MemCoRe: Recovering Evidence from Progressively Compressed Factual Knowledge for Agent Memory

Zhenyuan Zhang[1*] Xianzhang Jia[1*] Zhiqin Yang[1] Zhenbo Song[2] Wei Xue[1†] Sirui Han[1†] Yike Guo[1†]

[1] The Hong Kong University of Science and Technology [2] Nanjing University of Science and Technology

* Equal contribution † Corresponding authors

NEURAL INFORMATION PROCESSING SYSTEMS

NeurIPS 2026, Sydney

Paper Code WeChat

MOTIVATION

(a) Stored evidence can be missed

A-MEM · LoCoMo · 1,986 questions

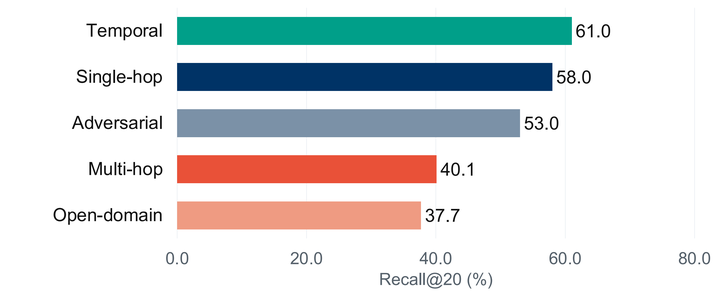

→ 53.8% question-weighted recall.
Multi-hop and open-domain queries show the lowest recall.

(b) Distractors dilute useful evidence

Qwen3-8B · 1,531 questions with gold evidence

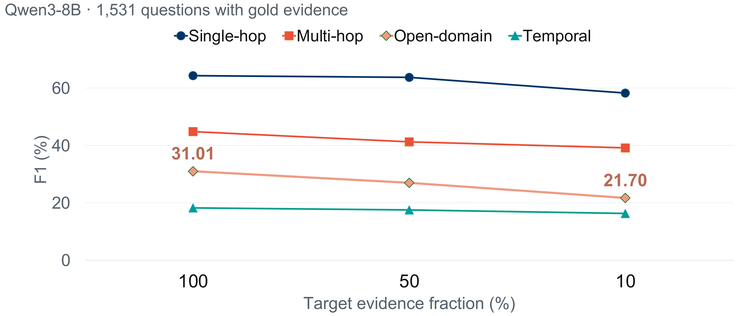

→ 30.0% relative F1 drop on open-domain.
Gold evidence remains in the context.

(c) Design rationale

Consolidation with evidence retention

Consolidate redundant contexts while retaining raw observations.

Retrieval across memory levels

Queries differ in their entity, thematic, and relational cues.
Keywords, topics, and links provide complementary access to evidence.

PROPOSED METHOD

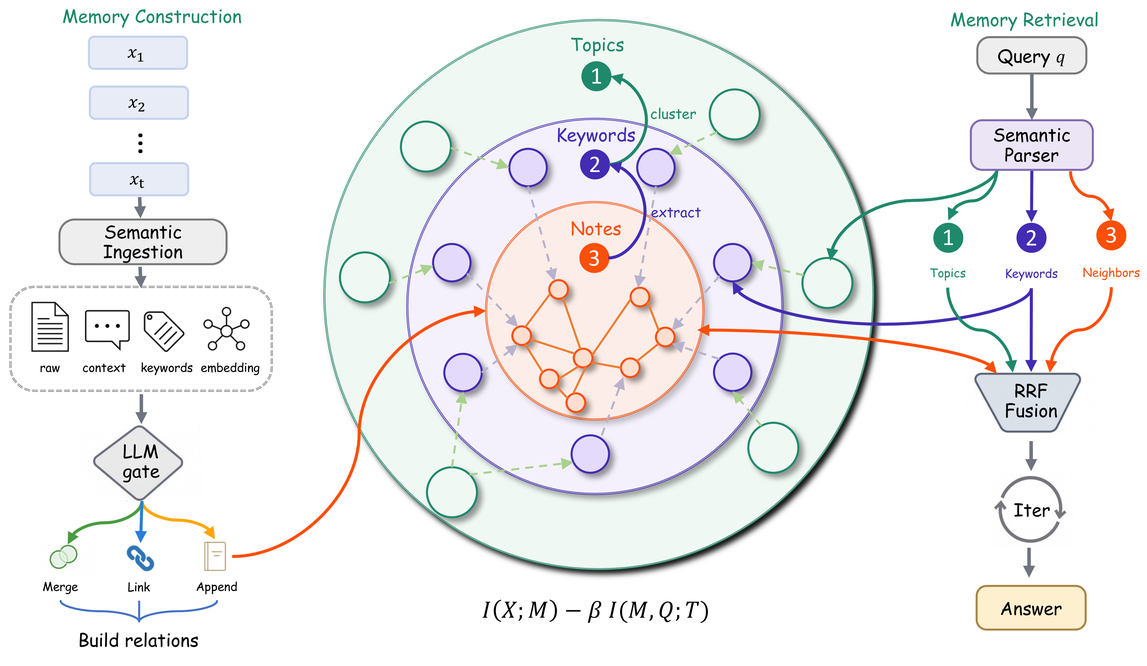

Construction consolidates contexts and builds keywords, topics, and links.
Retrieval uses these structures to recover evidence and iteratively update the evidence pool.

EXPERIMENTS

LoCoMo F1 (%)

1,986 questions, question-weighted

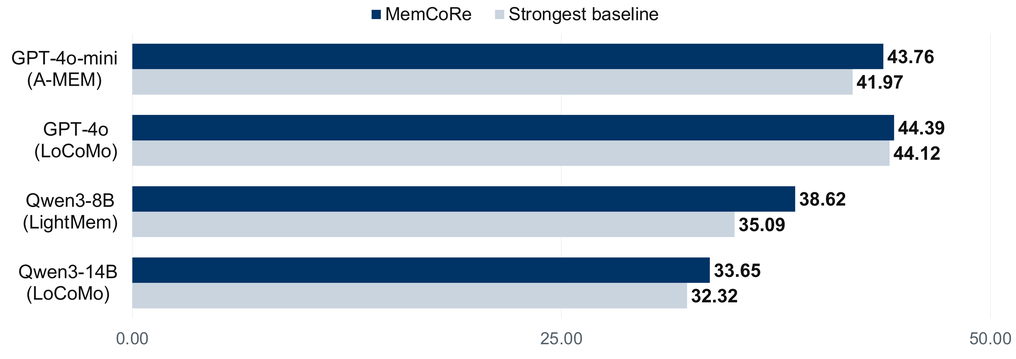

Highest reported F1 across four backbones.
Statistical evidence is clearest on Qwen3-8B.

Component ablations

LoCoMo, Qwen3-8B, F1 (%)

MemCoRe	**38.62**
No gated update	27.97
No Link	33.57
No refinement	32.94

LongMemEval-S

500 questions, Qwen3-8B
Judge accuracy (%)

LightMem	59.6
MemCoRe	**64.2**

zyzhang.bbetter@gmail.com

HKUST authors were supported by the Hong Kong Generative AI Research & Development Center. No competing interests.