0.424	0.482	0.292	0.221	<b>0.707</b>	<b>0.510</b>	<b>0.432</b>

Table 1: Main results on the LoCoMo benchmark. We compare AMA with representative memory-based baselines across four backbone models.

Model	Question Type	Full-ctx	RAG	LangMem	A-Mem	MemGPT	Mem0	Zep	Nemori	AMA
GPT-4o-mini	single-session-pref.	0.300	0.333	0.267	0.367	0.200	0.333	<b>0.533</b>	0.467	0.467
	single-session-asst.	0.818	0.714	0.777	0.804	0.857	0.875	0.750	0.839	<b>0.964</b>
	temporal-reasoning	0.365	0.280	0.353	0.398	0.451	0.399	0.541	<b>0.617</b>	0.444
	multi-session	0.406	0.254	0.424	0.451	0.549	0.481	0.474	0.511	<b>0.624</b>
	knowledge-update	0.769	0.385	0.578	0.602	0.410	0.654	0.744	0.615	<b>0.897</b>
	single-session-user	0.814	0.686	0.750	0.750	0.814	0.857	0.929	0.886	<b>0.986</b>
	<b>Average</b>	0.548	0.374	0.528	0.538	0.554	0.574	0.632	0.642	<b>0.698</b>

Table 2: Performance breakdown on LongMemEval<sub>s</sub>. We report category-wise results of AMA and memory-based baselines across six question types.

### 3.4 Refresher

Drawing inspiration from prior studies on dynamic memory maintenance (Wang et al., 2025; Zhong et al., 2024; Yan et al., 2025), we introduce the Refresher to ensure the logical validity of memory storage. This component is triggered exclusively when the Judge detects a conflict set  $C_{err}$ . Guided by a dedicated prompt  $P_{ref}$ , the Refresher follows a strict conditional branching strategy to resolve detected inconsistencies.

**Delete.** This operation is triggered only under two rigorous conditions: (1) in response to explicit user instructions to forget specific information, and (2) when the lifespan of a conflicting memory entry exceeds a predefined *maximum retention limit*. In such cases, the system permanently removes the entry to purge the storage space.

**Update.** For all remaining conflict scenarios, the Refresher defaults to an update operation. Specifically, it performs a state modification  $m_i \leftarrow$

$\mathcal{U}(m_i, u_t)$ , which selectively adjusts the attributes of  $m_i$  to align with the latest state implied by the current input  $u_t$  (e.g., updating outdated location data), thereby rectifying logical contradictions while preserving memory continuity.

The process yields a consistent memory state  $\mathcal{H}_t^* \leftarrow \text{Refresher}(C_{err}, \mathcal{H}_t \parallel P_{ref})$ . This conflict-free set  $\mathcal{H}_t^*$  is then immediately routed to the Constructor for memory synthesis and the downstream agent to ensure reliable response generation.

## 4 Experiment

### 4.1 Experimental Setup

**Datasets and Metrics.** We evaluate long-term memory capabilities on two established benchmarks: **LoCoMo** (Maharana et al., 2024) and **LongMemEval<sub>s</sub>** (Wu et al., 2024). Detailed statistics for both datasets are provided in Appendix A.1. For LoCoMo, we report F1 and BLEU-1 scores in

Method				Locomo			LongMemEval <sub>s</sub>	
RT	FK	EP	RF	LLM Score	F1	BLEU-1	Knowledge-update	Average
✓	×	×	✓	0.669	0.484	0.401	0.767	0.614
×	✓	×	✓	0.712	0.502	0.424	0.804	0.642
×	×	✓	✓	0.688	0.495	0.407	0.748	0.630
✓	✓	×	✓	0.752	0.538	0.448	0.863	0.686
✓	×	✓	✓	0.725	0.515	0.428	0.804	0.656
×	✓	✓	✓	0.741	0.529	0.435	0.842	0.680
✓	✓	✓	✓	<b>0.774</b>	<b>0.558</b>	<b>0.465</b>	<b>0.897</b>	<b>0.710</b>
✓	✓	✓	×	0.771	0.556	0.462	0.568	0.634

Table 3: Ablation studies on memory design. RT, FD, EP, and RF denote Raw Text Memory, Fact Knowledge Memory, Episode memory, and the Refresher.

addition to the LLM Score. For LongMemEval<sub>s</sub>, we specifically select the more challenging Pass@1 accuracy evaluated by an LLM judge to rigorously test performance. Following Achiam et al. (2023); Maharana et al. (2024), we employ GPT-4o-mini as the unified judge for all model-based evaluations.

**Baselines.** We compare AMA with various baselines, starting with FullContext and a standard Retrieval-Augmented Generation (RAG) implemented with 2048-token chunks (Lewis et al., 2020). We then evaluate representative memory frameworks including: LangMem (Chase, 2024), MemGPT (Packer et al., 2023a), Zep (Rasmussen et al., 2025), A-Mem (Xu et al., 2025), Mem0 (Chhikara et al., 2025), Nemori (Nan et al., 2025).

**Implementation Details.** We conduct experiments using both closed-source APIs (GPT-4o-mini, GPT-4.1-mini (Achiam et al., 2023)) and open-source models (Qwen3-8B-Instruct, Qwen3-30B-Instruct (Yang et al., 2025)), ensuring that the AMA framework utilizes the identical backbone model as the response generator. To ensure reproducibility, we fix the temperature to 0 for all experiments. For RAG, the retrieval top- $k$  is set to 10, while for AMA the maximum retrieval loop  $K_r$  is limited to 2. All memory embeddings are computed using OpenAI’s text-embedding-3-large model. Due to the commercial nature of Zep, Mem0, and MemGPT, we exclude these frameworks from evaluations involving open-source models. Additionally, given the substantial scale of LongMemEval<sub>s</sub> (approximately 58M tokens), we restrict its evaluation exclusively to GPT-4o-mini for computational feasibility. Prompts used in AMA and descriptions of the baselines are provided in Appendix A and C.

## 4.2 Main Results

We first evaluate AMA on the LoCoMo benchmark using closed-source backbones. As shown in Table 1, with GPT-4o-mini, AMA achieves an overall LLM Score of 0.774, substantially outperforming the strongest baseline Nemori (0.740) and

Method	Tokens	Latency (s)	LLM Score
FullContext	18625	7.206	0.717
RAG	5800	<b>2.983</b>	0.300
Nemori	2925	3.152	0.740
Zep	2461	3.255	0.580
Mem0	<b>1340</b>	3.739	0.608
A-Mem	2720	3.227	0.476
AMA ( $K_r=1$ )	2491	3.124	0.723
AMA ( $K_r=2$ )	3613	3.910	<b>0.774</b>

Table 4: Efficiency–performance trade-off on LoCoMo. We report input token usage, inference latency, and LLM Score for different methods.

all other memory-based methods by a clear margin. When scaled to the more capable GPT-4.1-mini, AMA further improves to 0.805. Notably, under this setting, AMA is the only approach that surpasses FullContext (0.786), demonstrating that AMA effectively distills raw history into critical facts and episodes, thereby filtering out noise to support reasoning beyond the raw context window.

We further assess robustness by extending the evaluation to open-source models on LoCoMo. With Qwen3-30B-Instruct, AMA attains a dominant LLM Score of 0.791, exceeding FullContext (0.733) by a large margin of 0.058. This advantage persists even with the smaller Qwen3-8B-Instruct backbone, where AMA (0.707) continues to outperform FullContext (0.696). These results demonstrate that AMA consistently enhances complex reasoning performance across backbones of varying capacity.

Finally, we evaluate the generalization of AMA on another benchmark, LongMemEval<sub>s</sub> (Table 2). AMA again achieves the highest average accuracy of 0.698, outperforming the two strongest baselines, Nemori and Zep, by 0.056 and 0.066, respectively. Notably, AMA attains near-perfect accuracy on single-session-user tasks (0.986) and shows a pronounced advantage on knowledge-update tasks (0.897), where dynamic knowledge maintenance and conflict resolution are critical. These consistent improvements across benchmarks with different data distributions indicate that AMA generalizes effectively to diverse long-term reasoning scenarios while robustly supporting dynamic knowledge evolution.

## 4.3 Ablation Studies

We conduct ablation studies on LoCoMo and LongMemEval<sub>s</sub>, with results in Table 3.

**Impact of Memory Granularities.** We first analyze the contribution of individual memory gran-




Figure 4: Effect of retrieval round limit  $K_r$ . The left and middle panels show that increasing  $K_r$  improves performance on LoCoMo and LongMemEval<sub>s</sub> with diminishing returns, while the right panel illustrates the corresponding growth in token consumption and inference latency.

ularities. Under single-granularity settings, **Fact Knowledge Memory** performs best, achieving an LLM Score of 0.712 on LoCoMo and the highest average accuracy of 0.642 on LongMemEval<sub>s</sub>, indicating the effectiveness of structured factual representations for long-term retrieval and reasoning. Moreover, jointly enabling Raw Text, Fact Knowledge, and Episodic Memory yields the strongest overall performance across both benchmarks, outperforming any single-granularity configuration. This result highlights the complementary nature of different memory forms and underscores the importance of multi-granularity collaboration.

**Effectiveness of the Refresher.** Beyond memory representation, we evaluate the role of the **Refresher** in maintaining long-term consistency. Under the full multi-granularity setting, enabling the Refresher substantially improves performance on knowledge-update scenarios in LongMemEval<sub>s</sub>, achieving an accuracy of 0.897. In contrast, removing the Refresher leads to a sharp drop to 0.568. This pronounced degradation indicates that accurate long-term memory requires not only multi-granularity storage, but also explicit mechanisms for conflict resolution and memory updating.

#### 4.4 Efficiency Analysis

We evaluate the efficiency–performance trade-off of AMA on the LoCoMo benchmark, with results reported in Table 4. Compared to FullContext, which processes 18625 input tokens with a latency of 7.21 seconds, AMA substantially reduces input length while maintaining strong performance. With the default setting  $K_r = 2$ , AMA requires only 3613 tokens, approximately 19% of FullContext, with a latency of 3.91 seconds, while achieving the highest LLM Score of 0.774 among all compared memory frameworks. This configuration represents a favorable balance between efficiency

and accuracy. Even under the more efficient setting ( $K_r = 1$ ), AMA attains an LLM Score of 0.723, only lower than Nemori (0.740), while operating within a comparable latency range to Nemori, Zep, Mem0, and A-Mem. These results show that AMA maintains competitive reasoning capability at low retrieval depth and offers a flexible trade-off between computational efficiency and performance.

#### 4.5 Analysis of the Retrieval Round Limit $K_r$

Figure 4 analyzes the impact of the retrieval round limit  $K_r$  on both model performance and computational cost. Increasing  $K_r$  from 1 to 3 yields consistent performance gains on LoCoMo and LongMemEval<sub>s</sub>, suggesting that additional retrieval rounds progressively surface useful historical information for long-term reasoning. However, the improvement exhibits clear diminishing returns, with performance largely saturating beyond  $K_r \geq 5$ . In contrast, both input token consumption and inference latency grow approximately linearly with  $K_r$ , reflecting the increasing overhead of deeper retrieval. Balancing these trends, we adopt  $K_r = 2$  as the default setting, which achieves near-optimal performance while substantially reducing token usage and latency, offering an effective trade-off between long-term reasoning quality and computational efficiency in practice.

In addition, we present a detailed case study in Appendix B, which demonstrates AMA’s capabilities in adaptive retrieval and conflict resolution.

## 5 Conclusion

In this work, we introduce AMA, a multi-agent memory framework for long-term interactions that integrates multi-granularity memory, adaptive routing, and principled memory maintenance. By decomposing the memory lifecycle into coordinated agent roles, AMA dynamically aligns retrieval

granularity with task demands while maintaining memory consistency over time. Extensive experiments demonstrate that AMA consistently outperforms strong baselines across challenging long-context benchmarks, validating the effectiveness of our design. Overall, this work underscores the importance of adaptive retrieval control and long-term memory management for building robust and scalable LLM agents.

## Limitations

Despite the significant performance gains, the multi-agent collaboration incurs a moderate computational overhead compared to static retrieval baselines. Additionally, the reliance on the backbone model’s reasoning capabilities suggests that the system’s efficiency on smaller architectures has room for further optimization. We aim to address these challenges in future work to further enhance the efficiency and universality of the framework.

## References

- Iman Abbasnejad, Xuefeng Liu, and Atanu Roy. 2025. Deciding the path: Leveraging multi-agent systems for solving complex tasks. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 4216–4225.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Harrison Chase. 2024. Langmem. <https://github.com/langchain-ai/langmem>. LangChain project. Accessed: 2025-07-20.
- Prateek Chhikara, Dev Khant, Saket Aryan, Taranjeet Singh, and Deshraj Yadav. 2025. Mem0: Building production-ready ai agents with scalable long-term memory. *arXiv preprint arXiv:2504.19413*.
- Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, and 1 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Sam Stevens, Boshi Wang, Huan Sun, and Yu Su. 2023. Mind2web: Towards a generalist agent for the web. *Advances in Neural Information Processing Systems*, 36:28091–28114.
- Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvasy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. 2025. The faiss library. *IEEE Transactions on Big Data*.
- Fatemeh Haji, Mazal Bethany, Maryam Tabar, Jason Chiang, Anthony Rios, and Peyman Najafirad. 2024. Improving llm reasoning with multi-agent tree-of-thought validator agent. *arXiv preprint arXiv:2409.11527*.
- Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, and 1 others. 2023. Metagpt: Meta programming for a multi-agent collaborative framework. In *The Twelfth International Conference on Learning Representations*.
- Yuanzhe Hu, Yu Wang, and Julian McAuley. 2025. Evaluating memory in llm agents via incremental multi-turn interactions. *arXiv preprint arXiv:2507.05257*.
- Rodnry Huddleston and Geqffry Pullum. 2005. The cambridge grammar of the english language. *Zeitschrift für Anglistik und Amerikanistik*, 53(2):193–194.
- Fantine Huot, Reinald Kim Amplayo, Jennimaria Palomaki, Alice Shoshana Jakobovits, Elizabeth Clark, and Mirella Lapata. 2024. Agents’ room: Narrative generation through multi-step collaboration. *arXiv preprint arXiv:2410.02603*.
- Dong-Ho Lee, Adyasha Maharana, Jay Pujara, Xiang Ren, and Francesco Barbieri. 2025. Realtalk: A 21-day real-world dataset for long-term conversation. *arXiv preprint arXiv:2502.13270*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, and 1 others. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:9459–9474.
- Guannan Liang and Qianqian Tong. 2025. Llm-powered ai agent systems and their applications in industry. *arXiv preprint arXiv:2505.16120*.
- Hehai Lin, Shilei Cao, Sudong Wang, Haotian Wu, Minzhi Li, Linyi Yang, Juepeng Zheng, and Chengwei Qin. 2025. Interactive learning for llm reasoning. *arXiv preprint arXiv:2509.26306*.
- Lei Liu, Xiaoyan Yang, Yue Shen, Binbin Hu, Zhiqiang Zhang, Jinjie Gu, and Guannan Zhang. 2023. Think-in-memory: Recalling and post-thinking enable llms with long-term memory. *arXiv preprint arXiv:2311.08719*.
- Zhiwei Liu, Weiran Yao, Jianguo Zhang, Liangwei Yang, Zuxin Liu, Juntao Tan, Prafulla K Choubey, Tian Lan, Jason Wu, Huan Wang, and 1 others. 2024.

- Agentlite: A lightweight library for building and advancing task-oriented llm agent system. *arXiv preprint arXiv:2402.15538*.
- Adyasha Maharana, Dong-Ho Lee, Sergey Tulyakov, Mohit Bansal, Francesco Barbieri, and Yuwei Fang. 2024. Evaluating very long-term conversational memory of llm agents. *arXiv preprint arXiv:2402.17753*.
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. When not to trust language models: Investigating effectiveness of parametric and non-parametric memories. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9802–9822.
- Kai Mei, Xi Zhu, Wujiang Xu, Wenyue Hua, Mingyu Jin, Zelong Li, Shuyuan Xu, Ruosong Ye, Yingqiang Ge, and Yongfeng Zhang. 2024. Aios: Llm agent operating system. *arXiv preprint arXiv:2403.16971*.
- Jiayan Nan, Wenquan Ma, Wenlong Wu, and Yize Chen. 2025. Nemori: Self-organizing agent memory inspired by cognitive science. *arXiv preprint arXiv:2508.03341*.
- Mike Owens and Grant Allen. 2010. *SQLite*. Apress LP New York.
- Charles Packer, Vivian Fang, Shishir G. Patil, Kevin Lin, Sarah Wooders, and Joseph E. Gonzalez. 2023a. **Memgpt: Towards llms as operating systems**. *CoRR*, abs/2310.08560.
- Charles Packer, Vivian Fang, Shishir\_G Patil, Kevin Lin, Sarah Wooders, and Joseph\_E Gonzalez. 2023b. **Memgpt: Towards llms as operating systems**.
- Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22.
- Chen Qian, Wei Liu, Hongzhang Liu, Nuo Chen, Yufan Dang, Jiahao Li, Cheng Yang, Weize Chen, Yusheng Su, Xin Cong, and 1 others. 2024. Chatdev: Communicative agents for software development. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15174–15186.
- Hongjin Qian, Zheng Liu, Peitian Zhang, Kelong Mao, Defu Lian, Zhicheng Dou, and Tiejun Huang. 2025. Memorag: Boosting long context processing with global memory-enhanced retrieval augmentation. In *Proceedings of the ACM on Web Conference 2025*, pages 2366–2377.
- Preston Rasmussen, Pavlo Paliychuk, Travis Beauvais, Jack Ryan, and Daniel Chalef. 2025. Zep: a temporal knowledge graph architecture for agent memory. *arXiv preprint arXiv:2501.13956*.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.
- Alireza Rezaeizadeh, Zichao Li, Wei Wei, and Yujia Bao. 2024. From isolated conversations to hierarchical schemas: Dynamic tree memory representation for llms. *arXiv preprint arXiv:2410.14052*.
- Theodore Sumers, Shunyu Yao, Karthik Narasimhan, and Thomas Griffiths. 2023. Cognitive architectures for language agents. *Transactions on Machine Learning Research*.
- Zhen Tan, Jun Yan, I-Hung Hsu, Rujun Han, Zifeng Wang, Long Le, Yiwen Song, Yanfei Chen, Hamid Palangi, George Lee, and 1 others. 2025. In prospect and retrospect: Reflective memory management for long-term personalized dialogue agents. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8416–8439.
- Lukas Thede, Karsten Roth, Matthias Bethge, Zeynep Akata, and Tom Hartvigsen. 2025. Understanding the limits of lifelong knowledge editing in llms. *arXiv preprint arXiv:2503.05683*.
- Bing Wang, Xinnian Liang, Jian Yang, Hui Huang, Shuangzhi Wu, Peihao Wu, Lu Lu, Zejun Ma, and Zhoujun Li. 2023. Enhancing large language model with self-controlled memory framework. *arXiv preprint arXiv:2304.13343*.
- Cunxiang Wang, Ruoxi Ning, Boqi Pan, Tonghui Wu, Qipeng Guo, Cheng Deng, Guangsheng Bao, Xi-angkun Hu, Zheng Zhang, Qian Wang, and 1 others. 2024a. Novelqa: Benchmarking question answering on documents exceeding 200k tokens. *arXiv preprint arXiv:2403.12766*.
- Song Wang, Yaochen Zhu, Haochen Liu, Zaiyi Zheng, Chen Chen, and Jundong Li. 2024b. Knowledge editing for large language models: A survey. *ACM Computing Surveys*, 57(3):1–37.
- Yu Wang and Xi Chen. 2025. Mirix: Multi-agent memory system for llm-based agents. *arXiv preprint arXiv:2507.07957*.
- Yu Wang, Ryuichi Takanobu, Zhiqi Liang, Yuzhen Mao, Yuanzhe Hu, Julian McAuley, and Xiaojian Wu. 2025. Mem- $\{\alpha\}$ : Learning memory construction via reinforcement learning. *arXiv preprint arXiv:2509.25911*.
- Yu Wang, Ruihan Wu, Zexue He, Xiushi Chen, and Julian McAuley. 2024c. Large scale knowledge washing. *arXiv preprint arXiv:2405.16720*.
- Di Wu, Hongwei Wang, Wenhao Yu, Yuwei Zhang, Kai-Wei Chang, and Dong Yu. 2024. Longmemeval: Benchmarking chat assistants on long-term interactive memory. *arXiv preprint arXiv:2410.10813*.

- Yaxiong Wu, Sheng Liang, Chen Zhang, Yichao Wang, Yongyue Zhang, Huifeng Guo, Ruiming Tang, and Yong Liu. 2025. From human memory to ai memory: A survey on memory mechanisms in the era of llms. *arXiv preprint arXiv:2504.15965*.
- Yiran Wu, Feiran Jia, Shaokun Zhang, Hangyu Li, Erkang Zhu, Yue Wang, Yin Tat Lee, Richard Peng, Qingyun Wu, and Chi Wang. 2023. Mathchat: Converse to tackle challenging math problems with llm agents. *arXiv preprint arXiv:2306.01337*.
- Wujiang Xu, Zujie Liang, Kai Mei, Hang Gao, Juntao Tan, and Yongfeng Zhang. 2025. A-mem: Agentic memory for llm agents. *arXiv preprint arXiv:2502.12110*.
- Sikuan Yan, Xiufeng Yang, Zuchao Huang, Ercong Nie, Zifeng Ding, Zonggen Li, Xiaowen Ma, Kristian Kersting, Jeff Z Pan, Hinrich Schütze, and 1 others. 2025. Memory-r1: Enhancing large language model agents to manage and utilize memories via reinforcement learning. *arXiv preprint arXiv:2508.19828*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Shaowei Zhang and Deyi Xiong. 2025. Debate4math: Multi-agent debate for fine-grained reasoning in math. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 16810–16824.
- Zeyu Zhang, Quanyu Dai, Xiaohe Bo, Chen Ma, Rui Li, Xu Chen, Jieming Zhu, Zhenhua Dong, and Ji-Rong Wen. 2025. A survey on the memory mechanism of large language model-based agents. *ACM Transactions on Information Systems*, 43(6):1–47.
- Wanjun Zhong, Lianghong Guo, Qiqi Gao, He Ye, and Yanlin Wang. 2024. Memorybank: Enhancing large language models with long-term memory. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19724–19731.

## A Experiment Details

In this appendix, we provide additional experimental details to support reproducibility and clarity. We first describe the datasets used in our experiments, followed by the evaluation metrics employed for performance assessment. We then introduce the baseline methods considered for comparison, and finally present the implementation details of the proposed AMA framework, including system components and experimental configurations.

### A.1 Datasets

**LoCoMo (Maharana et al., 2024)** It is a large-scale benchmark for evaluating very long-term

conversational memory of LLM agents, consisting of 10 conversations that span an average of 27.2 sessions and 21.6 turns per session, with each conversation containing approximately 16.6K tokens. In our experiments, we evaluate models on 1,540 question-answering samples, which are categorized into 841 single-hop retrieval questions, 282 multi-hop retrieval questions, 321 temporal reasoning questions, and 96 open-domain knowledge questions, all of which require accurate recall and reasoning over long-range conversational histories. Beyond question answering, LoCoMo further includes an event summarization task grounded in temporally structured event graphs, as well as a multimodal dialogue generation task involving natural image sharing behaviors, providing a comprehensive benchmark for assessing long-term memory, temporal understanding, and multimodal consistency in LLM-based agents.

**LongMemEval, (Wu et al., 2024)** It is a benchmark for assessing long-term memory in user–assistant interactions under a standardized long-context setting. Inspired by the “needle-in-a-haystack” paradigm, it compiles a coherent yet length-configurable chat history for each question, and provides a standard setting where each problem is paired with an interaction history of approximately 115k tokens. In our experiments, we use LongMemEval<sub>s</sub> and evaluate on 500 question-answering instances, including 70 single-session-user, 133 multi-session, 30 single-session-preference, 133 temporal-reasoning, 78 knowledge-update, and 56 single-session-assistant questions, covering diverse memory abilities such as extracting user-provided information, synthesizing evidence across sessions, reasoning with temporal references, handling updated user facts, and recalling assistant-provided information.

### A.2 Evaluation Metric

We evaluate the quality of generated answers using F1 score and BLEU-1, while employing cosine similarity as the similarity measure during the retrieval stage.

The F1 score represents the harmonic mean of precision and recall, providing a balanced metric that jointly considers both correctness and completeness of the predicted answers:

$$F1 = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}} \quad (1)$$

where

$$\text{precision} = \frac{\text{true positives}}{\text{true positives} + \text{false positives}} \quad (2)$$

$$\text{recall} = \frac{\text{true positives}}{\text{true positives} + \text{false negatives}} \quad (3)$$

In question-answering tasks, the F1 score is widely used to measure the overlap between predicted and reference answers, especially for short-form or span-based responses where exact matching is overly restrictive.

BLEU-1 further evaluates unigram-level precision between the generated response and the reference text:

$$\text{BLEU-1} = BP \cdot \exp\left(\sum_{n=1}^1 w_n \log p_n\right) \quad (4)$$

where

$$BP = \begin{cases} 1, & \text{if } c > r \\ e^{1-r/c}, & \text{if } c \leq r \end{cases} \quad (5)$$

$$p_n = \frac{\sum_i \sum_k \min(h_{ik}, m_{ik})}{\sum_i \sum_k h_{ik}} \quad (6)$$

Here,  $c$  and  $r$  denote the lengths of the candidate and reference sequences, respectively, while  $h_{ik}$  and  $m_{ik}$  represent unigram counts in the hypothesis and reference texts.

In the retrieval stage, we use cosine similarity as the similarity measure in the embedding space to identify relevant memory entries. Specifically, given a query embedding  $\mathbf{q}$  and a set of memory embeddings  $\{\mathbf{m}_i\}$ , we compute their cosine similarity as:

$$\text{CosineSim}(\mathbf{q}, \mathbf{m}_i) = \frac{\mathbf{q} \cdot \mathbf{m}_i}{\|\mathbf{q}\|_2 \|\mathbf{m}_i\|_2} \quad (7)$$

and retrieve the top- $K$  memory entries with the highest similarity scores. Cosine similarity measures the angular similarity between vectors and is invariant to vector magnitude, making it well suited for embedding-based semantic retrieval.

### A.3 Baselines

**RAG (Lewis et al., 2020)** As a strong and widely adopted baseline, we implement Retrieval-Augmented Generation (RAG), which enhances language models by retrieving external textual evidence and conditioning generation on the retrieved content. RAG decomposes the generation process into a retrieval stage and a generation stage, where a dense retriever is first used to select the top- $k$  most

relevant memory entries given the input query, and the retrieved texts are then concatenated with the query as context for the generator. In our implementation, all memory entries are embedded into a shared vector space, and cosine similarity is used to retrieve the top- $k$  relevant records for each query, with  $k$  fixed across experiments for fair comparison. The retrieved memory is directly appended to the input prompt without additional refinement, filtering, or memory updating, reflecting a standard single-shot retrieval paradigm. While RAG has demonstrated strong effectiveness in knowledge-intensive question answering, its retrieval process remains static and non-iterative, lacking mechanisms for retrieval refinement, conflict resolution, or long-term memory maintenance, which limits its effectiveness in long-horizon and multi-session reasoning scenarios.

**LangMem (Chase, 2024)** We include LangMem as a baseline that explicitly models long-term memory for LLM-based agents. LangMem maintains an external memory store to record historical user–assistant interactions, where past information is organized into structured textual representations and indexed for retrieval. During inference, LangMem retrieves memory entries that are semantically relevant to the current query and incorporates them into the model input to support response generation, enabling the model to leverage information accumulated over extended interaction histories.

**MemGPT (Packer et al., 2023b)** MemGPT is a baseline designed to address long-term context limitations through an OS-inspired memory management mechanism. It treats the LLM context window as a constrained resource and augments it with a hierarchical memory architecture that separates in-context memory from external persistent storage. Through predefined function calls, the model can autonomously write important information to external memory, retrieve relevant out-of-context records, and dynamically move information into or out of the active context window during inference. This design enables MemGPT to maintain and access information across extended interactions by paging memory between the limited context window and persistent storage.

**Zep (Rasmussen et al., 2025)** Zep is a memory layer designed for LLM-based agents that represents long-term conversational memory using a temporally aware knowledge graph. It ingests both

unstructured conversational messages and structured data, and incrementally constructs a dynamic graph composed of episodic nodes, semantic entities, and higher-level community abstractions. Temporal information is explicitly modeled, allowing facts and relationships to be associated with validity intervals and updated as new information arrives. During inference, Zep retrieves relevant graph elements through a combination of semantic similarity search, full-text search, and graph traversal, and converts the retrieved nodes and edges into structured textual context for response generation.

**Mem0 (Chhikara et al., 2025)** Mem0 is a memory-centric architecture designed to provide scalable long-term memory for LLM-based agents. It continuously processes conversational interactions and extracts salient factual information that is deemed useful for future reasoning. The extracted memories are stored in an external memory store as compact natural-language representations, each associated with semantic embeddings to support efficient similarity-based retrieval. Mem0 incorporates an explicit memory update mechanism that evaluates newly extracted information against existing memories, allowing the system to add new entries, update existing ones with refined content, or remove outdated or contradictory information. During inference, the model retrieves a small set of semantically relevant memory entries and conditions response generation on the retrieved memories, enabling consistent access to accumulated knowledge across extended and multi-session interactions.

**A-Mem (Xu et al., 2025)** A-Mem is an agentic memory system designed for LLM-based agents that enables dynamic organization and evolution of long-term memory. Inspired by the Zettelkasten method, A-Mem represents each interaction as an atomic memory note enriched with multiple structured attributes, including contextual descriptions, keywords, tags, timestamps, and dense semantic embeddings. When new memories are added, the system autonomously analyzes existing memory notes to identify semantically related entries and establishes meaningful links among them, forming an interconnected memory network. In addition, newly integrated memories can trigger updates to the contextual representations of existing notes, allowing the memory structure to continuously evolve over time. During inference, A-Mem retrieves semantically relevant memory notes based

on embedding similarity and augments them with linked memories to construct context for response generation, enabling agents to leverage structured and evolving long-term memory.

**Nemori (Nan et al., 2025)** Nemori is a self-organizing memory architecture for LLM-based agents inspired by principles from cognitive science. It autonomously structures long conversational streams into semantically coherent episodic units through a top-down boundary detection mechanism, avoiding arbitrary or fixed-granularity segmentation. Each episode is transformed into a structured narrative representation and stored as episodic memory, while a complementary semantic memory is incrementally distilled through a predict-calibrate process that identifies and integrates novel information from prediction gaps. During inference, Nemori retrieves relevant episodic and semantic memories using dense similarity search and incorporates them into the model context, enabling effective utilization of long-term interaction history.

#### A.4 Framework Implementation Details

Our framework adopts a modular system design to support efficient long-term memory storage, retrieval, and dynamic utilization. At the implementation level, the system is composed of three core components: a structured memory storage module, a vector-based retrieval module, and a configurable inference controller that adapts execution behavior based on user requirements. This design balances scalability, efficiency, and engineering simplicity.

For memory storage, we employ **SQLite** as a lightweight relational database to manage structured memory content (Owens and Allen, 2010). SQLite is responsible for persistently storing processed memory entries along with their associated metadata, including textual content, timestamps, session identifiers, memory types, and auxiliary indexing fields. Owing to its zero-configuration nature, transactional support, and efficient local read/write performance, SQLite enables reliable long-term memory persistence without introducing additional service dependencies, making it well suited for experimental and single-node deployment scenarios.

For semantic retrieval, we utilize the **FAISS** vector retrieval library to enable efficient similarity search over large memory collections (Douze et al., 2025). Memory entries are first encoded into dense

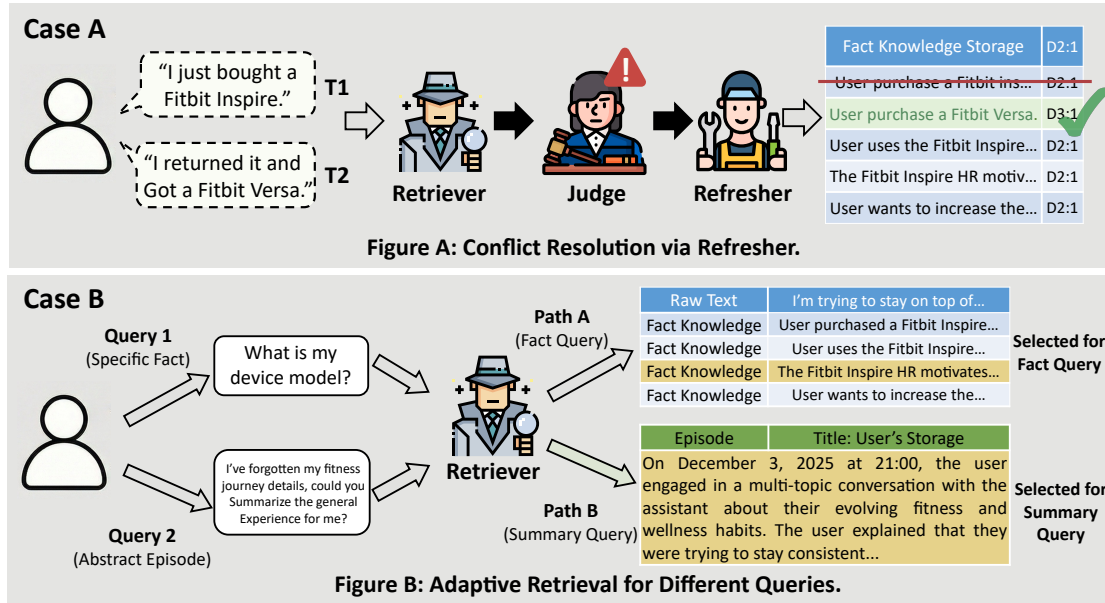


Figure 5: **Case Study.** (1) The upper part of the figure shows conflict resolution, where outdated factual memories are updated to maintain consistency. (2) The lower part of the figure shows adaptive retrieval, routing queries to different memory types based on intent.

vector representations and indexed using FAISS. During retrieval, the incoming query is mapped into the same embedding space, and the system performs similarity-based search to retrieve the Top- $K$  most relevant memory records. The use of FAISS significantly reduces retrieval latency under long interaction histories and large memory scales, while maintaining retrieval accuracy.

The inference pipeline further supports dynamic execution modes based on user requirements. Specifically, the framework provides both a *retrieval-only mode* and a *full inference mode*. In the retrieval-only mode, the system executes memory retrieval and directly returns the most relevant memory entries, which is useful for information lookup and debugging. In contrast, the full inference mode integrates the retrieved memories with the current user input and conditions the language model on the combined context to generate a final response. This configurable design allows the framework to flexibly trade off computational cost and response completeness across different application scenarios.

Overall, by combining **SQLite** for structured memory management and **FAISS** for efficient vector-based retrieval, together with configurable inference modes, the framework provides a robust and extensible implementation foundation for long-term memory modeling in LLM-based agents.

## B Case Study

In Figure 5, we conduct a qualitative case study to provide an in-depth illustration of how the proposed framework operates over long interaction histories. This case study demonstrates two core capabilities of AMA. (1) The upper part shows conflict resolution via the Refresher module: when the user provides contradictory device information across turns, the Retriever recalls relevant facts and the Judge identifies the inconsistency, triggering the Refresher to update former entries in the Fact Knowledge Memory to ensure temporal consistency. (2) The lower part illustrates adaptive retrieval across query intents, where fact-oriented queries are routed to the Fact Knowledge Memory for precise factual recall, while abstract summarization queries retrieve corresponding episodic memory chunks to provide high-level summaries, supporting coherent long-term reasoning.

## C Prompt Templates

In the AMA framework, multiple prompt templates are employed at different stages of the system. Specifically,  $P_{\text{con}}$  is used for memory construction,  $P_{\text{tri}}$  and  $P_{\text{epi}}$  are used for episode triggering and episode synthesis, respectively,  $P_{\text{ret}}$  is used for query rewriting and retrieval routing,  $P_{\text{jud}}$  is used for memory verification and consistency checking, and  $P_{\text{ref}}$  is used for memory updating and mainte-

nance. In addition, during evaluation, a separate prompt  $P_{llm}$  is adopted for LLM-as-Judge to automatically assess and compare model outputs. The concrete prompt templates used in each stage are provided in this appendix.

### C.1 Prompt Template of Constructor ( $P_{con}$ )

As shown in Figure 6, the Constructor prompt  $P_{con}$  guides the model to transform the current user input into structured and atomic memory representations. It enforces strict syntactic and semantic constraints to ensure that the constructed memories are stable, parsable, and suitable for long-term storage.

### C.2 Prompt Template of Episode Triggering ( $P_{tri}$ )

As illustrated in Figure 7, the episode triggering prompt  $P_{tri}$  is used to determine whether the current interaction should activate episodic memory construction. This prompt enables the system to selectively trigger high-level abstraction based on dialogue dynamics and contextual signals.

### C.3 Prompt Template of Episode Generation ( $P_{epi}$ )

Figure 8 presents the episode synthesis prompt  $P_{epi}$ , which is responsible for generating an abstract summary once episodic memory is activated. The prompt encourages concise and semantically coherent representations that capture the high-level meaning of a dialogue segment.

### C.4 Prompt Template of Retriever ( $P_{ret}$ )

As shown in Figure 9, the retriever prompt  $P_{ret}$  guides the model to rewrite the input query and infer retrieval intents. It produces structured signals that support dynamic routing to appropriate memory granularities during retrieval.

### C.5 Prompt Template of Judge ( $P_{jud}$ )

Figure 10 illustrates the judge prompt  $P_{jud}$ , which enables LLM-based verification of retrieved memory candidates. This prompt is used to assess relevance and consistency, producing validated memory sets and control decisions for subsequent system actions.

### C.6 Prompt Template of Refresher ( $P_{ref}$ )

As depicted in Figure 11, the refresher prompt  $P_{ref}$  is applied when memory conflicts are detected. It

guides the model to update or remove inconsistent memory entries in order to maintain long-term coherence.

### C.7 Prompt Template of LLM-as-Judge for Evaluation ( $P_{llm}$ )

Figure 12 shows the evaluation prompt  $P_{llm}$ , which is used to implement LLM-as-Judge during experimental evaluation. This prompt enables automatic assessment and comparison of model outputs under a unified evaluation protocol.

**You are a Constructor Agent for structured memory construction.**

Given the current user input  $u_t$ , decompose it into **atomic, parsable facts** for long-term memory. All facts must strictly follow **canonical SVO-based patterns** and be derived **only from the current input**.

**[Syntactic Elements Definition]**

1. S = Subject
2. V = Verb
3. O = Object
4. C = Complement (attribute or state)
5. L = Location or explicit time

**[Sentence Pattern Constraint]**

Each fact **MUST** follow exactly one of the following forms:

1. S-V
2. S-V-O
3. S-V-C
4. S-V-O-C
5. S-V-O-L

Any fact not matching these patterns is invalid.

**[Atomicity Rules]**

- Each fact represents **one single relation only**.
- Do **NOT** merge actions, attributes, locations, or roles.
- Descriptive or prepositional phrases must be split into separate facts.

**[Appositive Rule]**

If the input contains appositive or implicit equivalence (e.g., “my friend John”, “my home country, Sweden”), you **MUST** convert it into an explicit fact (e.g., “John is her friend”, “Sweden is her home country”). Do not omit appositive relations.

**[Source Constraint]**

- Facts must originate strictly from the current user input  $u_t$ .
- Memory window and retrieved context are for reference only.
- Do **NOT** introduce new facts from memory or retrievals.

**[Timestamp Rule]**

- Extract a timestamp **ONLY** if an explicit date or time is present.
- Normalize to “YYYY-MM-DD HH:MM”.
- Otherwise output “timestamp”: “empty”.
- Do **NOT** infer time from relative expressions.

**[Output Format]**

Return exactly one JSON object:

```
{
  "facts": [
    {"content": "<atomic fact>"},
    ...
  ],
  "source": "user",
  "related_id": ["<D_s,t>", ...],
  "timestamp": "<normalized time>" or "empty"
}
```

Figure 6: The prompt template for the Constructor Agent.

**You are an Episode Boundary Judge.**

Your task is to determine whether the newly added dialogue triggers the start of a new episode. This corresponds to a binary decision  $T_t \in \{0, 1\}$ .

**Input:**

Conversation history  $W_t$  (may be empty):

{conversation\_history}

Newly added messages  $u_t$ :

{new\_messages}

**Decision Criteria:**

Set  $T_t = 1$  if any of the following conditions is satisfied:

1) **Topic Shift**

The new messages introduce a topic, event, or task that is semantically independent from the current conversation, or the dialogue moves to an unrelated question or scenario.

2) **Explicit Transition Intent**

The user explicitly signals a transition to a new topic or request, such as starting a new question or changing the discussion focus.

3) **Context Saturation**

The current topic has been sufficiently discussed or concluded, or the dialogue scope exceeds what is reasonable for a single coherent episode.

Otherwise, set  $T_t = 0$ .

**Decision Principles:**

- Prioritize topic independence over conversational continuity.
- When uncertain, favor splitting and set  $T_t = 1$ .
- Each episode should correspond to one coherent topic or event.

**Output Format:**

Return exactly one JSON object:

```
{
  "T_t": 0 or 1,
  "reason": "<brief justification>",
  "confidence": 0.0-1.0,
  "topic_summary": "<summary ... if T_t = 1, otherwise empty>"
}
```

**Notes:** If the conversation history is empty, return  $T_t = 0$ . Output only the JSON object and nothing else.

Figure 7: The prompt template for the Episode Triggering.

**You are an Episodic Memory Constructor.**

Your task is to convert a completed dialogue episode into a structured episodic memory entry, written as a coherent third-person narrative.

**Input:**

Conversation episode  $E_t$ :

{conversation}

Episode boundary trigger reason:

{boundary\_reason}

**Construction Principles:**

- 1) The episode represents one coherent event or topic.
- 2) The description must preserve factual accuracy and chronological order.
- 3) Do not introduce information that is not present in the conversation.

**Time Handling Rules:**

- Identify the episode time from explicit timestamps in the dialogue.
- If relative time expressions appear in the conversation, convert them into absolute dates based on available context, and keep the converted time consistent throughout the episode.
- If no reliable time information exists, infer a reasonable episode time from the dialogue context.
- The episode timestamp represents when the episode occurred, not the current system time.

**Content Requirements:**

- Write in third-person narrative form.
- Clearly describe:
  - who participated,
  - what was discussed,
  - what decisions or conclusions were reached,
  - any expressed intentions, plans, or outcomes.
- Include time information explicitly in the narrative.
- Maintain causal and temporal coherence.

**Output Format:**

Return exactly one JSON object:

```
{  
  "title": "<concise episode title summarizing the core event or topic>",  
  "content": "<third-person episodic narrative describing what happened>",  
  "timestamp": "YYYY-MM-DDTHH:MM:SS"  
}
```

} Output only the JSON object and nothing else.

Figure 8: The prompt template for the Episodic Memory Generation.

**You are a Retriever.**

Given the current user input  $u_t$  and the short-term context window  $W_t$ , your task is to perform query rewriting and intent-based routing to the appropriate memory granularity.

**Input:**

Short-term context window  $W_t$  (may be empty): {memory\_window}

Current user input  $u_t$ : {user\_input}

**Task Definition:**

Simultaneously generate:

- 1) a standalone rewritten query  $u'_t$  with resolved references,
- 2) a four-dimensional binary intent vector  $B$ ,
- 3) a dynamic retrieval budget  $K_{dyn}$ ,
- 4) a target memory type for retrieval.

**Intent Vector Definition:**

$B = [b_{fine}, b_{abs}, b_{event}, b_{atomic}]$

- $b_{fine} = 1$  if the query requires fine-grained or exact details.
- $b_{abs} = 1$  if the query is abstract or summary-oriented.
- $b_{event} = 1$  if the query involves cross-turn, cross-time, or event-level semantics.
- $b_{atomic} = 1$  if the query is short, single-point, and factual.

Each dimension must be either 0 or 1.

**Memory Routing Rule:**

Select the target memory type  $M$  based on  $B$ :

- If  $b_{fine} = 1$ , route to **Raw Text Memory**.
- Else if  $b_{abs} = 1$  or  $b_{event} = 1$ , route to **Episode Memory**.
- Otherwise, route to **Fact Knowledge Memory**.

**Query Rewriting Rule:**

- $u'_t$  must be self-contained and unambiguous.
- Resolve references using  $W_t$  when necessary.
- Preserve explicit entities and identities.
- Do not introduce information not present in  $u_t$  or  $W_t$ .

**Retrieval Budget Rule:**

Set  $K_{dyn}$  as a small integer reflecting the expected retrieval scope. Broader or cross-event queries should use a larger  $K_{dyn}$ .

**Output Format:**

Return exactly one JSON object:

```
{
  "rewrite_query": "<u'_t>",
  "intent_vector": {
    "b_fine": 0 or 1,
    "b_abs": 0 or 1,
    "b_event": 0 or 1,
    "b_atomic": 0 or 1
  },
  "memory_type": "raw" | "fact" | "episode",
  "K_dyn": <int>
}
```

Output only the JSON object and nothing else.

Figure 9: The prompt template for the Retriever, incorporating intent-based memory routing.

**You are a Judge.**

Your task is to verify whether the retrieved candidate information  $H_t$  is sufficient to answer the current user input  $u_t$ , and to detect whether any factual conflict exists. Based on this verification, you must select one action from: **Pass, Retry, or Refresh.**

**Input:**

Candidate retrieved information  $H_t$  (may be empty): {information}

Current user input  $u_t$ : {user\_input}

**Verification Rules:**

1) **Relevance and Sufficiency**

The retrieved information is sufficient only if it satisfies the required granularity of the query  $u_t$ .

- **Entity requirement (who / what):** The exact entity or attribute asked in the query must be explicitly stated.
- **Location requirement (where):** The location must be concrete and match the requested level (e.g., city, country, or region).
- **Temporal requirement (when):** The time information must meet or exceed the precision required by the query. Relative expressions are acceptable only if they can be normalized to a unique value.
- **Quantity requirement (how many):** Numeric values or explicitly bounded ranges are required.

If any required element is missing or underspecified, the information is considered insufficient.

2) **No Proxy or Vague Answers**

Generic or proxy expressions such as “home country”, “recently”, “around that time”, or “somewhere” do not satisfy sufficiency. If only such information is available, treat the result as insufficient.

3) **Controlled Inference**

Only minimal one-step inference that does not reduce specificity is allowed. Do not infer unstated states, temporal validity, or current conditions.

4) **Conflict Detection**

A conflict exists if two or more pieces of information assert incompatible facts about the same event, entity, or timeframe. Differences in wording or tone alone do not constitute a conflict.

**Action Selection:**

- If a conflict is detected, set **Action = Refresh.**
- Else if the information is insufficient, set **Action = Retry.**
- Else, set **Action = Pass.**

**Output Format:**

Return exactly one JSON object:

```
{
  "Action": "Pass" | "Retry" | "Refresh",
  "reason": "<brief justification>",
  "confidence": 0.0-1.0
}
```

Output only the JSON object and nothing else.

Figure 10: The prompt template for the Judge, responsible for sufficiency checking and conflict detection.

**You are a Refresher.**

Your task is to maintain the temporal and logical consistency of stored memory. Given detected conflicts  $C_{err}$  and the corresponding memory entries  $H_t$ , decide whether to update or delete memory records.

You may operate on any memory granularity: **Raw Text Memory, Fact Knowledge Memory, or Episodic Memory.**

**Input:**

Conflicting memory entries  $H_t$  (may be empty): {memory\_entries}

Current user input  $u_t$ : {user\_input}

**Editing Principles:****1) Scope Constraint**

Only modify memory entries that refer to the same entity or event as  $u_t$ . If the input concerns unrelated content, do nothing.

**2) Update Rule**

Perform **Update** when:

- The same entity or event is referenced;
- A factual inconsistency exists in state, time, quantity, or outcome;
- The new input provides a more up-to-date or correct value.

Update only the conflicting attributes, and preserve all other information.

**3) Delete Rule**

Perform **Delete** only when:

- The user explicitly requests removal, cancellation, or invalidation;
- Or the memory entry is no longer valid by explicit user instruction.

**4) Conservative Default**

If the situation is ambiguous, perform no modification.

**Timestamp Rule:**

- Extract a timestamp only if the user input contains an explicit date or time.
- Normalize the timestamp if possible.
- If no explicit time is provided, set "timestamp" to "empty".

**Output Format:**

Return exactly one JSON object:

```
{
  "Action": "Update" | "Delete" | "No-Op",
  "memory_type": "raw" | "fact" | "episode",
  "dataList": [
    {
      "id": "<memory_id>",
      "new_content": "<updated content...>"
    }
  ],
  "timestamp": "<normalized date/time>" or "empty",
  "reason": "<brief justification>"
}
```

Output only the JSON object and nothing else.

Figure 11: The prompt template for the Refresher, handling memory updates and conflict resolution.

**You are an Accuracy Judge.**

Your task is to label a generated answer as '**CORRECT**' or '**WRONG**' by comparing it against a ground truth ('gold') answer.

**Input Data:**

- **Question ( $Q$ ):** A question posed by one user to another based on prior conversations.
- **Gold Answer ( $A_{gold}$ ):** The concise ground truth (e.g., specific object, date, or event).
- **Generated Answer ( $A_{gen}$ ):** The model's response to be evaluated.

**Current Instance:**

Question: {question}

Gold Answer: {gold\_answer}

Generated Answer: {generated\_answer}

**Evaluation Criteria:**

1) **Semantic Generosity (Topic Matching)**

The generated answer  $A_{gen}$  may be verbose. As long as it touches on the **same core topic or object** as  $A_{gold}$  (e.g., "A shell necklace" vs. "I think it was a necklace made of shells from the trip"), mark it as **CORRECT**.

2) **Temporal Flexibility**

For time-related questions:

- **Formats:** Accept different date formats (e.g., "May 7th" vs "7 May").
- **Relative References:** Accept relative terms (e.g., "last Tuesday", "next month") if they logically refer to the same time period as  $A_{gold}$ .

If the time reference resolves to the same value, mark it as **CORRECT**.

3) **Binary Labeling**

You must output exactly one label. Do NOT include both 'CORRECT' and 'WRONG' in the reasoning to avoid parsing errors.

**Output Format:**

First, provide a short (one sentence) explanation of your reasoning, then return the label in JSON format:

```
{
  "reasoning": "<one sentence explanation>",
  "label": "CORRECT" | "WRONG"
}
```

Output only the JSON object.

Figure 12: The prompt template for the LLM-as-a-Judge, used to evaluate the factual accuracy of answers.