

TELEMEM: Building Long-Term and Multimodal Memory for Agentic AI

Chunliang Chen*, Ming Guan*, Xiao Lin*, Jiaxu Li*, Qiyi Wang*,
Xiangyu Chen, Jixiang Luo, Changzhi Sun†, Dell Zhang#, Xuelong Li#
Institute of Artificial Intelligence (TeleAI), China Telecom

<https://github.com/TeleAI-UAGI/TeleMem>

Abstract

Large language models (LLMs) excel at many NLP tasks but struggle to sustain long-term interactions due to limited attention over extended dialogue histories. Retrieval-augmented generation (RAG) mitigates this issue but lacks reliable mechanisms for updating or refining stored memories, leading to schema-driven hallucinations, inefficient write operations, and minimal support for multimodal reasoning. To address these challenges, we propose TELEMEM, a unified long-term and multimodal memory system that maintains coherent user profiles through narrative dynamic extraction, ensuring that only dialogue-grounded information is preserved. TELEMEM further introduces a structured writing pipeline that batches, retrieves, clusters, and consolidates memory entries, substantially improving storage efficiency, reducing token usage, and accelerating memory operations. Additionally, a multimodal memory module combined with ReAct-style reasoning equips the system with a closed-loop observe, think, and act process that enables accurate understanding of complex video content in long-term contexts. Experimental results show that TELEMEM surpasses the state-of-the-art Mem0 baseline with 19% higher accuracy, 43% fewer tokens, and a $2.1\times$ speedup on the ZH-40 long-term role-play gaming benchmark.

1. Introduction

Large language models (LLMs) have demonstrated remarkable performance across a wide range of natural language processing tasks [5, 29, 37]. However, their effectiveness in long-term interactions remains fundamentally constrained by the context window limitations inherent to Transformer architectures. Even though contemporary long-context models can process hundreds of thousands of tokens [11], simply enlarging the context window does not fully address the issue: as interaction rounds accumulate, models struggle

to allocate attention to distant tokens, making it difficult to accurately retrieve user-specific information embedded deep in the dialogue history. As a result, they may overlook key details or even forget important personalized information over time [13], significantly hindering their practical utility in scenarios requiring sustained, long-term relationships with users.

To address the context-window limitations, retrieval-augmented generation (RAG) has become a widely adopted strategy for managing long dialogue histories. RAG systems typically encode past interactions into vector embeddings stored in external datastores and retrieve relevant information via semantic search [7, 21, 44]. Although effective at retrieving information beyond the model’s native context window, traditional RAG pipelines fundamentally lack mechanisms for updating or refining stored memories. Once written into the datastore, memories cannot be reliably modified or deleted, making it difficult to handle evolving user preferences, resolve contradictions, or maintain long-term consistency.

This inherent limitation has motivated the development of more advanced memory architectures capable of structured storage, organization, and dynamic updating of information [36, 40]. For example, MemoryBank [45] introduces a decay mechanism inspired by the Ebbinghaus forgetting curve to selectively preserve salient information. Mem0 [2] proposes a meaning-aware memory framework in which the LLM itself extracts atomic facts, retrieves semantically related memories, and decides whether to add, update, delete, or ignore them in a streaming, turn-by-turn process. Rsum [25] employs hierarchical summarization to distinguish short-term details from long-term abstractions, while Zep [22] represents dialogue history as a temporal knowledge graph to capture causal and temporal dependencies. Collectively, these systems extend RAG by incorporating notions of forgetting, abstraction, and structure, offering more principled approaches for managing long-term conversational memory.

However, despite recent progress, existing memory systems still face three major challenges in real-world de-

* Equal contribution. † Project lead. # Corresponding authors.

ployment. **First**, maintaining a consistent user profile remains difficult. Current systems rely on dozens of pre-defined fields, while real conversations typically provide only sparse information. This mismatch forces the model to hallucinate unsupported details and leaves many fields empty, creating unnecessary computational overhead and increasing the likelihood of structural errors. **Second**, current memory systems exhibit low write efficiency. Each dialogue turn triggers a retrieval step and sends the top-k memories to an LLM for *add*, *delete*, *update*, or *no-op* decisions, resulting in frequent datastore writes and inefficient API round-trips that degrade throughput and latency. **Third**, multimodal reasoning capabilities remain limited. Most systems are predominantly text-centric and struggle to integrate or reason over memories derived from images, audio, and other modalities, significantly restricting their use in complex real-world settings.

To tackle these challenges, we propose TELEMEM, a unified long-term and multimodal memory system. By extracting only dialogue-supported narrative units, TELEMEM avoids schema-driven hallucinations and enables lightweight, precise updates, resulting in a compact and reliable long-term memory. To further address storage and retrieval inefficiencies, TELEMEM employs a novel writing pipeline that batches summaries, retrieves related memories, clusters semantically similar entries, and applies an LLM-driven decision step before committing them to persistent storage. This process pre-aggregates and deduplicates fragmented information, substantially improving throughput while reducing token usage and storage overhead. In addition, TELEMEM incorporates a multimodal memory module that transforms raw video streams into event and object memories. Combined with a ReAct-style reasoning framework [38], this design enables a closed-loop observe–think–act process for precise reasoning over complex video content. In summary, our contributions are following:

1. We introduce TELEMEM, a unified long-term and multimodal memory framework that leverages narrative dynamic extraction to maintain coherent, hallucination-free user profiles.
2. We design a structured writing pipeline that batches, retrieves, clusters, and consolidates memory entries, substantially improving storage efficiency, reducing token overhead, and accelerating memory operations.
3. We develop a multimodal memory module with ReAct-style reasoning, enabling end-to-end observe–think–act capabilities for complex video understanding.
4. Extensive experiments demonstrate that TELEMEM outperforms the state-of-the-art Mem0 baseline by 19% in accuracy on the ZH-4O benchmark, while reducing token usage by 43% and achieving a $2.1\times$ speedup.

2. TELEMEM System

TELEMEM integrates text-based and multimodal memory into a unified long-term memory system. The text memory module includes profile memory, which captures stable user attributes through lightweight summarization and vector retrieval, and event memory, which models dynamic interaction details using summarization, retrieval, clustering, and consolidation (Sec. 2.1). Complementing these, the multimodal memory module processes video inputs through segmentation, captioning, and cross-modal embedding, and supports video-grounded reasoning via ReAct-style tool use and temporal querying (Sec. 2.2). Together, these components enable TELEMEM to maintain consistent, structured, and query-efficient long-term memory across both linguistic and visual modalities.

2.1. Text Memory Module

Overall, the text memory module centers around two complementary operations: memory writing, which transforms incoming dialogue into structured long-term representations through summarization, retrieval, clustering, and refinement (Sec. 2.1.1); and memory reading, which retrieves and integrates relevant profile and event information to support downstream reasoning (Sec. 2.1.2). Together, these operations enable TELEMEM to continuously accumulate, organize, and utilize long-term textual knowledge in a coherent and query-efficient manner.

Formally, we denote the dialogue at turn t as D_t . We maintain three memory stores: a user profile memory $\mathcal{M}_{\text{user}}$, a bot profile memory \mathcal{M}_{bot} both capturing stable character attributes and relations from their respective perspectives, and an event memory \mathcal{M}_e , which contains dynamic, temporally grounded interaction summaries.

2.1.1. Memory Writing

The profile and event writing mechanisms enable TELEMEM to construct a coherent, dynamically evolving memory space that faithfully reflects both stable role attributes and unfolding narrative developments.

Profile Memory To capture stable, character-centric information, TELEMEM constructs role-specific profile summaries from each dialogue turn D_t . Because a single turn contains contributions from both interlocutors, the system produces two separate summaries p_t^{user} and p_t^{bot} , which are written into the corresponding profile memory stores $\mathcal{M}_{\text{user}}$ and \mathcal{M}_{bot} . Unlike schema-driven methods that rely on pre-defined fields or incremental slot updates, TELEMEM imposes no fixed profile template. Instead, each p_t^{user} and p_t^{bot} is a standalone textual description grounded solely in information explicitly supported by the dialogue. This lightweight formulation avoids schema constraints, prevents unsupported inferences, and keeps profile memory

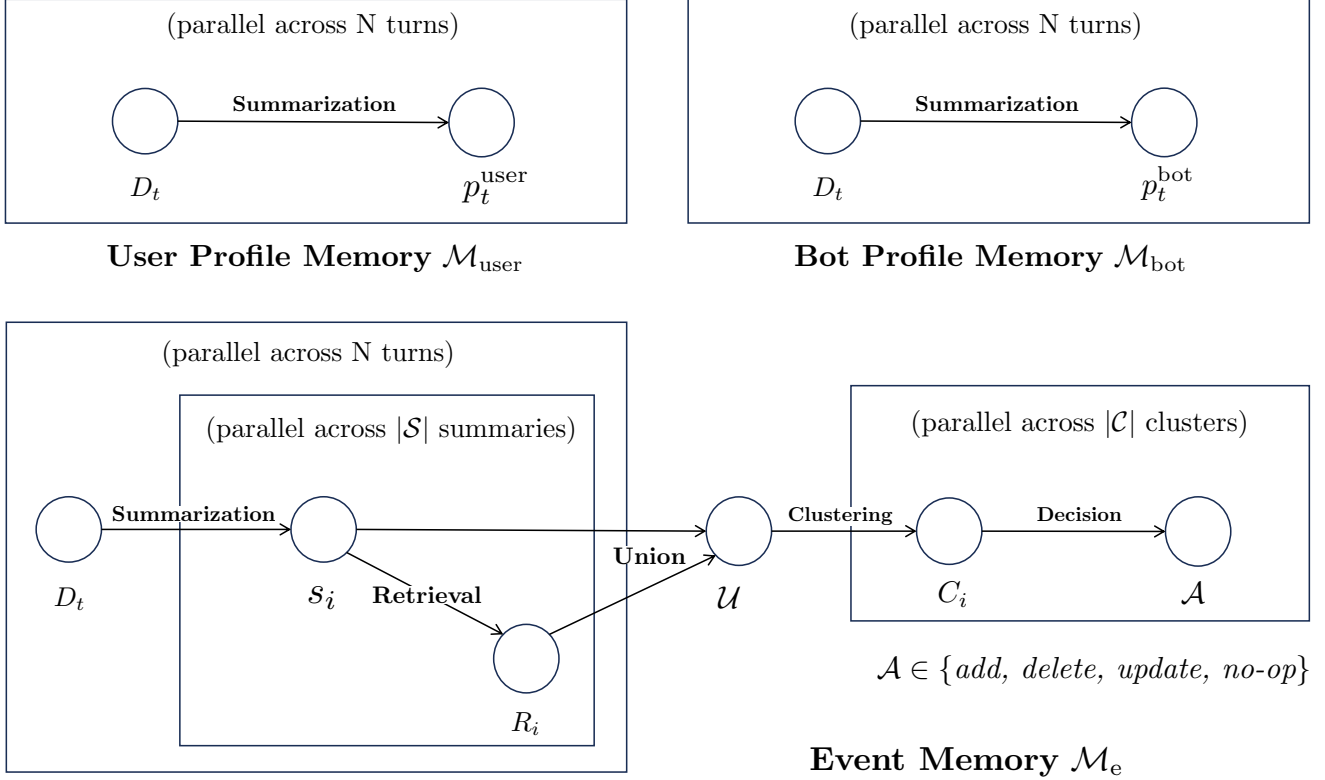


Figure 1. Overview of the text memory writing pipeline in TELEMEM. The system maintains three types of long-term memory: user profile memory $\mathcal{M}_{\text{user}}$, bot profile memory \mathcal{M}_{bot} , and event memory \mathcal{M}_e . 1). For each dialogue turn D_t , the model generates role-specific profile summaries p_t^{user} and p_t^{bot} , which are written into their respective profile memory stores in parallel. 2). Event memory is constructed through a four-stage batch pipeline. Each turn D_t is summarized into one or more textual summaries s_i , forming the set \mathcal{S} . For every summary s_i , the system retrieves the most relevant event memories from \mathcal{M}_e , yielding retrieval sets R_i . All summaries and retrieved items are merged into a unified candidate pool \mathcal{U} , which is globally clustered into semantic groups $C_i \in \mathcal{C}$. For each cluster C_i , entries are chronologically ordered and passed to an LLM that assigns an action $A \in \{\text{add}, \text{delete}, \text{update}, \text{no-op}\}$, producing consolidated event memory updates that are written back into \mathcal{M}_e .

flexible, grounded, and easy to retrieve throughout long-term interactions. Moreover, this role-specific formulation naturally generalizes to settings with additional participants, allowing profile memories to scale seamlessly to multi-party interactions.

```
{
  "summary": "Characters discussed the
↪ upcoming action plan.",
  "sample_id": "session_001",
  "round_index": 3,
  "timestamp": "2024-01-01T00:00:00Z",
  "user": "Jordan"
}
```

Figure 2. Example JSON snippet of a text profile memory entry.

Event Memory Event memory is responsible for capturing the dynamic, interaction-driven information that

emerges throughout a multi-turn dialogue. Different from streaming methods [2] that update memory turn by turn and often incur redundant retrievals and fragmented writes, our batch formulation enables joint processing of multiple dialogue turns, improving coherence, efficiency, and consolidation quality. Assuming that the system receives a batch of dialogue turns $\{D_t\}_{t=1}^N$. To transform these into long-term event memory, TELEMEM applies a four-stage writing pipeline consisting of summarization, retrieval, clustering, and decision (Fig. 1).

- **Summarization (parallel across turns):** Each dialogue turn D_t is independently summarized into a set of one or more textual summaries, denoted as S_t . We denote the collection of all summaries produced from a batch of dialogue turns as

$$\mathcal{S} = \bigcup_{t=1}^N S_t,$$

- **Retrieval (parallel across summaries):** For each sum-

mary $s_i \in \mathcal{S}$, the system performs vector retrieval from the memory store \mathcal{M}_e , obtaining the top- k related memories:

$$R_i = \text{Top-}k(\text{sim}(s_i, \mathcal{M}_e)).$$

- **Clustering (global, non-parallel):** A unified candidate set is constructed as

$$\mathcal{U} = \left(\bigcup_{i=1}^N R_i \right) \cup \mathcal{S},$$

upon which global semantic clustering is applied to obtain clusters

$$\mathcal{C} = \{C_1, C_2, \dots, C_m\}, \quad C_i \subseteq \mathcal{U},$$

where each cluster groups entries describing semantically related events or topics.

- **Decision (parallel across clusters):** For each cluster C_i , the system orders all entries temporally and feeds them to an LLM, which determines for each item whether to perform an action $\mathcal{A} \in \{\text{add}, \text{delete}, \text{update}, \text{no-op}\}$. This process resolves redundancies and inconsistencies and produces refined memory entries that are subsequently re-embedded and written back into \mathcal{M}_e .

```
{
  "summary": "Jordan took the subway to
  ↳ work again today, believing that
  ↳ it's faster than driving. He left
  ↳ at 7 o'clock, and the subway wasn't
  ↳ crowded.",
  "sample_id": "session_001",
  "round_index": 0,
  "timestamp":
  ↳ "2025-12-05T03:18:10.423379+00:00",
  "original_messages": [
    {
      "role": "user",
      "content": "Jordan, did you
      ↳ take the subway to work
      ↳ again today?"
    },
    {
      "role": "assistant",
      "content": "Yes, James. The
      ↳ subway is much faster than
      ↳ driving. I leave at 7
      ↳ o'clock and it's just not
      ↳ crowded."
    }
  ]
}
```

Figure 3. Example JSON snippet of a text event memory entry.

2.1.2. Memory Reading

Given a query q , TELEMEM retrieves relevant information from three memory stores: the user profile memory $\mathcal{M}_{\text{user}}$,

the bot profile memory \mathcal{M}_{bot} , and the event memory \mathcal{M}_e . For each store, the embedding of q is used to retrieve the top 5 most similar entries. These retrieved items provide stable user attributes, stable bot attributes, and dynamic interaction context. They are then inserted into the LLM prompt as external memory, enabling retrieval-augmented reasoning that remains consistent with long-term roles and past events.

2.2. Multimodal Memory Module

Beyond long-term textual inputs, we extend the memory architecture to support multimodal video memory. Our design targets two key requirements of long-horizon video understanding: (1) retaining fine-grained visual grounding while (2) enabling efficient retrieval and reasoning over long sequences. To this end, we build a multimodal memory module that captures both event-level semantics and entity-level information (Sec. 2.2.1), and supports flexible, tool-driven access through an agentic reading mechanism (Sec. 2.2.2).

2.2.1. Memory Writing

We first divide the video stream into 10-second clips and uniformly sample two frames from each clip as the basis for memory extraction. The writing pipeline produces two complementary memory types: event memory and key-value object memory.

Event Memory For each clip, a VLM generates a concise caption using the sampled frames. These clip-level event memories capture localized semantic events and remain fixed once written, forming a chronological record of the video \mathcal{M}_e .¹

```
{
  "caption": "A boy is sitting on a bench
  ↳ watching the sea. Words appear on
  ↳ screen reading 'A report by Luca
  ↳ Schmitt-Walz.'"
}
```

Figure 4. Example JSON snippet of a video event memory entry.

Key-Value Object Memory For each clip, a VLM generates person and object entities and store them in a key-value structure. Each key specifies stable identity attributes including name, appearance, identity cues, and first occurrence, while each value captures the entity’s behavior or state in the clip. To respect the LLM’s context length constraints, we avoid consolidating all object entries at once. Instead, raw key-value records are accumulated and updated in batches by a text LLM, which incrementally

¹We slightly overload the notation here for simplicity of presentation.

merges redundant information and produces temporally coherent entity summaries. After all batches are processed, we obtain the final consolidated object memory \mathcal{M}_{obj} .

```
{
  "name": "Unidentified young man",
  "appearance": [
    "wearing a red sleeveless jersey",
    "has dark hair",
    "wear floppy shades"
  ],
  "identity": ["preparing to tie a bun"],
  "first_seen": "00:01:00"
}
```

Figure 5. Example JSON snippet of an object memory entry.

2.2.2. Memory Reading

We adopt a ReAct-style agent to read from the multimodal memory, where the agent iteratively selects among three tools to retrieve, refine, or verify information. These tools operate over the text-based memory stores (\mathcal{M}_e and \mathcal{M}_{obj}) or the original video clips (Alg. 1).

video.retrieval This tool performs vector-based retrieval between a textual query and the text-based memory store (\mathcal{M}_e and \mathcal{M}_{obj}). It returns the clip indices, specifically the start and end timestamps, that correspond to the most relevant memory entries. This allows the agent to efficiently localize the video segments associated with the query.

video.rag This tool also conducts vector retrieval against the text-based memory store, but instead of returning clip boundaries, it selects the top- k most relevant memory entries and feeds them into the LLM prompt. The LLM then synthesizes an aggregated, query-grounded response using these retrieved memory elements.

video.qa This tool performs vision-language question answering over a specified video clip. After the agent identifies a relevant clip, VQA is used to extract fine-grained visual details such as appearance attributes, spatial relations, or object states that may not be available in textual memory.

3. Experiments

Dataset We conduct experiments on two ultra-long dialogue datasets to assess agent memory capabilities. ZH-40 [1] serves as a Chinese role-playing benchmark comprising 28 authentic human-LLM sessions (avg. 600 turns), annotated with 1,068 multiple-choice probing questions

Algorithm 1: Multimodal Memory Reading.

Input: Initial query q , $\text{LLM}(\cdot)$, $\mathcal{M} = \mathcal{M}_e \cup \mathcal{M}_{\text{obj}}$
Output: Response

- 1 Initialize history $H = [q]$;
- 2 **while** not exceed max iterations **do**
- 3 action, args = $\text{LLM}(H)$;
- 4 **if** action is *video.retriever* **then**
- 5 $q \leftarrow$ args ;
- 6 results \leftarrow video.retriever(q, \mathcal{M});
- 7 **else if** action is *video.rag* **then**
- 8 $q \leftarrow$ args ;
- 9 results \leftarrow video.rag(q, \mathcal{M});
- 10 **else if** action is *video.qa* **then**
- 11 $q, t_{\text{start}}, t_{\text{end}} \leftarrow$ args ;
- 12 results \leftarrow video.qa($q, t_{\text{start}}, t_{\text{end}}$);
- 13 **else if** action is *finish* **then**
- 14 break;
- 15 $H = H + [(action, args, results)]$;
- 16 **return** $\text{LLM}(H)$.

to test memory recall. Complementarily, we utilize LoCoMo [16] as an English dataset, which consists of 10 long-context sessions paired with 1,540 question-answer pairs covering single-hop, multi-hop, open-domain, and temporal reasoning. Notably, we exclude the LoCoMo adversarial subset due to the absence of ground-truth answers.

Baselines We compare our framework against five baselines, categorized into general paradigms and specialized memory systems. First, we consider two foundational approaches: **Long context LLM** utilizes the entire conversation history to establish a full-context reference, while **RAG** retrieves the top- k semantically relevant segments to augment generation. Second, we include three advanced memory architectures: **Memobase** [24] maintains structured user profiles and event memories; **A-Mem** [35] utilizes an agentic framework to autonomously link memory notes; and **Mem0** [2], a modular memory system designed for scalable deployment with explicit in-context memory operations.

Implementation To ensure a fair and consistent evaluation, we re-implemented all baselines using **Qwen3-8B** [37] (configured in “no-think” mode) as the backbone LLM, paired with **Qwen3-8B-embedding** [37] for vector representations. Crucially, a unified prompt template is applied for response generation across all settings to strictly control for variance.

Method	Overall (%)
RAG	62.45
Mem0	70.20
MOOM	72.60
A-Mem	73.78
Memobase	76.78
Long context LLM	84.92
TELEMEM	86.33

Table 1. Performance comparison on the ZH-4O benchmark. The metric represents QA Accuracy (%) across 1,068 probing questions. The best performance is highlighted in **bold**.

Evaluation Our evaluation strategies are tailored to each dataset: for ZH-4O [1], we quantify memory fidelity via QA Accuracy, where the model selects the single correct option for each multiple-choice query against ground-truth labels. For LoCoMo [16], following prior work [2], we employ an LLM-as-a-Judge (J) approach using **GPT-4o** [5], which assesses the factual accuracy, relevance, and completeness of the generated responses.

3.1. Results on ZH-4O

Main Results Table 1 presents the performance of different memory paradigms on the ZH-4O benchmark, revealing clear distinctions in their ability to support long-horizon, multi-turn question answering. We make the following observations.

- **Retrieval-only methods are insufficient.** RAG achieves the lowest accuracy (62.45%), indicating that flat semantic retrieval without temporal ordering or relational structure fails to support multi-turn, memory-intensive reasoning required by ZH-4O.
- **Explicit memory mechanisms consistently improve performance.** All memory-augmented approaches outperform RAG, demonstrating the necessity of maintaining persistent and updatable memory states for long-horizon dialogue understanding. Among them, Memobase performs best (76.78%), likely due to its structured user profiling that aligns well with the role-playing characteristics of the benchmark.
- **Long-context modeling alone has inherent limitations.** The Long Context LLM baseline achieves 84.92% accuracy by leveraging the full dialogue history, but its reliance on raw context makes it susceptible to noise, redundancy, and attention dilution as interactions grow longer.
- **Coordinated read-write memory yields the best results.** Our proposed TELEMEM attains the highest accuracy of **86.33%**, outperforming both long-context and prior memory-based architectures. This demonstrates that selectively compressing salient information and enabling

context-aware memory access is more effective than unstructured context accumulation or loosely coupled memory modules.

Read-Write Scaling Law Table 2 reveals a clear memory read-write scaling law by varying the model sizes of the memory write LLM and the memory read LLM. Overall performance improves monotonically as either component scales up, indicating that both memory writing and memory reading contribute substantially to downstream QA accuracy. However, the gains are not symmetric across the two dimensions. Increasing the size of the write LLM leads to consistent improvements across almost all read settings, suggesting that stronger writers produce more informative, compact, and robust memory representations that benefit readers of different capacities. Scaling the read LLM yields even more pronounced gains, particularly when paired with medium-to-large write models. This indicates that memory retrieval and reasoning capacity plays a critical role in exploiting stored information, especially when the memory content is sufficiently well-structured. Notably, performance saturates when the reader significantly outpaces the writer, implying that retrieval capacity alone cannot compensate for low-quality memory writing. The best results are achieved when both the write and read LLMs are scaled jointly, highlighting the complementary and interdependent nature of memory writing and reading. These findings suggest that effective long-term memory systems should be designed with coordinated read-write capacity rather than over-optimizing either component in isolation.

3.2. Results on LoCoMo

Main Results Table 3 compares different methods on the LoCoMo benchmark across four reasoning categories. Overall, the Long Context LLM achieves the strongest performance (70.71%), demonstrating the effectiveness of processing full interaction histories for complex memory-centric tasks. This advantage is particularly pronounced on temporal questions, where access to complete chronological context is crucial. In contrast, retrieval-based baselines perform poorly across all categories, highlighting the limitations of static semantic retrieval for long-term conversational reasoning. Memory systems exhibit diverse strengths across reasoning types. Mem0 performs well on single-hop questions (54.96%) and temporal reasoning (60.28%), suggesting that its modular memory abstraction is effective for fact recall and time-sensitive information. Memobase excels on multi-hop reasoning (66.04%), likely benefiting from its structured memory representations that better support compositional inference. A-Mem shows more balanced but moderate performance, indicating that agentic memory control alone is insufficient without strong memory structuring. TELEMEM achieves competitive performance

Write ↓	Read →	0.6B	1.7B	4B	8B	14B	32B	Avg.
0.6B		52.72	64.23	73.69	75.47	74.06	72.19	68.73
1.7B		61.14	71.54	79.12	78.56	80.81	77.72	74.82
4B		63.48	74.25	81.09	83.33	84.55	83.33	78.34
8B		65.54	77.15	84.36	86.33	85.77	84.93	80.68
14B		67.04	78.28	84.55	86.33	87.55	85.49	81.54
32B		68.91	80.81	83.99	86.70	86.61	85.96	82.16
Avg.		63.14	74.38	81.13	82.79	83.17	81.60	—

Table 2. Memory read–write scaling law. Performance (QA Accuracy, %) when using different model sizes for the memory write LLM (rows) and the memory read LLM (columns). Results show consistent performance gains as either component scales up, with larger read–write combinations yielding the strongest performance.

on single-hop (64.53%) and temporal reasoning (78.47%), approaching the Long Context LLM in these categories, which suggests that its memory abstraction effectively captures salient factual and temporal information. However, its weaker performance on multi-hop questions indicates remaining challenges in supporting complex relational reasoning over stored memories. Overall, these results highlight that different memory mechanisms favor different reasoning patterns, and no single approach uniformly dominates all categories, underscoring the importance of task-aware memory design for long-horizon QA.

4. Related Work

Text LTM Recent research on equipping Large Language Models (LLMs) with text long-term memory has primarily evolved through system-level abstractions, cognitive-inspired mechanisms, and structured agentic frameworks. To transcend fixed context windows, operating system paradigms like MemGPT [20] and MemOS [8] leverage virtual memory abstractions and hierarchical scheduling to orchestrate data flow between active context and external storage. Complementing these architectural innovations, bio-inspired models such as MemoryBank [45], Light-Mem [3], and Nemori [19] integrate human cognitive theories—ranging from the Ebbinghaus forgetting curve to the Atkinson-Shiffrin model—to dynamically prioritize, segment, and decay information for efficient retrieval. On the structural front, Mem0 [2], A-Mem [35], and MIRIX [30] propose modular systems that enable self-evolving knowledge consolidation and multi-agent coordination.

Graph LTM Recent literature on long-term memory for LLM agents highlights a transition from unstructured storage to graph-based architectures that enhance reasoning, consistency, and scalability. Foundational works like *Mem0^g* [2] and Zep [22] establish production-ready and temporal knowledge graph systems to manage agent mem-

ory efficiently. To address the complexity of long-term interactions, G-Memory [41], LiCoMemory [4], and SG-MEM [32] proposes hierarchical graph structures that organize information at varying granularities, from sentence-level details to high-level insights, thereby facilitating self-evolution and reducing fragmentation. Enhancing the reasoning capabilities over these structures, GraphCogent [26] and D-SMART [6] integrate working memory models and dynamic reasoning trees to support complex graph understanding and dialogue consistency, MemQ [34] focuses on optimizing knowledge graph reasoning. Furthermore, [33] introduces trainable graph memories that abstract agent trajectories into strategic meta-cognition. Collectively, these studies demonstrate that structured, hierarchical, and dynamic graph memories are essential for developing autonomous agents capable of coherent, long-term strategic planning and reasoning.

Parametric LTM Parametric memory aims to encode knowledge or contextual information directly into the weights or persistent hidden states of neural networks [12, 14, 18], distinguishing itself from non-parametric approaches that rely on external vector databases. Wang et al. [27] propose a decoupled memory mechanism, employing a residual side network to cache long-term context while keeping the base LLM frozen. Similarly, MLPMemory [31] internalizes retrieval by training multilayer perceptrons to approximate k-nearest neighbor distributions as differentiable mappings. Pushing the internalization of retrieval further, Tay et al. [23] introduce the Differentiable Search Index, which eliminates external indices by training the model to map queries directly to document identifiers via its parameters. regarding personalized adaptation, Zhang et al. [42] combine LoRA with Bayesian optimization to inject dialogue history into model weights. Complementary to these injection methods, Meng et al. [17] explore the interpretability of parametric storage, proposing ROME to locate and directly edit factual associations within the Trans-

Method	Single-Hop	Multi-Hop	Open Domain	Temporal	Overall
RAG	20.91	32.39	35.41	48.03	39.03
Mem0	54.96	31.15	40.62	60.28	52.01
A-Mem	44.32	35.82	33.33	56.59	48.57
Memobase	50.70	66.04	37.50	58.97	57.59
Long context LLM	64.89	47.66	42.70	84.66	70.71
TELEMEM	64.53	20.56	40.62	78.47	61.49

Table 3. Performance comparison on the LoCoMo benchmark. All metrics represent accuracy (%) evaluated via LLM-as-a-Judge across 1,540 question-answer pairs.

former’s MLP layers. Overall, parametric memory methods offer scalability and unified reasoning but face challenges regarding update costs and potential misalignment with the base model.

Multimodal LTM Multimodal memory [28, 43] has drawn increasing attention as modern agents must store and reason over long-horizon visual and textual information. Mem0 [2] introduces a multimodal image interface but ultimately reduces visual inputs to captions and continues to operate purely in the textual space. M3-Agent [15] extends memory into the multimodal domain by maintaining entity-level representations across audio–visual streams and enabling agentic retrieval over long video sequences. Inspired by hippocampal mechanisms, HippoMM [9] proposes cross-modal event encoding and temporal consolidation to support richer multimodal recall. MemVerse [10] further explores lifelong multimodal memory through a hierarchical retrieval framework with periodic distillation for compactness. In contrast, VisMem [39] emphasizes preserving visual latent representations rather than relying solely on textual abstractions. Overall, existing multimodal memory systems highlight the need for multi-granular representations and adaptive retrieval, while our approach advances this direction with a lightweight hierarchical design tailored for continuous multimodal streams.

5. Conclusion

We introduced TELEMEM, a unified long-term and multimodal memory system that overcomes key limitations of existing RAG-based approaches. By extracting narrative-grounded information and employing a structured writing pipeline for batching, clustering, and consolidation, TELEMEM maintains coherent user profiles while greatly improving storage and token efficiency. Its multimodal memory module with ReAct-style reasoning further enables accurate observe–think–act processing for complex video content. Experiments on the ZH-4O benchmark show that TELEMEM substantially outperforms Mem0 in accuracy,

efficiency, and speed, underscoring its effectiveness.

Acknowledgments. The development of TELEMEM has been greatly influenced by the contributions of open-source communities and prior research efforts. We express our sincere appreciation to the teams and projects whose work has inspired and supported this system, including Mem0, Memobase, MOOM, DVD, and Memento. Their innovations have provided valuable foundations upon which this work builds.

References

- [1] Weishu Chen, Jinyi Tang, Zhouhui Hou, Shihao Han, Mingjie Zhan, Zhiyuan Huang, Delong Liu, Jiawei Guo, Zhicheng Zhao, and Fei Su. Moom: Maintenance, organization and optimization of memory in ultra-long role-playing dialogues. *arXiv preprint arXiv:2509.11860*, 2025. 5, 6
- [2] Prateek Chhikara, Dev Khant, Saket Aryan, Taranjeet Singh, and Deshraj Yadav. Mem0: Building production-ready ai agents with scalable long-term memory. *arXiv preprint arXiv:2504.19413*, 2025. 1, 3, 5, 6, 7, 8
- [3] Jizhan Fang, Xinle Deng, Haoming Xu, Ziyang Jiang, Yuqi Tang, Ziwen Xu, Shumin Deng, Yunzhi Yao, Mengru Wang, Shuofei Qiao, et al. Lightmem: Lightweight and efficient memory-augmented generation. *arXiv preprint arXiv:2510.18866*, 2025. 7
- [4] Zhengjun Huang, Zhoujin Tian, Qintian Guo, Fangyuan Zhang, Yingli Zhou, Di Jiang, and Xiaofang Zhou. Licomemory: Lightweight and cognitive agentic memory for efficient long-term reasoning. *arXiv preprint arXiv:2511.01448*, 2025. 7
- [5] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024. 1, 6
- [6] Xiang Lei, Qin Li, and Min Zhang. D-smart: Enhancing llm dialogue consistency via dynamic structured memory and reasoning tree. *arXiv preprint arXiv:2510.13363*, 2025. 7
- [7] Zhuowan Li, Cheng Li, Mingyang Zhang, Qiaozhu Mei, and Michael Bendersky. Retrieval augmented generation or long-context llms? a comprehensive study and hybrid approach.

- In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 881–893, 2024. 1
- [8] Zhiyu Li, Shichao Song, Chenyang Xi, Hanyu Wang, Chen Tang, Simin Niu, Ding Chen, Jiawei Yang, Chunyu Li, Qingchen Yu, et al. Memos: A memory os for ai system. *arXiv preprint arXiv:2507.03724*, 2025. 7
- [9] Yueqian Lin, Qinsi Wang, Hancheng Ye, Yuzhe Fu, Hai Li, Yiran Chen, et al. Hippomm: Hippocampal-inspired multimodal memory for long audiovisual event understanding. *arXiv preprint arXiv:2504.10739*, 2025. 8
- [10] Junming Liu, Yifei Sun, Weihua Cheng, Haodong Lei, Yirong Chen, Licheng Wen, Xuemeng Yang, Daocheng Fu, Pinlong Cai, Nianchen Deng, et al. Memverse: Multimodal memory for lifelong learning agents. *arXiv preprint arXiv:2512.03627*, 2025. 8
- [11] Jiaheng Liu, Dawei Zhu, Zhiqi Bai, Yancheng He, Huanxuan Liao, Haoran Que, Zekun Wang, Chenchen Zhang, Ge Zhang, Jiebin Zhang, et al. A comprehensive survey on long context language modeling. URL <https://api.semanticscholar.org/CorpusID/277271533>, 2025. 1
- [12] Nuowei Liu, Jiahao Kuang, Yanting Liu, Tao Ji, Changzhi Sun, Man Lan, and Yuanbin Wu. Protein design with dynamic protein vocabulary. *arXiv preprint arXiv:2505.18966*, 2025. 7
- [13] Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12: 157–173, 2024. 1
- [14] Yanting Liu, Tao Ji, Changzhi Sun, Yuanbin Wu, and Xiaoling Wang. Generation with dynamic vocabulary. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 18931–18948, 2024. 7
- [15] Lin Long, Yichen He, Wentao Ye, Yiyuan Pan, Yuan Lin, Hang Li, Junbo Zhao, and Wei Li. Seeing, listening, remembering, and reasoning: A multimodal agent with long-term memory. *arXiv preprint arXiv:2508.09736*, 2025. 8
- [16] Adyasha Maharana, Dong-Ho Lee, Sergey Tulyakov, Mohit Bansal, Francesco Barbieri, and Yuwei Fang. Evaluating very long-term conversational memory of llm agents. *arXiv preprint arXiv:2402.17753*, 2024. 5, 6
- [17] Kevin Meng, David Bau, Alex Andonian, and Yonatan Beilinkov. Locating and editing factual associations in gpt. *Advances in neural information processing systems*, 35:17359–17372, 2022. 7
- [18] Daiye Miao, Yufang Liu, Jie Wang, Changzhi Sun, Yunke Zhang, Demei Yan, Shaokang Dong, Qi Zhang, and Yuanbin Wu. Taso: Task-aligned sparse optimization for parameter-efficient model adaptation. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 22746–22758, 2025. 7
- [19] Jiayan Nan, Wenquan Ma, Wenlong Wu, and Yize Chen. Nemori: Self-organizing agent memory inspired by cognitive science. *arXiv preprint arXiv:2508.03341*, 2025. 7
- [20] Charles Packer, Vivian Fang, Shishir-G Patil, Kevin Lin, Sarah Wooders, and Joseph E Gonzalez. Memgpt: Towards llms as operating systems. 2023. 7
- [21] Zhuoshi Pan, Qianhui Wu, Huiqiang Jiang, Xufang Luo, Hao Cheng, Dongsheng Li, Yuqing Yang, Chin-Yew Lin, H Vicky Zhao, Lili Qiu, et al. Secom: On memory construction and retrieval for personalized conversational agents. In *The Thirteenth International Conference on Learning Representations*, 2025. 1
- [22] Preston Rasmussen, Pavlo Paliychuk, Travis Beauvais, Jack Ryan, and Daniel Chalef. Zep: a temporal knowledge graph architecture for agent memory. *arXiv preprint arXiv:2501.13956*, 2025. 1, 7
- [23] Yi Tay, Vinh Tran, Mostafa Dehghani, Jianmo Ni, Dara Bahri, Harsh Mehta, Zhen Qin, Kai Hui, Zhe Zhao, Jai Gupta, et al. Transformer memory as a differentiable search index. *Advances in Neural Information Processing Systems*, 35:21831–21843, 2022. 7
- [24] Memobase Development Team. Beyond rag: Memobase - a user profile-based long-term memory system. Online Blog & GitHub Repository, 2025. 5
- [25] Qingyue Wang, Yanhe Fu, Yanan Cao, Shuai Wang, Zhiliang Tian, and Liang Ding. Recursively summarizing enables long-term dialogue memory in large language models. *Neurocomputing*, 639:130193, 2025. 1
- [26] Rongzheng Wang, Shuang Liang, Qizhi Chen, Yihong Huang, Muquan Li, Yizhuo Ma, Dongyang Zhang, Ke Qin, and Man-Fai Leung. Graphcogent: Mitigating llms’ working memory constraints via multi-agent collaboration in complex graph understanding. *arXiv preprint arXiv:2508.12379*, 2025. 7
- [27] Weizhi Wang, Li Dong, Hao Cheng, Xiaodong Liu, Xifeng Yan, Jianfeng Gao, and Furu Wei. Augmenting language models with long-term memory. *Advances in Neural Information Processing Systems*, 36:74530–74543, 2023. 7
- [28] Xiaohan Wang, Yuhui Zhang, Orr Zohar, and Serena Yeung-Levy. Videoagent: Long-form video understanding with large language model as agent. In *European Conference on Computer Vision*, pages 58–76. Springer, 2024. 8
- [29] Xinyu Wang, Changzhi Sun, Lian Cheng, Yuanbin Wu, Dell Zhang, Xiaoling Wang, and Xuelong Li. Logic-regularized verifier elicits reasoning from llms. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 32617–32630, 2025. 1
- [30] Yu Wang and Xi Chen. Mirix: Multi-agent memory system for llm-based agents. *arXiv preprint arXiv:2507.07957*, 2025. 7
- [31] Rubin Wei, Jiaqi Cao, Jiarui Wang, Jushi Kai, Qipeng Guo, Bowen Zhou, and Zhouhan Lin. Mlp memory: A retriever-pretrained memory for large language models. *arXiv preprint arXiv:2508.01832*, 2025. 7
- [32] Yaxiong Wu, Yongyue Zhang, Sheng Liang, and Yong Liu. Sgmem: Sentence graph memory for long-term conversational agents. *arXiv preprint arXiv:2509.21212*, 2025. 7
- [33] Siyu Xia, Zekun Xu, Jiajun Chai, Wentian Fan, Yan Song, Xiaohan Wang, Guojun Yin, Wei Lin, Haifeng Zhang, and Jun Wang. From experience to strategy: Empowering llm agents with trainable graph memory. *arXiv preprint arXiv:2511.07800*, 2025. 7

- [34] Mufan Xu, Gewen Liang, Kehai Chen, Wei Wang, Xun Zhou, Muyun Yang, Tiejun Zhao, and Min Zhang. Memory-augmented query reconstruction for llm-based knowledge graph reasoning. *arXiv preprint arXiv:2503.05193*, 2025. 7
- [35] Wujiang Xu, Zujie Liang, Kai Mei, Hang Gao, Juntao Tan, and Yongfeng Zhang. A-mem: Agentic memory for llm agents. *arXiv preprint arXiv:2502.12110*, 2025. 5, 7
- [36] Xiwei Xu, Robert Mao, Quan Bai, Xuewu Gu, Yechao Li, and Liming Zhu. Everything is context: Agentic file system abstraction for context engineering, 2025. 1
- [37] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025. 1, 5
- [38] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *The eleventh international conference on learning representations*, 2022. 2
- [39] Xinlei Yu, Chengming Xu, Guibin Zhang, Zhangquan Chen, Yudong Zhang, Yongbo He, Peng-Tao Jiang, Jiangning Zhang, Xiaobin Hu, and Shuicheng Yan. Vismem: Latent vision memory unlocks potential of vision-language models. *arXiv preprint arXiv:2511.11007*, 2025. 8
- [40] Dell Zhang, Yue Feng, Haiming Liu, Changzhi Sun, Jixiang Luo, Xiangyu Chen, and Xuelong Li. Conversational agents: From rag to ltm. In *Proceedings of the 2025 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, pages 447–452, 2025. 1
- [41] Guibin Zhang, Muxin Fu, Guancheng Wan, Miao Yu, Kun Wang, and Shuicheng Yan. G-memory: Tracing hierarchical memory for multi-agent systems. *arXiv preprint arXiv:2506.07398*, 2025. 7
- [42] Kai Zhang, Yejin Kim, and Xiaozhong Liu. Personalized llm response generation with parameterized memory injection. *arXiv preprint arXiv:2404.03565*, 2024. 7
- [43] Xiaoyi Zhang, Zhaoyang Jia, Zongyu Guo, Jiahao Li, Bin Li, Houqiang Li, and Yan Lu. Deep video discovery: Agentic search with tool use for long-form video understanding. *arXiv preprint arXiv:2505.18079*, 2025. 8
- [44] Qingfei Zhao, Ruobing Wang, Yukuo Cen, Daren Zha, Shicheng Tan, Yuxiao Dong, and Jie Tang. Longrag: A dual-perspective retrieval-augmented generation paradigm for long-context question answering. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 22600–22632, 2024. 1
- [45] Wanjun Zhong, Lianghong Guo, Qiqi Gao, He Ye, and Yanlin Wang. Memorybank: Enhancing large language models with long-term memory. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 19724–19731, 2024. 1, 7