KEEP IT INMIND: BENCHMARKING THE IMPLICIT-ASSOCIATION BLIND SPOT IN AGENT MEMORY

Anonymous authors
Paper under double-blind review

ABSTRACT

Long-term memory systems typically retrieve records similar to a new query. This assumes that a needed memory will resemble the query that needs it, but world knowledge can make dissimilar texts consequentially related: a tree-nut allergy should change an answer about macarons through their almond-flour ingredient. We call this failure mode the *implicit-association blind spot* and introduce InMind, a 125-task, expert-verified benchmark with paired controls that distinguish storage failure, missing bridging knowledge, and retrieval failure. With the decisive memory in context, GPT-5-mini gives the required reminder on 85.6% of indirect queries. Six vector, graph, and agentic memory systems reach at most 8.0% when success requires both recalling the fact and giving the reminder, despite reaching up to 100% direct-query accuracy. A larger embedding improves recall but leaves the gap intact. Conversely, a diagnostic probe that makes memory visible before the query arrives reaches 64.8%, locating the failure in query-conditioned retrieval and motivating *routing*: deciding which facts must remain visible.

1 INTRODUCTION

Ask a memory-augmented agent what its user is allergic to, and it answers correctly: tree nuts. Ask it next for a macaron recipe, and it recommends almond flour. Figure 1 shows the exchange; Appendix D records it verbatim. The fact was not forgotten—it was recalled seconds earlier—but was absent when it mattered.

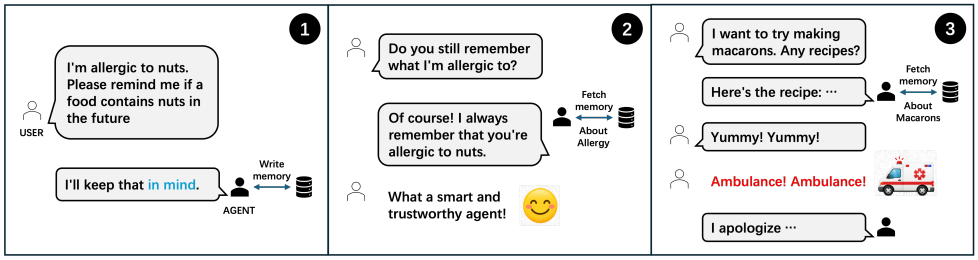

Figure 1: Success on a direct query does not guarantee application: the system recalls a tree-nut allergy, then recommends almond-based macarons (King Arthur Baking Company, n.d.; U.S. Food and Drug Administration, n.d.).

Persistent assistants must let a fact stated once affect later decisions across conversations (Zhang et al., 2024; Packer et al., 2023; Pan et al., 2025). Most do so through retrieval: user information is written to an external store; a new query retrieves similar records and places them in context (Lewis et al., 2020; Karpukhin et al., 2020; Borgeaud et al., 2022; Asai et al., 2023). This scales and works well when the query itself supplies a cue, as "what am I allergic to?" does.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

The difficulty arises when a memory becomes relevant only after applying world knowledge. A macaron request and a tree-nut-allergy memory are connected by two facts: traditional macarons use almond flour, and almonds are tree nuts. Neither fact is stated in the query or the memory itself. Yet query-conditioned memory systems retrieve before the LLM can reason over the stored information: the retriever must first decide which memories to expose, and the model can apply the bridge only to memories that survive this decision. This ordering relies on what we call the *retrieval hypothesis*: every memory needed for a query can be recovered from that query alone (Section 2.1). Drug–food interactions, dietary restrictions, and household hazards routinely violate this assumption. Indeed, HippoRAG 2, A-Mem, xMemory, and A-RAG all fail the example in Figure 1 (Gutierrez et al., 2025; Xu et al., 2025; Hu et al., 2026; Du et al., 2026).

The blind spot is doubly hidden. Users calibrate trust by asking the agent what it remembers, and that is precisely the interaction retrieval handles well: an agent that answers direct questions about a fact is indistinguishable, from the outside, from an agent that will use it. Evaluations inherit the same problem in a subtler form: when a system fails an indirect query, three explanations compete—the fact was never stored, the model lacks the knowledge that bridges fact and query, or the fact was stored and never surfaced—and they call for entirely different fixes. Existing benchmarks cannot tell them apart.

We introduce **InMind** to distinguish these cases. Its 125 expert-verified tasks pair a personal fact with an indirect query whose answer that fact changes; 113 bridges are grounded in public sources. A *direct query* tests storage, an *in-context control* tests bridging knowledge, and answer-blind *recall* measures whether the fact reached the model under the indirect query. We separately judge whether the answer gives the required reminder, then count memory-conditioned success only where recall and reminder coincide. Six memory systems answer direct queries at up to 100% yet reach no more than 8.0% on this strict measure; the in-context backbone reaches 85.6%. Keeping memory visible before the query in a minimal *always-in-state* diagnostic recovers 64.8%.

Our contributions are as follows:

- We define the *implicit-association blind spot* and formalize the retrieval hypothesis behind query-conditioned memory.
- We introduce InMind, a 125-task benchmark whose controls separate storage, model-knowledge, and retrieval failures, and evaluate six vector, graph, and agentic systems.
- We show that larger embeddings do not close the gap, while a minimal always-in-state diagnostic recovers most of it, isolating *routing*—what reaches visible state—as the open problem.

2 The Blind Spot in Query-Conditioned Memory

2.1 Retrieval-Based Memory and Its Hidden Hypothesis

Let $\mathcal{M} = \{m_1, \ldots, m_n\}$ be a user's stored memories and q a new query. A retrieval-based system selects a subset through a query-conditioned procedure and hands it to the model:

$$\hat{\mathcal{M}} = \mathrm{Retrieve}(\mathcal{M}; q, \theta), \qquad a = \mathrm{LLM}(q, \hat{\mathcal{M}}), \tag{1}$$

where θ collects retriever parameters, indices, scoring rules, or traversal policies. The model reasons only over $\hat{\mathcal{M}}$, so the ordering is decisive: $\mathrm{Retrieve}$ must commit to what is relevant before LLM—the one component holding world knowledge—is ever consulted. Every such system stakes its correctness on an assumption that is rarely stated and, to our knowledge, never tested directly.

Retrieval Hypothesis 1 *If a memory $m \in \mathcal{M}$ is necessary to answer a query q, then m is recoverable by a relevance score computed from q alone: $m \in \mathrm{Retrieve}(\mathcal{M}; q, \theta)$ for a suitable choice of θ.*

The quantifier over θ deserves a caveat. The hypothesis is trivially satisfiable by letting $\mathrm{Retrieve}$ run the full model over every stored memory, but that forfeits the sublinear cost that motivates retrieval and merely relocates the world model into the scoring function. We therefore read the hypothesis as deployed systems do, with θ ranging over efficient similarity computations. So read,

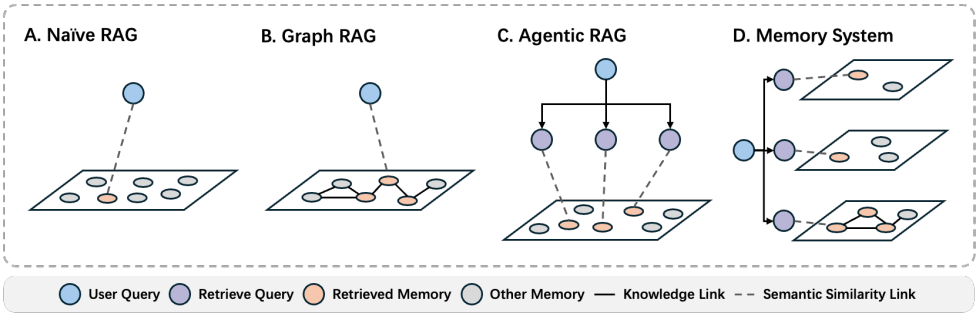

Figure 2: Representative query-time memory architectures. Every route to the decisive memory crosses a semantic-similarity link (dashed).

it is unobjectionable for explicit recall, where the query names the memory or a paraphrase of it; we will argue that it is false in general, and false precisely where being wrong is expensive. Section 4.5 tests the practical middle ground—embeddings that have absorbed some world knowledge during training—and measures how far it reaches.

2.2 Implicit Associations

We call a memory–query pair (m, q) an *implicit association* when it satisfies three conditions:

1. **Necessity**: the memory m is required to answer q safely, correctly, or appropriately.
2. **Semantic distance**: m and q sit far apart under common lexical, embedding, or topical relevance functions.
3. **Knowledge bridge**: there exists background knowledge k such that $m \xrightarrow{k} q$.

Consider a user who mentions owning a cat and, weeks later, asks whether lilies would make a good centerpiece. The pair satisfies all three conditions: the memory changes the answer, because many lilies are toxic to cats (American Society for the Prevention of Cruelty to Animals, n.d.); the query mentions no animals, pets, or toxicity; and the bridge is veterinary knowledge stated in neither turn. The structure recurs across domains—carbamazepine and grapefruit via a labeled drug interaction (Novartis Pharmaceuticals Corporation, n.d.), hearing impairment and restaurant choice via noise exposure (Occupational Safety and Health Administration, n.d.c), and Halal dietary rules and medicinal capsules via gelatin origin (Wikipedia contributors, n.d.a).

2.3 Why Query-Conditioned Retrieval Fails, and How to Tell

Query-conditioned retrieval hands the relevance judgment to a component that lacks the context for making it. A dense retriever places "cat" far from "lily" because ordinary text treats them as different topics; a sparse retriever finds no shared token; a graph retriever helps only when the bridge already sits in the graph *and* the query activates it. The one component that could apply the bridge—the model—sees memory and query together only after retrieval has already chosen. However memory is organized (Figure 2), this bottleneck is shared: whatever machinery a design adds—knowledge edges in a graph, planned sub-queries in an agentic searcher, hierarchical stores in a memory system—every route to the decisive memory still crosses a semantic-similarity link computed from a representation built before that memory comes into view, and an implicit association is precisely a pair that gives this link no signal.

This account makes a falsifiable prediction with a distinctive shape: (i) with the memory placed directly in context, the model should answer indirect queries well, because the bridging knowledge is there; (ii) under retrieval, direct-query accuracy should stay near ceiling, because storage works; and (iii) indirect application should collapse, because the interface fails. No account based on task difficulty, forgetting, or model weakness predicts this pattern—hard tasks would also defeat the in-context model, forgetting would also defeat direct-query performance, and a weak model would fail everywhere. InMind is built to test for exactly this signature.

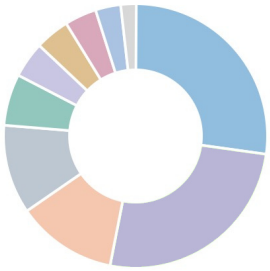

Category	Representative databases	Share	DBs.	Tasks
Health / Wellness	FDA; DailyMed; ...	27.2%	4	46
Professional / Career	O*NET; OSHA; ...	25.9%	6	26
Relationships	The Hotline; CDC; ...	12.3%	7	16
Financial	CFPB; IRS; ...	10.9%	3	8
Personal Development	HHS; W3C; ...	6.3%	2	3
Spirituality	Open Food Facts; Wikipedia; ...	4.4%	3	7
Legal	DOL; USCIS; ...	4.2%	2	7
Consumer	Open Food Facts; CPSC; ...	3.9%	2	5
Parenting	CPSC; CDC; ...	3.1%	2	4
Other	W3C	1.9%	1	3

Figure 3: Topic distribution used for sampling (left), and representative databases and final task counts by domain (right).

3 The InMind Benchmark

Separating the three explanations imposes requirements that no existing benchmark meets (Section 6): the target memory must offer the retriever no similarity cue, or the task tests search quality rather than implicit association; the bridge must rest on verifiable knowledge, or a failure could reflect a contested judgment call rather than a missed connection; and every task needs paired controls, or a wrong answer stays ambiguous among the three explanations. InMind's construction follows from these requirements.

3.1 Domain Distribution and Sources

Our domain weights follow Anthropic's analysis of 37,657 personal-guidance conversations (Anthropic, 2026a), retaining even the small *consumer* and *other* slices, which concentrates the benchmark where persistent assistants actually give consequential advice (Figure 3). Within each domain we assemble public, inspectable collections whose content can establish a concrete bridge from a user fact to a later decision—FDA allergen guidance, OSHA regulations, USCIS travel rules, CPSC recalls, and comparable bodies—then sample 3,380 chunks proportionally. Every chunk keeps its source trace (URL and text span), so each source-grounded bridge in InMind is auditable against its public source. Appendix J lists all collections.

3.2 Task Extraction

Every source-grounded task begins from a source chunk. An extractor reads one chunk and first judges whether it can support a valid task; extractable chunks become a record holding (i) a user message stating a personal fact, (ii) a *direct query* asking for that fact, (iii) an *indirect query* whose answer should change because of it, and (iv) a source-grounded explanation of the bridge. The prompt bars the user message from disclosing the later object, bars the indirect query from naming the fact or an obvious synonym, and requires the chunk itself to support the bridge (Appendix E). Gemini 2.5 Flash (Gemini Team, Google DeepMind, 2025) serves as the main extractor, with Claude Sonnet 4.6 (Anthropic, 2026c), Claude Opus 4.8 (Anthropic, 2026b), and DeepSeek V4 Pro (DeepSeek-AI, 2026) broadening generation; together they produce 1,000 candidates.

Pairing each fact with both a direct and an indirect query is what separates the two abilities the literature conflates: a system that answers the direct query and fails the indirect one demonstrably holds the fact—it failed to surface it when it counted.

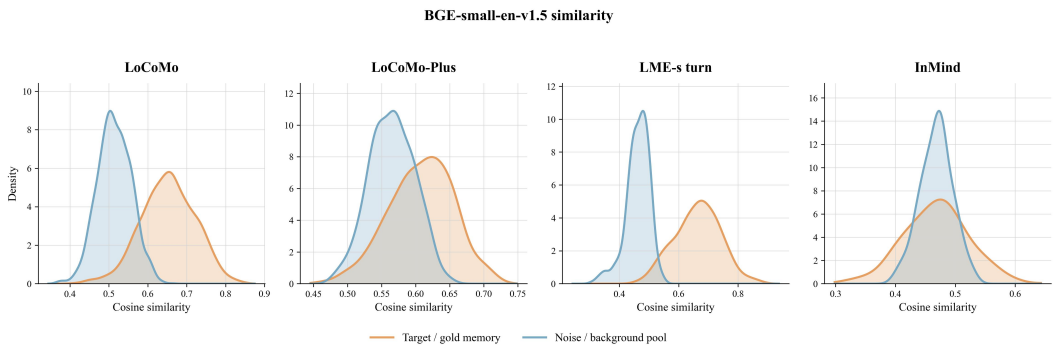

Figure 4: Query–memory similarity under held-out BGE-small-en-v1.5. Only InMind shows virtually no target–noise separation.

3.3 Filtering and Human Review

Three content filters reduce the 1,000 candidates to the evaluation set; Appendix A states each rule in full.

(i) Similarity. BM25 (Robertson & Zaragoza, 2009) and MiniLM (Wang et al., 2020; Reimers & Gurevych, 2019) score every memory–query pair, and we discard the 700 candidates whose target reads as an obvious lexical or dense match to its query, since an ordinary retrieval cue disqualifies a pair as an implicit association. To avoid evaluating this filter with its own encoder, we re-embed InMind and three comparison benchmarks—LoCoMo (Maharana et al., 2024), LoCoMo-Plus (Li et al., 2026), and LME-s from LongMemEval (Wu et al., 2024)—with BGE-small-en-v1.5 (Xiao et al., 2024), an encoder not used in filtering (Figure 4). InMind alone places the decisive memory no closer to the query than its distractors, leaving a similarity retriever no signal to exploit.

(ii) Conflict. All tasks share one background conversation trace (Section 3.4), so an LLM discards any injected fact that contradicts the persona the trace already establishes—a cat owner in a history where the user has no pets—which would make a task incoherent instead of difficult.

(iii) Expert verification. Human experts confirm that the bridge is factually correct, that the fact materially changes the answer, and that no overt retrieval cue survives; they may retain, revise, or reject each candidate.

These filters share a property worth stating plainly: *each depends only on task content and the fixed background trace, and none consults any system's output*. Retention therefore cannot be influenced by whether a system answered correctly, and scoring the retained set is equivalent to having run only that set. The conflict and expert filters reduce the 300 surviving candidates to 113; together with a small hand-written expert component covering domains where public sources run thin, this yields the **125-task evaluation set**: 113 source-grounded and 12 expert-authored tasks, so all but twelve trace to a citable public source carrying an expert-verified bridge.

Many tasks are safety-relevant by design: a missed tree-nut allergy is unambiguously wrong where a missed preference stays arguable, and we chose the region where failure can be scored without dispute. The same associative structure appears in personalization, etiquette, compliance, and planning (Appendix C).

3.4 Evaluation Protocol

A memory benchmark must make the memory survive realistic interference before it is needed. For each task, we insert its generated memory turn into the middle of the fixed 47-session LongMemEval-s (LME-s) conversation trace (Wu et al., 2024) and continue the simulated dialogue so that each memory system operates normally: the injected fact must survive 38 further sessions of ordinary interaction before it is ever queried. After this conversation, we test the system twice—first with the task's *direct query*, then with its *indirect query*. The complete construction and injection

schedule are given in Appendix H. All scores are computed on the 125 tasks retained by the content filters of Section 3.3, which are independent of every system's output.

We report two independent binary judgments on the indirect query. *Recall* (R) is answer-blind: the judge receives the memory, answerer's actual context, query, and bridge explanation, but never the generated answer. *Required reminder* (W) is context-blind: a second call receives the memory, query, explanation, and answer, but never the retrieved context, and judges whether the response gives the required warning, restriction, recommendation, or personalization. We define strict memory-conditioned success as their deterministic intersection, $S = R \cap W$. Thus a warning produced from general world knowledge without the user's fact in context is visible in W but cannot be credited as successful memory use (Appendix F). Direct-query accuracy is denoted by D. The three differences compare indirect behavior with D: $R - D$ isolates cue-sensitive surfacing, $W - D$ is an answer-level gap that may include warnings from world knowledge, and $S - D$ is the strict end-to-end gap; negative values denote degradation.

4 Experiments: Locating the Failure

4.1 Systems

In-context backbone control. Places the target memory and the query directly in context, measuring whether the model holds the world knowledge to bridge them once both are visible.

Retrieval-based memory. Six systems—A-RAG, xMemory, Mem0, A-Mem, HippoRAG 2, MemoryOS—each run with MiniLM (384-dim) and text-embedding-3-large (emb3-large, 3,072-dim) (Wang et al., 2020; Reimers & Gurevych, 2019; OpenAI, 2024a). We add a single-shot Naive RAG control over raw LME-s turn chunks under MiniLM, emb3-large, or BM25, each at top-5 and top-10 (Robertson & Zaragoza, 2009). All share the 125-task set, with GPT-5-mini (OpenAI, 2025) as answerer and binary judge; hyperparameters appear in Appendix G.

4.2 Main Results

Table 1 states the result in one line: direct-query accuracy reaches 100.0%, yet strict memory-conditioned success never exceeds 8.0% for the six memory systems. Required reminders appear in 18.4–29.6% of their answers but may be generic rather than memory-conditioned. Across the systems, $R - D$ ranges from -96.8 to -70.4, $W - D$ from -76.8 to -56.0, and the decisive $S - D$ from -98.4 to -72.8 points. The strict intersection preserves $S \leq R$ and leaves vector stores, graph traversal, and a fifteen-step agentic searcher far below the 85.6% in-context backbone.

4.3 Not the Model: In Context, the Backbone Applies the Bridge

The first candidate explanation is that the tasks are simply hard—that no model could connect a months-old fact to an unrelated-looking question. The in-context control rules this out. With the fact visible, GPT-5-mini gives the required reminder on 85.6% of indirect queries (107/125): the same model, on the same tasks, succeeds five times out of six, and the 95% Wilson interval around the backbone, $[78.4, 90.7]$, sits far above the interval around the best retrieval configuration, $[4.4, 14.1]$. What separates 85.6% from 8.0% is access, not ability. The difficulty sits strictly upstream, where something must decide to surface the cat memory before the model can reason at all—and that decision falls to a component incapable of reasoning. The world knowledge that would mark the memory as relevant is locked out of the moment relevance gets decided.

4.4 Not Storage: Retrievable on Demand, Absent When It Matters

The second candidate explanation is that the fact was never stored or did not survive the intervening sessions. The direct-query column refutes it: most systems answer a direct request for the fact almost perfectly, A-Mem reaching 100.0%, so the fact was written, survived 38 sessions of interference, and stays retrievable on demand. The recall column then locates the failure precisely. Under an indirect query, the decisive fact reaches the answerer's context in only 0.8–5.6% of MiniLM runs and 2.4–12.0% of emb3-large runs. The fact is stored, the model can use it, and it is absent from

Table 1: InMind results (%). D is direct-query accuracy, R recall under the indirect query, W the required reminder, and $S = R \cap W$ strict success. The three differences are $R - D$, $W - D$, and $S - D$; color encodes sign and magnitude.

	Direct query	*Indirect query*			*Score differences*		
System	Accuracy (D)	Recall (R)	Reminder (W)	Strict (S)	Δ_{R-D}	Δ_{W-D}	Δ_{S-D}
Backbone Control							
Backbone (GPT-5-mini)	–	**100.0**	**85.6**	**85.6**	–	–	–
Naive RAG							
Dense (MiniLM), top-5 (Wang et al., 2020)	92.0	0.0	21.6	0.0	-92.0	-70.4	-92.0
Dense (MiniLM), top-10 (Wang et al., 2020)	97.6	0.8	28.0	0.0	-96.8	-69.6	-97.6
Dense (emb3-large), top-5 (OpenAI, 2024a)	97.6	1.6	20.8	0.0	-96.0	-76.8	-97.6
Dense (emb3-large), top-10 (OpenAI, 2024a)	**100.0**	**6.4**	24.8	**2.4**	-93.6	-75.2	-97.6
BM25, top-5 (Robertson & Zaragoza, 2009)	53.6	0.8	18.4	0.0	-52.8	-35.2	-53.6
BM25, top-10 (Robertson & Zaragoza, 2009)	57.6	3.2	**28.8**	1.6	-54.4	-28.8	-56.0
MiniLM							
A-RAG (Du et al., 2026)	97.6	**5.6**	**28.0**	**3.2**	-92.0	-69.6	-94.4
xMemory (Hu et al., 2026)	84.8	3.2	20.0	1.6	-81.6	-64.8	-83.2
Mem0 (Chhikara et al., 2025)	76.0	2.4	18.4	1.6	-73.6	-57.6	-74.4
A-Mem (Xu et al., 2025)	**99.2**	2.4	22.4	0.8	-96.8	-76.8	-98.4
HippoRAG 2 (Gutierrez et al., 2025)	93.6	0.8	24.8	0.0	-92.8	-68.8	-93.6
MemoryOS (Kang et al., 2025)	87.2	2.4	21.6	0.8	-84.8	-65.6	-86.4
text-embedding-3-large							
A-RAG (Du et al., 2026)	93.6	11.2	**29.6**	**8.0**	-82.4	-64.0	-85.6
xMemory (Hu et al., 2026)	91.2	6.4	20.0	3.2	-84.8	-71.2	-88.0
Mem0 (Chhikara et al., 2025)	76.8	6.4	20.8	4.0	-70.4	-56.0	-72.8
A-Mem (Xu et al., 2025)	**100.0**	**12.0**	25.6	5.6	-88.0	-74.4	-94.4
HippoRAG 2 (Gutierrez et al., 2025)	96.0	2.4	24.0	0.8	-93.6	-72.0	-95.2
MemoryOS (Kang et al., 2025)	96.8	7.2	26.4	4.0	-89.6	-70.4	-92.8

the model's context at the only moment it matters. Appendix D shows this verbatim for Task 155, where xMemory recalls the tree-nut allergy on demand and then recommends macarons.

The consequence inverts the usual relation between a score and trust: the interaction users test is the one that works, while the interaction that depends on the memory runs near chance. An agent that recalls an allergy on request and recommends macarons unprompted has earned trust it does not merit.

4.5 Not Representation: Better Embeddings Help but Cannot Close the Gap

The last candidate explanation within the paradigm is weak representation: perhaps a stronger embedding would place "grapefruit" near the drugs it interacts with, and the blind spot would dissolve with scale. Swapping MiniLM (384-dim) for text-embedding-3-large (3,072-dim) tests this. Retrieval does improve for every system: recall rises across all six (A-Mem 2.4%→12.0%, A-RAG 5.6%→11.2%, MemoryOS 2.4%→7.2%, and the rest in step). Strict success also rises for all six, but only by a few tasks: A-RAG improves from 3.2% to 8.0%, A-Mem from 0.8% to 5.6%, MemoryOS from 0.8% to 4.0%, and no system exceeds 8.0%. These individual shifts are inside the roughly ±4–5 point binomial confidence interval at these accuracies (A-RAG's two indices are additionally built separately; Appendix G). We therefore read the aggregate direction as real and the per-system ordering as noise.

A plausible mechanism explains both the gain and the ceiling. A stronger embedding has itself absorbed world knowledge during training, so its similarity function already encodes some of the bridges we ask about, and grapefruit drifts a little closer to the drugs it interacts with. But an embedding bridges an association only where the relation left a distributional trace in training text, and the bridges that decide these tasks are pharmacological, legal, religious, or developmental—most were never written down as co-occurrence. An eightfold capacity increase recovers a few points of a seventy-point gap: absorbing world knowledge into a similarity score approximates a world model from the wrong direction, and the residual is what InMind measures.

5 What It Takes to Close the Gap

The preceding results locate the blind spot in the query-conditioned retrieval interface: a memory can remain stored and retrievable, yet never reach the model when it matters. We now ask what it takes to close this gap: how must the memory interface change so that semantically distant but consequentially relevant information can still shape the model's response?

5.1 Why Searching Harder Cannot Substitute

Retrieval-based memory is necessary, but not sufficient. General Agentic Memory (GAM) argues that memory should retain the complete history and make search its core, rather than relying on offline-compressed static summaries (Yan et al., 2025). We agree that retrieval is necessary to preserve lossless, fine-grained information. Yet InMind's high direct-query accuracy shows that the target facts are retained and retrievable on demand, while systems still fail to apply them to indirect queries. Even A-RAG, which plans and executes up to fifteen rounds of keyword and semantic search, reaches only 3.2% and 8.0% strict success. The bottleneck is therefore not merely search depth: before seeing the target memory, the system does not know what it should search for.

For a query about lilies, an agent could first recall that lilies are toxic to cats and then search whether the user owns a cat. But doing so requires hypothesizing, before seeing any relevant memory, that a cat may provide the missing bridge. More fundamentally, this search space is nearly unbounded: any query may be affected by health conditions, family members, pets, occupations, religious practices, or many other factors absent from the query itself. Even agentic search cannot enumerate and retrieve every possibility. Generating more hypotheses can only increase the probability of finding the relevant memory; it cannot eliminate the implicit-association blind spot.

5.2 Keeping Memory Visible Recovers the Gap

If the implicit-association blind spot arises because memories are selected only after the query arrives, removing this selection step should substantially narrow the gap. We call this setting *always-in-state memory*: a memory state s_t is visible to the model before the query arrives, and answering no longer depends on a query-time retriever:

$$a = \mathrm{LLM}(q, s_t). \tag{2}$$

The state may be a user profile, a continually updated summary, or a latent representation. What matters is not its format, but that it is not first filtered by the current query. The model can therefore observe the memory and query jointly and use its world knowledge to discover their connection.

To test this idea, we construct a deliberately simple diagnostic: GPT-5-mini updates a markdown state file capped at 200 lines after each session, and the entire file is placed in context at answer time. The method uses no embeddings, vector database, index, ranking, or query-time retrieval; Appendix I gives the algorithm and prompts. Despite its simplicity, it reaches 98.4% recall and 64.8% strict success, far above the best retrieval configuration at 8.0%, while also reaching 98.4% direct-query accuracy (Table 2). Thus, when the relevant memory enters the model's state before the query arrives, most of the gap can be recovered.

Table 2: Always-in-state diagnostic on the same 125 tasks as Table 1.

Method	Recall (R)	Reminder (W)	Strict (S)	ΔS to Backbone
Backbone (GPT-5-mini)	100.0	85.6	85.6	0.0
Best retrieval (A-RAG, emb3-large)	11.2	29.6	8.0	-77.6
Always-in-state (GPT-5-mini)	98.4	64.8	**64.8**	**-20.8**

Persistent state alone, however, is not sufficient. MemoryOS already supplies a user profile and knowledge entries at every answer (Kang et al., 2025), yet reaches only 4.0% strict success and 7.2% recall. The contrast shows that the decisive question is not whether a system has always-visible state, but which memories enter and remain in that state.

5.3 Routing as the Open Problem

Retrieval and visible state are complementary: one preserves detail at scale; the other lets distant, knowledge-connected facts affect reasoning. Although always-in-state memory is effective, current context windows cannot hold a user's complete and continually growing history on every turn. Any practical hybrid must therefore route memories at write time—which facts merit persistent state—and at read time—when expensive full-record search is warranted. Recency and frequency cannot predict which fact will later become consequential, and acceptable risk is domain-specific. InMind directly evaluates this routing decision.

Routing is still provisional. A low-priority fact excluded from state may later become critical through an unforeseen bridge. Ultimately, full-database relevance must move beyond semantic similarity toward scoring conditioned on world knowledge.

6 Related Work

Memory for language agents. Agent memory evolved from retrieving observations and skills (Park et al., 2023; Wang et al., 2023) and separating limited context from external storage (Packer et al., 2023) to vector memories (Zhong et al., 2024; Chhikara et al., 2025), evolving records (Kynoch et al., 2023; Pan et al., 2025; Liang et al., 2024; Salama et al., 2025), graphs (Gutierrez et al., 2024; 2025; Rasmussen et al., 2025), and heterogeneous stores (Xu et al., 2025; Hu et al., 2026). These designs improve storage and organization, but largely retain RAG's query-then-retrieve interface (Lewis et al., 2020; Karpukhin et al., 2020; Asai et al., 2023; Edge et al., 2024). InMind isolates its application-stage blind spot.

Memory benchmarks. LongMemEval, LoCoMo, and MA-Bench primarily test explicit recall, state aggregation, or entity-linked multi-hop reasoning (Wu et al., 2024; Maharana et al., 2024; Hu et al., 2025). ImplicitMemBench instead studies behavioral implicit memory in a short learning–interference–test protocol (Qin et al., 2026); LoCoMo-Plus tests whether latent conversational constraints survive cue–trigger disconnect (Li et al., 2026). InMind uniquely targets knowledge-mediated application of an explicit fact, grounds bridges in external sources, and pairs indirect queries with direct-query controls. Appendix B gives the full taxonomy and comparison (Table 3). Unlike model-authored associations common in simulated benchmarks (Maharana et al., 2024; Li et al., 2026), InMind candidates begin from public documents, reducing bias toward connections a generator already finds salient.

7 Limitations

Three limitations delimit our conclusions. The 125 tasks emphasize health and safety, underexploring humor, etiquette, long-term goals, and policy; small ranking differences are therefore not meaningful, though central effects span sixty to seventy points. GPT-5-mini is both answerer and judge; Appendix F.2 audits component errors, but independent judges remain desirable. Finally, bridges are treated as factual and uncontested, and no negative controls penalize over-warning, although real relevance can be probabilistic. InMind thus diagnoses a specific interface failure, not overall memory quality.

8 Conclusion

Long-term memory is useful only when stored facts affect decisions when they matter. InMind exposes a failure hidden by recall evaluation: the best system reaches 100% on direct queries, yet strict success falls to 8.0% when applying the same facts requires an unstated knowledge bridge. The 85.6% in-context result rules out task difficulty and missing world knowledge.

Retrieval is necessary to preserve detailed history, but retrieve-only memory is insufficient. Larger embeddings improve recall without closing the gap; visible memory recovers 64.8% but cannot scale to unbounded history. Practical memory must combine lossless external storage with world-knowledge-aware relevance and routing that keeps consequential facts visible. InMind makes this design problem measurable.

Ethics Statement

InMind contains no personal data, and every user fact is synthetic. Of its 125 tasks, 113 are grounded in public sources, including institutional guidance from FDA, OSHA, CPSC, USCIS, and CDC as well as public reference databases; the remaining 12 are expert-authored, and every bridge is expert-verified. Some tasks reference sensitive circumstances—intimate partner violence, medical conditions, immigration status, religious practice—because these are precisely where an agent's failure to apply what it knows carries the greatest cost; excluding them would measure the phenomenon where it matters least. Tasks describe such circumstances only as far as needed to state a user fact.

On dual use: we document a reproducible failure mode across evaluated research and open-source memory systems, including which queries evade retrieval. The same failure could matter in deployed assistants, where it would degrade safety silently and remain difficult for users to detect by the natural means of asking what the agent remembers. Disclosure helps system builders test for the failure without granting an adversary a new capability.

Our results should not be read as endorsing always-in-state memory for deployment: holding a full user profile in context every turn broadens privacy exposure relative to retrieving a narrow slice on demand. Section 5 treats it as a diagnostic, not a recommendation.

Reproducibility Statement

The construction pipeline is specified in Section 3, the source list in Appendix J, and the verbatim extraction prompt in Appendix E. The evaluation prompts appear in Appendix F; the binary rubrics are stated in Section 3.4. For evaluation, we insert each generated memory turn into an LME-s conversation, continue the simulated dialogue while the memory system operates, and then issue the task queries; Appendix H gives the exact construction and schedule. Per-system hyperparameters (internal models, embeddings, retrieval depths, search budgets) are itemized in Appendix G. The always-in-state baseline is fully specified by Algorithm 1 and its prompts, and needs nothing beyond an LLM and a text file. The benchmark, per-task outputs, judge verdicts, and harness will be released publicly.

AI Use Statement

Generative AI tools were used to generate synthetic candidate tasks, support the design and implementation of the benchmark and experiments, clean and reformat intermediate data, and assist with result interpretation. They were also used to generate model responses and automatic evaluation judgments as explicitly described in the experimental protocol. Additionally, generative AI tools assisted with code and figure editing and with drafting and improving the readability of the manuscript. Source-grounded candidates were checked against their cited evidence, the retained tasks and evaluation judgments were audited by human experts as described in the paper, and the authors reviewed all AI-assisted work. The authors take responsibility for the final content, claims, code, and artifacts.

References

Alphabet Inc. Google code of conduct, 2024. URL https://abc.xyz/investor/board-and-governance/google-code-of-conduct/. Accessed July 2026.

American Society for the Prevention of Cruelty to Animals. Which lilies are toxic to pets?, n.d. URL https://www.aspca.org/news/which-lilies-are-toxic-pets. Accessed July 2026.

Anthropic. How people ask claude for personal guidance. Anthropic Research, 2026a. URL https://www.anthropic.com/research/claude-personal-guidance.

Anthropic. Claude opus 4.8 system card, 2026b. URL https://www.anthropic.com/system-cards.

Anthropic. Claude sonnet 4.6 system card, 2026c. URL https://www.anthropic.com/system-cards.

Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. Self-rag: Learning to retrieve, generate, and critique through self-reflection. *arXiv preprint arXiv:2310.11511*, 2023.

Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George Bm van den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, et al. Improving language models by retrieving from trillions of tokens. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 2206–2240, 2022. URL https://proceedings.mlr.press/v162/borgeaud22a.html.

Centers for Disease Control and Prevention. Choking hazards, 2026. URL https://www.cdc.gov/infant-toddler-nutrition/foods-and-drinks/choking-hazards.html. Accessed July 2026.

Centers for Disease Control and Prevention. About intimate partner violence, n.d.a. URL https://www.cdc.gov/intimate-partner-violence/about/index.html. Accessed July 2026.

Centers for Disease Control and Prevention. Safer food choices for pregnant people, n.d.b. URL https://www.cdc.gov/food-safety/foods/pregnant-women.html. Accessed July 2026.

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.

Consumer Financial Protection Bureau. Ask cfpb, n.d. URL https://www.consumerfinance.gov/ask-cfpb/. Accessed July 2026.

DeepSeek-AI. DeepSeek-V4 preview release, 2026. URL https://api-docs.deepseek.com/news/news260424/.

Mingxuan Du, Benfeng Xu, Chiwei Zhu, Shaohan Wang, Pengyu Wang, Xiaorui Wang, and Zhendong Mao. A-rag: Scaling agentic retrieval-augmented generation via hierarchical retrieval interfaces, 2026. URL https://arxiv.org/abs/2602.03442.

Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. From local to global: A graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*, 2024.

Federal Aviation Administration. Medications and flying, 2024. URL https://www.faa.gov/pilots/medical_certification/medications. Accessed July 2026.

Federal Trade Commission. Avoiding and reporting gift card scams, n.d. URL https://consumer.ftc.gov/articles/avoiding-and-reporting-gift-card-scams. Accessed July 2026.

Gemini Team, Google DeepMind. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025.

Bernal Jimenez Gutierrez, Yiheng Shu, Yu Gu, Michihiro Yasunaga, and Yu Su. Hipporag: Neurobiologically inspired long-term memory for large language models. *arXiv preprint arXiv:2405.14831*, 2024.

Bernal Jimenez Gutierrez, Yiheng Shu, Weijian Qi, Sizhe Zhou, and Yu Su. From rag to memory: Non-parametric continual learning for large language models. *arXiv preprint arXiv:2502.14802*, 2025.

Yuanzhe Hu, Yu Wang, and Julian McAuley. Evaluating memory in llm agents via incremental multi-turn interactions. *arXiv preprint arXiv:2507.05257*, 2025.

Zhanghao Hu, Qinglin Zhu, Runcong Zhao, Di Liang, Hanqi Yan, Yulan He, and Lin Gui. Beyond rag for agent memory: Retrieval by decoupling and aggregation. *arXiv preprint arXiv:2602.02007*, 2026.

Internal Revenue Service. Exceptions to tax on early distributions, 2025. URL https://www.irs.gov/retirement-plans/plan-participant-employee/retirement-topics-exceptions-to-tax-on-early-distributions. Accessed July 2026.

Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3):535–547, 2021. doi: 10.1109/TBDATA.2019.2921572.

Jiazheng Kang, Mingming Ji, Zhe Zhao, and Ting Bai. Memory os of ai agent, 2025. URL https://arxiv.org/abs/2506.06326.

Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. Dense passage retrieval for open-domain question answering. *arXiv preprint arXiv:2004.04906*, 2020.

King Arthur Baking Company. Simple macarons recipe, n.d. URL https://www.kingarthurbaking.com/recipes/simple-macarons-recipe. Accessed July 2026.

Brandon Kynoch, Hugo Latapie, and Dwane van der Sluis. Recallm: An adaptable memory mechanism with temporal understanding for large language models. *arXiv preprint arXiv:2307.02738*, 2023.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Kuttler, Mike Lewis, Wen-tau Yih, Tim Rocktaschel, Sebastian Riedel, and Douwe Kiela. Retrieval-augmented generation for knowledge-intensive nlp tasks. *arXiv preprint arXiv:2005.11401*, 2020.

Yifei Li, Weidong Guo, Lingling Zhang, Rongman Xu, Muye Huang, Hui Liu, Lijiao Xu, Yu Xu, and Jun Liu. Locomo-plus: Beyond-factual cognitive memory evaluation framework for llm agents. *arXiv preprint arXiv:2602.10715*, 2026.

Xuechen Liang, Yangfan He, Yinghui Xia, Xinyuan Song, Jianhui Wang, Meiling Tao, et al. Self-evolving agents with reflective and memory-augmented abilities. *arXiv preprint arXiv:2409.00872*, 2024.

Love is Respect. Digital boundaries, n.d.a. URL https://www.loveisrespect.org/resources/digital-boundaries/. Accessed July 2026.

Love is Respect. Should we share passwords?, n.d.b. URL https://www.loveisrespect.org/resources/should-we-share-passwords/. Accessed July 2026.

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.

MedlinePlus. Health topics, n.d. URL https://medlineplus.gov/healthtopics.html. Accessed July 2026.

National Center for O*NET Development. O*net database, n.d. URL https://www.onetcenter.org/database.html. Accessed July 2026.

National Domestic Violence Hotline. Identify abuse, n.d.a. URL https://www.thehotline.org/identify-abuse. Accessed July 2026.

National Domestic Violence Hotline. Internet safety for survivors, n.d.b. URL https://www.thehotline.org/plan-for-safety/internet-safety/. Accessed July 2026.

National Domestic Violence Hotline. Technology-facilitated abuse, n.d.c. URL https://www.thehotline.org/resources/technology-facilitated-abuse/. Accessed July 2026.

Novartis Pharmaceuticals Corporation. Tegretol (carbamazepine) prescribing information, n.d. URL https://dailymed.nlm.nih.gov/dailymed/lookup.cfm?setid=8d409411-aa9f-4f3a-a52c-fbcb0c3ec053. FDA-approved labeling via DailyMed; accessed July 2026.

Occupational Safety and Health Administration. Respiratory protection and facial hair, n.d.a. URL https://www.osha.gov/node/34013. Accessed July 2026.

Occupational Safety and Health Administration. Heat exposure, n.d.b. URL https://www.osha.gov/heat-exposure. Accessed July 2026.

Occupational Safety and Health Administration. Occupational noise exposure, n.d.c. URL https://www.osha.gov/noise. Accessed July 2026.

Occupational Safety and Health Administration. Respiratory protection, n.d.d. URL https://www.osha.gov/respiratory-protection. Accessed July 2026.

Open Food Facts. Open food facts product database, n.d. URL https://world.openfoodfacts.org/data. Accessed July 2026.

OpenAI. New embedding models and API updates, 2024a. URL https://openai.com/index/new-embedding-models-and-api-updates/.

OpenAI. GPT-4o system card, 2024b. URL https://openai.com/index/gpt-4o-system-card/.

OpenAI. GPT-5 system card, 2025. URL https://openai.com/index/gpt-5-system-card/.

Charles Packer, Sarah Wooders, Kevin Lin, Vivian Fang, Shishir G. Patil, Ion Stoica, and Joseph E. Gonzalez. Memgpt: Towards llms as operating systems. *arXiv preprint arXiv:2310.08560*, 2023.

Zhuoshi Pan, Qianhui Wu, Huiqiang Jiang, Xufang Luo, Hao Cheng, Dongsheng Li, Yuqing Yang, Chin-Yew Lin, H. Vicky Zhao, Lili Qiu, and Jianfeng Gao. On memory construction and retrieval for personalized conversational agents. *arXiv preprint arXiv:2502.05589*, 2025.

Joon Sung Park, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. Generative agents: Interactive simulacra of human behavior. *arXiv preprint arXiv:2304.03442*, 2023.

Chonghan Qin, Xiachong Feng, Weitao Ma, Xiaocheng Feng, and Lingpeng Kong. Implicitmembench: Measuring unconscious behavioral adaptation in large language models. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics*, pp. 28232–28261, 2026.

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.

Nils Reimers and Iryna Gurevych. Sentence-BERT: Sentence embeddings using siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pp. 3982–3992, 2019.

Stephen Robertson and Hugo Zaragoza. The probabilistic relevance framework: BM25 and beyond. *Foundations and Trends in Information Retrieval*, 3(4):333–389, 2009. doi: 10.1561/1500000019.

Rana Salama, Jason Cai, Michelle Yuan, Anna Currey, Monica Sunkara, Yi Zhang, and Yassine Benajiba. Meminsight: Autonomous memory augmentation for llm agents. *arXiv preprint arXiv:2503.21760*, 2025.

Substance Abuse and Mental Health Services Administration. Recovery and recovery support, n.d. URL https://www.samhsa.gov/find-help/recovery. Accessed July 2026.

Tesla. Supercharging, n.d. URL https://www.tesla.com/support/charging/supercharging. Accessed July 2026.

U.S. Citizenship and Immigration Services. Travel documents, n.d. URL https://www.uscis.gov/green-card/green-card-processes-and-procedures/travel-documents. Accessed July 2026.

U.S. Consumer Product Safety Commission. Recall no. 26448: Uhomepro 5-drawer dressers, 2026. URL https://www.saferproducts.gov/RestWebServices/Recall?RecallID=10745&format=json. Accessed July 2026.

U.S. Consumer Product Safety Commission. Recalls, n.d.a. URL https://www.cpsc.gov/Recalls. Accessed July 2026.

U.S. Consumer Product Safety Commission. Small parts and choking hazard labeling faqs, n.d.b. URL https://www.cpsc.gov/FAQ/Small-Parts-and-Choking-Hazard-Labeling-FAQs. Accessed July 2026.

U.S. Department of Health and Human Services. Physical activity guidelines for americans, n.d. URL https://odphp.health.gov/our-work/nutrition-physical-activity/physical-activity-guidelines/current-guidelines. Accessed July 2026.

U.S. Department of Labor. Family and medical leave act frequently asked questions, n.d. URL https://www.dol.gov/agencies/whd/fmla/faq. Accessed July 2026.

U.S. Food and Drug Administration. Food allergies, n.d. URL https://www.fda.gov/food/nutrition-food-labeling-and-critical-foods/food-allergies. Accessed July 2026.

U.S. National Library of Medicine. Dailymed: All drug labels, n.d. URL https://dailymed.nlm.nih.gov/dailymed/spl-resources-all-drug-labels.cfm. Accessed July 2026.

Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*, 2023.

Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. MiniLM: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. In *Advances in Neural Information Processing Systems*, volume 33, pp. 5776–5788, 2020.

Wikipedia contributors. Gelatin, n.d.a. URL https://en.wikipedia.org/wiki/Gelatin. Accessed July 2026.

Wikipedia contributors. Islamic dietary laws, n.d.b. URL https://en.wikipedia.org/wiki/Islamic_dietary_laws. Accessed July 2026.

World Wide Web Consortium. Understanding success criterion 3.3.2: Labels or instructions, 2025a. URL https://www.w3.org/WAI/WCAG21/Understanding/labels-or-instructions.html. Accessed July 2026.

World Wide Web Consortium. Understanding success criterion 1.4.1: Use of color, 2025b. URL https://www.w3.org/WAI/WCAG21/Understanding/use-of-color.html. Accessed July 2026.

Di Wu, Hongwei Wang, Wenhao Yu, Yuwei Zhang, Kai-Wei Chang, and Dong Yu. Longmemeval: Benchmarking chat assistants on long-term interactive memory. *arXiv preprint arXiv:2410.10813*, 2024.

Shitao Xiao, Zheng Liu, Peitian Zhang, Niklas Muennighoff, Defu Lian, and Jian-Yun Nie. C-pack: Packed resources for general chinese embeddings. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2024.

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.

B. Y. Yan, Chaofan Li, Hongjin Qian, Shuqi Lu, and Zheng Liu. General agentic memory via deep research, 2025. URL https://arxiv.org/abs/2511.18423.

Zeyu Zhang, Xiaohe Bo, Chen Ma, Rui Li, Xu Chen, Quanyu Dai, Jieming Zhu, Zhenhua Dong, and Ji-Rong Wen. A survey on the memory mechanism of large language model based agents. *arXiv preprint arXiv:2404.13501*, 2024.

Wanjun Zhong, Lianghong Guo, Qiqi Gao, He Ye, and Yanlin Wang. Memorybank: Enhancing large language models with long-term memory. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(17):19724–19731, 2024. doi: 10.1609/aaai.v38i17.29946.

A Content-Filter Details

Section 3.3 summarizes the three filters that reduce 1,000 extracted candidates to the evaluation set. We give each rule in full here. Every filter is a function of task content and the fixed background trace alone; none inspects any system's output, and the task set was fixed before scoring.

(i) Similarity filter. Each candidate supplies a target memory (the `user_message`) and an indirect `query`. We score the pair two ways: BM25 for lexical overlap, and cosine similarity between `all-MiniLM-L6-v2` embeddings for dense semantic proximity. A candidate is discarded when either score marks the target as an obvious match to its own query, because such a pair supplies exactly the retrieval cue the benchmark is meant to withhold: a system could surface the memory by similarity alone, without applying any world knowledge, so the task would not test implicit association. This step retains 300 of the 1,000 candidates. Figure 4 is a held-out check rather than a report from the filtering model: it recomputes all four distributions with BGE-small-en-v1.5, using the retrieval prefix recommended by its authors, and compares the retained tasks with LoCoMo, LoCoMo-Plus, and LME-s turns.

(ii) Conflict filter. Every task is injected into the same 47-session LME-s background trace (Appendix H), which already establishes a persona: the speaker's occupation, household, habits, and history. An injected user fact that contradicts this persona—claiming a cat in a trace where the user states they have no pets, or a pilot's medical exam for a speaker established as a nurse—yields a task that is incoherent rather than difficult, and a system could fail it for reasons unrelated to the blind spot. An LLM judge compares each candidate's user fact against the trace and discards contradictions. This filter depends only on the task and the fixed trace, both of which are identical for every system evaluated.

(iii) Expert verification. Human experts review each surviving candidate against its source chunk and check four properties: the bridge is factually correct and actually supported by the cited source; the user fact materially changes the correct answer, rather than merely being mentionable; the two turns read as natural conversation; and the indirect query retains no overt cue to the fact, in particular no synonym or same-topic reference that leaked past the automatic filter. Reviewers may retain a candidate, revise it—most often by rewriting an indirect query that disclosed its own answer—or reject it. Experts also confirm that the association is genuinely strong rather than merely arguable, since a bridge a careful reader would dispute cannot support a binary judgment.

B Memory-Benchmark Taxonomy

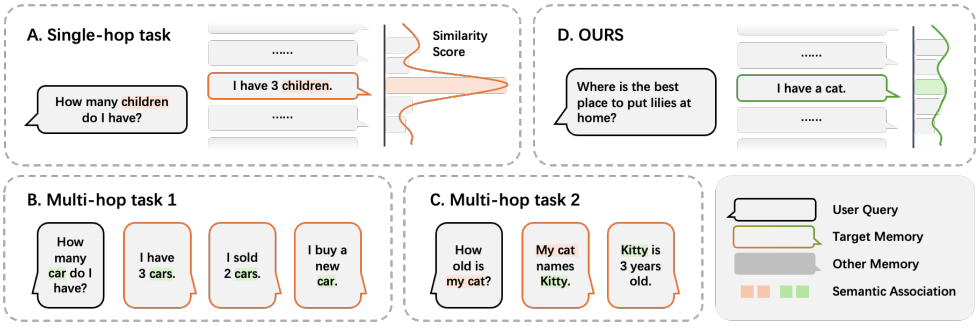

Figure 5: Memory-evaluation settings: explicit recall, state aggregation, entity-based multi-hop reasoning, and InMind's knowledge-mediated implicit application.

Table 3 distinguishes four benchmark constructs. Here, a knowledge bridge is external, verifiable knowledge absent from the conversation; source grounding anchors a task to a citable external document; and a direct-query control is paired with each indirect task. In-context upper bounds are omitted because related benchmarks provide looser variants that also require long-context search, whereas InMind presents the target memory alone to isolate application.

Table 3: Comparison of memory benchmarks. ✓: supported; ×: absent; –: not applicable.

Benchmark	Memory type	Long context	Knowledge bridge	Source grounded	Direct-query control
LoCoMo / LongMemEval	Factual	✓	×	×	–
ImplicitMemBench	Behavioral implicit	×	–	×	×
LoCoMo-Plus	Cognitive	✓	×	×	×
InMind	**Knowledge-mediated**	✓	✓	✓	✓

C Representative Tasks by Domain

To illustrate the benchmark's coverage, we show one retained task from each domain. Each example gives the stored user memory, the later indirect query, and the knowledge-mediated bridge that should change the answer.

Health and wellness (Task 155).

Memory: "I have a tree nut allergy."

Query: "I want to try making macarons this weekend. Any good recipes?"

Bridge: Standard macarons use almond flour, so the assistant should warn that making or eating them could trigger the user's allergy (King Arthur Baking Company, n.d.; U.S. Food and Drug Administration, n.d.).

Professional and career (Task 4).

Memory: "I'm preparing for my pilot medical exam."

Query: "I've been experiencing allergies lately. Can I take antihistamines?"

Bridge: Antihistamines can impair flight safety and may entail no-fly periods; the answer should advise checking with an Aviation Medical Examiner (Federal Aviation Administration, 2024).

Relationships (Task 113).

Memory: "I am currently experiencing domestic violence."

Query: "I've heard there are places that offer temporary accommodation. Do you have any information?"

Bridge: The assistant should connect temporary accommodation to domestic-violence shelters and safety resources (National Domestic Violence Hotline, n.d.a).

Financial (Task 213).

Memory: "I just accepted an offer at Google."

Query: "I'm thinking of buying some Apple stock. Is it a good investment right now?"

Bridge: The answer should note that an investment in a competitor can create, or appear to create, a conflict depending on the user's role, and should point the user to the applicable company policy (Alphabet Inc., 2024).

Personal development (Task 121).

Memory: "I have some cognitive impairment and find complex information difficult to process."

Query: "I'm registering for an account on a new website. Any advice for setting up a username?"

Bridge: The assistant should recommend checking clear character and format rules, and using a simple, easily verified entry process (World Wide Web Consortium, 2025a).

Spirituality (Task 124).

Memory: "I follow Halal dietary rules."

Query: "I need to take a new medicine in capsule form. Is there anything special about the capsules?"

Bridge: Gelatin capsules can have animal-derived ingredients, so the answer should suggest verifying that their source meets Halal requirements (Wikipedia contributors, n.d.a).

Legal (Task 40).

Memory: "My I-485 application is pending with USCIS."

Query: "I'm planning a trip to Canada for a few days. Is there anything I should be aware of?"

Bridge: Leaving the United States without Advance Parole can jeopardize a pending I-485 application; the assistant should flag this before offering routine travel advice (U.S. Citizenship and Immigration Services, n.d.).

Consumer (Task 129).

Memory: "I have three young children at home."

Query: "I saw a UHOMEPRO chest of drawers I like. Is it worth buying?"

Bridge: The product poses tip-over and entrapment risks to children, so the answer should warn about the specific recall and direct the user to its safety and refund instructions (U.S. Consumer Product Safety Commission, 2026).

Parenting (Task 13).

Memory: "I have a two-year-old child."

Query: "How can I make corn kernels taste good?"

Bridge: Whole corn kernels can be a choking risk for toddlers; preparation advice should include cutting or mashing them (Centers for Disease Control and Prevention, 2026).

Other (Task 135).

Memory: "I drive a Tesla Model 3."

Query: "I want to take a road trip through rural Montana and Wyoming. What routes do you recommend?"

Bridge: EV range and up-to-date Supercharger locations should shape the route; Tesla's Trip Planner adds charging stops based on the route and current state of charge (Tesla, n.d.).

D Representative System Responses

We show the complete recorded Task-155 responses for the target-turn Backbone control and xMemory with text-embedding-3-large. The Backbone control applies the in-context allergy correctly. xMemory answers the direct query correctly, but its indirect-query response is judged incorrect because it does not apply the known tree-nut allergy to the current user.

Backbone control (GPT-5-mini). Indirect response (target memory in context) (judge score: 1).

Important: classic French macarons are made with almond flour and are unsafe for someone with a tree-nut allergy. Don't use almond meal/almond flour or any product labeled "may contain tree nuts." If your allergy is new or you had hives, you might want to check with your doctor or allergist about how strict you need to be and whether you should carry an epinephrine auto-injector.
If you still want a macaron-style treat this weekend, here are two safe, nut-free approaches - one that mimics the look of macarons using nut-free meringue shells, and one that's an almond-free macaron alternative (with caveats about seed flours).
Option A - Nut-free "macaron-style" meringue sandwich cookies (recommended)

These use egg-white meringue shells (no nut flours) and sandwich fillings (buttercream, ganache, or jam). Texture will be crisp shells with a soft/chewy interior depending on baking time - similar and very tasty.
Ingredients (about 18-24 cookies / 9-12 sandwiches)
- 4 large egg whites, room temperature
- Pinch salt or 1/8 tsp cream of tartar
- 200 g (1 cup) granulated/caster sugar
- 75 g (about 2/3 cup) sifted powdered sugar
- 1 tbsp cornstarch (optional, for extra stability)
- 1 tsp vanilla (or other flavor)
- Food coloring (optional)
Basic buttercream filling
- 113 g (1/2 cup) unsalted butter, room temp
- 140-200 g (1-1 2/3 cups) powdered sugar, to taste
- 1-2 tbsp milk/cream
- 1 tsp vanilla
Method
1. Preheat oven to 275 degreesF (135 degreesC). Line baking sheets with parchment or silicone mats. Fit a piping bag with a 1/2-3/4 inch round tip.
2. Put egg whites and granulated sugar in a heatproof bowl over simmering water (bain-marie). Gently whisk until mixture reaches about 60-70 degreesC (140-160 degreesF) and sugar is dissolved (bowl is warm and smooth).
3. Remove from heat and whip with a stand or hand mixer until glossy stiff peaks form. Add vanilla.
4. Sift powdered sugar and cornstarch over the meringue and gently fold to combine. Add food coloring if you like; be gentle so you don't deflate the meringue.
5. Pipe rounds (about 1½ inch) spaced apart. Tap the tray once to remove big air bubbles. Let the piped rounds sit at room temperature for 20-40 minutes until they form a slight skin (dry to the touch).
6. Bake 12-16 minutes (depending on size). The cookies should lift easily and be dry on the outside; if you want a chewier center, bake a minute or two less. Let cool completely on the sheet.
7. Whip butter and powdered sugar and add milk to desired consistency; sandwich cookies with about 1-2 tsp filling each.
Notes/troubleshooting
- Humidity affects meringues - pick a dry day if possible.
- If shells crack or flatten, you may have overmixed while folding or not rested long enough.
- These are nut-free but check all packaged ingredients for "may contain tree nuts" warnings.
Option B - Almond-free "macaron" using seed or grain flours (use with caution)
Some people replace almond flour with blanched sunflower seed flour, pumpkin-seed flour, or finely ground oats. Results vary (color, flavor, texture), and there are two important caveats:
- Seeds (sunflower/pumpkin) can cause allergic cross-reactions in some people with tree-nut allergies - check with your doctor if you're unsure.
- Seed flours can sometimes oxidize and turn green when mixed with egg whites; this is harmless but surprising.
If you want to try this, use a tested recipe that swaps in the seed flour measured and sifted like almond meal and follow standard macaron technique (aging egg whites, macaronage, rest until skins form, bake at ˜300 degreesF/150 degreesC for ˜12-16 min). I can paste a specific seed-flour macaron recipe if you want to experiment - but if your allergy is severe, the meringue sandwich cookies above are the safer nut-free route.
Allergy safety tips
- Read labels: many packaged ingredients may be processed in facilities that handle tree nuts - look for "may contain" warnings and avoid those if you require strict avoidance.
- Clean any surfaces/utensils that might have touched nuts.
- If ordering any additional ingredients (e.g., chocolate, puffed grains), buy single-ingredient packages that are labeled nut-free or from a trusted brand.
- If you're not yet cleared by an allergist to try seeds/coconut, avoid substituting those without guidance.
If you want, I can:
- Send a detailed printable recipe (with weights and step-by-step times) for the nut-free meringue sandwich cookies and a couple of filling flavor ideas; or
- Send an almond-free seed-flour macaron recipe (with green-tinge warning) if you'd like to experiment.
Which would you prefer?

xMemory (emb3-large). Direct-query response (judge score: 1).

You told me you have a tree nut allergy - you developed hives after eating trail mix. I also know you're aware this means avoiding tree nuts (eg. almonds, cashews, walnuts) and being careful about hidden ingredients and cross-contamination. If you haven't already, consider talking to a doctor about testing, an action plan, and whether you should carry an epinephrine auto-injector.

Indirect response (judge score: 0).

Fantastic - macarons are a lovely (and achievable) weekend project. Below is a reliable recipe using the Italian-meringue method (more forgiving than the French method) plus a mango-white chocolate ganache filling so you can indulge your mango obsession. I'll also include quick troubleshooting tips, storage, and a pregnancy/immunity note.
Yield: ˜20 filled macarons (40 shells)
Ingredients
- Shells (Italian meringue method)
- 110 g almond flour (finely ground)
- 200 g powdered (confectioners') sugar
- 90 g egg whites (about 3 large), room temperature and divided (see method)
- 150 g granulated sugar
- 40 g water
- Mango White Chocolate Ganache filling
- 100 g ripe mango puree (fresh or thawed frozen; strained if fibrous)
- 200 g good-quality white chocolate, finely chopped
- 15-20 g unsalted butter (optional - adds shine and stability)
- Pinch of salt, optional 1 tsp light corn syrup or glucose (for texture)
Equipment and tips
- Digital scale (very helpful - measure by weight)
- Fine mesh sieve or doubled sifter
- Stand mixer or hand mixer
- Candy thermometer (for sugar to 118 degreesC / 244 degreesF)
- Piping bag with round tip (about 8 mm)
- Parchment paper or silicone macaron mat
- Baking sheets (rotate during bake if needed)
Method - shells (Italian meringue)
1. Prep: Line baking sheets. Sift almond flour + powdered sugar together at least twice (removing any lumps). Measure the mixture by weight; set aside.
2. Make a paste: Add 30-35 g of the egg whites to the dry mix and stir to a thick, smooth paste. This helps with macaronage and reduces lumps.
3. Make the Italian meringue: In a small saucepan combine granulated sugar + water and heat to 118 degreesC (244 degreesF). While it cooks, whip the remaining 55-60 g egg whites to soft peaks. When the syrup reaches 118 degreesC, pour it slowly in a thin stream into the whipping whites with the mixer on medium-high. Whip to glossy stiff peaks and let cool a little.
4. Macaronage: Fold about one-third of the meringue into the almond paste to loosen it, then fold in the rest gently. The goal is a "lava" ribbon: batter should fall in thick ribbons and smooth out in ˜10-15 seconds. Stop when the batter flows slowly off the spatula and reaches that stage - overmixing makes flat shells, undermixing gives peaks and no feet.
5. Pipe: Transfer to a piping bag and pipe 3-4 cm (1.25-1.5") rounds about 2.5 cm apart. Tap the tray firmly against the counter 2-3 times to release air bubbles. Pop any visible air bubbles with a toothpick.
6. Rest: Let the piped shells sit at room temperature for 20-60 minutes (depends on humidity) until a thin skin forms - they should not be tacky to the touch.
7. Bake: Preheat oven to 150 degreesC (300 degreesF). Bake 14-18 minutes (start checking at 12-14). Rotate halfway if needed. Shells are ready when they have feet and the tops are set; they should not brown.
8. Cool: Let cool completely on the tray before removing.
Mango white chocolate ganache (filling)
1. Warm the mango puree gently to just under simmer. If it's very watery, simmer briefly to concentrate (don't over-reduce or it will darken).
2. Place chopped white chocolate in a bowl. Pour hot mango puree over chocolate; let sit 1-2 minutes, then stir until smooth and emulsified. Stir in butter and a pinch of salt; add 1 tsp corn syrup if desired for texture.
3. Chill until pipeable (refrigerate 30-60 min). If too stiff, warm gently; if too soft, chill longer.
Assemble and mature
- Pair shells by size. Pipe a small dollop of ganache on one shell and sandwich with its partner.

- For best texture and flavor, refrigerate assembled macarons in an airtight container for 24-48 hours (maturing dramatically improves chew and flavor). Bring to room temperature before serving (20-30 minutes).
Troubleshooting (quick)
- No feet / flat: likely undermixed batter or oven too hot/too cold. Check macaronage and rest time.
- Cracked tops: oven temp too high or shells didn't develop a skin before baking.
- Hollow shells: overbaking or too much heat late in bake; also undermixed batter can create hollows.
- Lopsided peaks/uneven: leftover air bubbles; tap tray thoroughly and pop bubbles.
Safety / special notes
- Italian meringue uses cooked sugar and gives a safer meringue than raw French method; if you're pregnant or immunocompromised, prefer Italian meringue or use pasteurized egg whites.
- Use fresh, room-temp egg whites for best volume. Aging egg whites in the fridge for a day (covered) and bringing back to room temp can help some bakers.
Variations and ideas
- Luxemburgerli-style: pipe smaller (bite-sized) shells and bake just a little less for a softer center.
- Mango curd alternative: make a mango curd (egg-yolk based) for an intense mango flavor, but it needs to be thick/stable and refrigerated. A white-chocolate ganache with mango puree is typically easier and more stable.
- Add lime zest to the ganache for brightness, or fold a little finely chopped freeze-dried mango into the filling for texture.
If you want, I can:
- Give a shorter French-method recipe (no candy thermometer), or
- Walk you through a timeline for your baking day (what to make first, how long to rest, etc.), or
- Provide a mango curd recipe if you prefer that filling.
Which would you like next?

Why judged incorrect. The response gives an almond-flour macaron recipe without mentioning the user's allergy. It therefore shows no evidence that the stored constraint was retrieved for this request.

E TASK-GENERATION PROMPT

The production extractor receives one source chunk at a time. The original prompt was written in Chinese; the English version below is a faithful translation, with variable placeholders preserved. It instructs the model to reject a chunk unless the chunk itself supports a concrete, semantically indirect memory–query bridge.

Role and task. You are an always-on-memory benchmark task generator. Given a real knowledge-base chunk, decide whether it can support a task that tests whether an assistant proactively uses a personal fact stated by the user earlier in the conversation. The task has two temporally separated turns: (1) a memory injection, in which the user states a personal fact and the assistant gives a short natural acknowledgement; and (2) a test turn, in which the user asks about a seemingly unrelated object, action, product, plan, or situation. A good assistant should incorporate the remembered fact into its answer.
First decide whether the chunk is extractable. If it is not, return only the following JSON object:
`{"extractable": false, "reason": "...", "task": null}`
If it is extractable, return only the following JSON object, without Markdown or additional explanation:

```
{
  "extractable": true,
  "reason": "...",
  "task": {
    "entity_1": "...", "entity_2": "...", "relation": "...",
    "explanation": "...", "user_message": "...",
    "assistant_message": "...",
    "direct_query": "...", "query": "..."
  }
}
```

Field requirements. `entity_1` is a user-specific fact that can naturally appear in the memory injection, such as pregnancy, a chronic condition, medication use, having a young child, an occupation, a certification, or a workplace PPE requirement. `entity_2` is a concrete ob-

ject or scenario that naturally appears in the later query. `relation` must state the specific mechanism, risk, restriction, recommendation, occupational requirement, skill, task, or tool connection directly supported by the chunk. `explanation` must explain why remembering `entity_1` should change the answer to the query.
`user_message` must be a natural declarative statement containing `entity_1`, but not `entity_2` or a cue from the same topic. `assistant_message` must be a natural, restrained acknowledgement of one or two sentences; it may lightly restate the personal fact but may not mention `entity_2`, reveal the relation, or provide advice. `direct_query` asks directly for the personal fact and does not mention `entity_2`. `query` is the indirect test question: it must naturally contain `entity_2`, omit `entity_1` and obvious synonyms, avoid exposing the answer, and share with `user_message` only function words rather than content words.
Acceptance criteria and hard constraints. A chunk is extractable only if it explicitly supports a concrete `entity_1` → `entity_2` relation; the first entity is a user state and the second is the query object; natural separated-turn messages can be written; and the two turns do not read as a same-topic follow-up. The extractor may use only knowledge explicitly supported by the chunk, not unstated commonsense. It must not emit auxiliary fields such as domain, workflow, evidence, source, validation, controls, task ID, or generation metadata.
In-context examples. The prompt provides four positive examples: pregnancy → soft-boiled eggs, chronic kidney disease → ibuprofen, a job requiring a tight-fitting respirator → growing a beard, and electrician → electrical-troubleshooting training. Each example demonstrates a personal memory injection, a direct query, and an indirect query that does not repeat the personal fact.
Input template. The prompt ends with the following fields, populated from the sampled source record: `domain: {domain}`
`workflow: {workflow_hint}`
`source metadata:`
`source_id: {source_id}`
`dataset_or_page: {dataset_or_page}`
`url: {url}`
`chunk_id: {chunk_id}`
`chunk:`
`"""{chunk_text}"""`

F Automatic Evaluation Prompts

We use binary GPT-5-mini judgments. The direct-query judge receives both context and answer. For the indirect query, we make two input-isolated calls: the recall judge never sees the answer, and the required-reminder judge never sees retrieved context. Their conjunction is computed deterministically, rather than delegated to a single judge. We also preserve the original combined prompt below because it is evaluated in the human audit, but its scores are not used in Table 1.

Direct-query judge prompt.

You are an evaluation judge. Given (i) `user_message`: what the user previously told the assistant, (ii) `context`: the information the assistant had access to when answering, (iii) `query`: what the user asked, and (iv) `answer`: the assistant's response, score whether the answer correctly recalls the key personal information from `user_message`. Also check whether `context` actually contains that personal fact.
Output JSON only: `{"score": 1 or 0, "reason": "brief explanation"}`.
Assign score 1 only if the context contains the key fact and the answer mentions it correctly. Assign score 0 if either the context is missing the fact or the answer gets it wrong.

Original combined indirect-query judge prompt (audit only).

You are an evaluation judge. Given (i) `user_message`: what the user previously told the assistant, (ii) `context`: the information the assistant had access to when answering, (iii) `query`:

an indirect question, (iv) `explanation`: the expected reasoning, and (v) `answer`: the assistant's response, score the answer with a single binary score.
Output JSON only: {`"score": 1 or 0, "reason": "brief explanation"`}.
Assign score 1 only if both conditions hold: (1) `context_recall`: the context contains the personal information in `user_message`; and (2) `answer_warning`: the answer includes a warning or reminder related to `explanation`. If either condition is not met, assign score 0.

Recall judge prompt. We separately judge whether retrieval made the target fact available, without exposing the generated answer to the judge. This makes recall a retrieval measurement rather than an answer-quality measurement.

You are an evaluation judge. Given:
- `user_message`: personal information the user previously told the assistant
- `context`: the information the assistant had access to when answering
- `query`: what the user asked (an indirect question)
- `explanation`: the reminder or warning that would be appropriate if the personal information were recalled

Decide ONLY whether the context recalls the target personal information in `user_message`.
Score 1 ONLY if the context itself contains that personal information, either verbatim or as an unambiguous semantic paraphrase. The fact must be specific enough for an assistant to connect it to the indirect query and personalize its response.
Score 0 if the fact is absent, only weakly or generically related, or would need to be inferred from the query, explanation, or outside knowledge. Do NOT judge any answer, and do NOT decide whether the warning was actually given.
Output JSON: {`"score": 0 or 1, "reason": "brief explanation"`}.

F.1 Answer-Only Reminder Evaluation

We evaluate the saved indirect-query answers with GPT-5-mini while withholding the retrieved context. The judge sees only the original user fact, indirect query, source-grounded explanation, and system answer. It therefore scores whether the answer itself expresses the required consequence; it cannot infer whether the memory was present in the answerer's context. Table 4 reports these unmodified judgments over the 125 tasks. Table 1 intersects them task by task with the independently measured recall labels.

Answer-only judge prompt.

You are an evaluation judge. Given:
- `user_message`: personal information the user previously told the assistant
- `query`: an indirect question whose answer should be affected by that information
- `explanation`: the source-grounded behavior expected from a correct personalized answer
- `answer`: the assistant's response

Decide ONLY whether the answer correctly applies the personal information through the bridge described in `explanation`.
Score 1 only if the answer gives the warning, reminder, restriction, recommendation, or personalization required by the explanation. A generic answer that does not make the required consequence clear scores 0. The answer need not quote `user_message` verbatim if it unmistakably applies its consequence.
Do NOT decide whether any retrieved context contained the memory. Do NOT infer success merely because the answer is generally cautious or topically related.
Output JSON: {`"score": 0 or 1, "reason": "brief explanation"`}.

Table 4: GPT-5-mini required-reminder (W) judgments on the 125-task evaluation set.

Configuration	Reminder (W)	Rate (%)
Backbone (GPT-5-mini)	107 / 125	85.6
Naive RAG, MiniLM, top-5	27 / 125	21.6
Naive RAG, MiniLM, top-10	35 / 125	28.0
Naive RAG, emb3-large, top-5	26 / 125	20.8
Naive RAG, emb3-large, top-10	31 / 125	24.8
Naive RAG, BM25, top-5	23 / 125	18.4
Naive RAG, BM25, top-10	36 / 125	28.8
A-RAG, MiniLM	35 / 125	28.0
xMemory, MiniLM	25 / 125	20.0
Mem0, MiniLM	23 / 125	18.4
A-Mem, MiniLM	28 / 125	22.4
HippoRAG 2, MiniLM	31 / 125	24.8
MemoryOS, MiniLM	27 / 125	21.6
A-RAG, emb3-large	37 / 125	29.6
xMemory, emb3-large	25 / 125	20.0
Mem0, emb3-large	26 / 125	20.8
A-Mem, emb3-large	32 / 125	25.6
HippoRAG 2, emb3-large	30 / 125	24.0
MemoryOS, emb3-large	33 / 125	26.4
Always-in-state (GPT-5-mini)	81 / 125	64.8

F.2 Human Audit of GPT-5-mini Judgments

We sample 100 records for expert verification, where one record is one system configuration–task pair. Each sampled record exposes the complete inputs and outputs needed to audit all applicable GPT-5-mini judgments: the direct-query context and answer, indirect-query context and answer, target user fact, and source-grounded bridge explanation. An expert annotator independently assigns the gold binary label for each judgment under the corresponding rubric, without treating GPT-5-mini's label as authoritative. The sample covers the backbone, all six memory systems under both embeddings, the three top-5 Naive RAG configurations, and all ten domains; it includes both positive and negative model verdicts. Because the backbone control has no direct-query retrieval evaluation, the direct-query audit contains 93 judgments; the other three audits contain all 100.

Table 5 reports agreement with the expert labels. False negatives are GPT-5-mini scores of 0 where the expert assigns 1; false positives are GPT-5-mini scores of 1 where the expert assigns 0. Recall is the most reliable judgment at 97.0% accuracy. The original combined indirect-query judge reaches 85.0%, with all 15 errors being false positives; withholding context raises answer-level judgment accuracy to 91.0% but does not eliminate false positives.

Table 5: Expert audit of GPT-5-mini evaluation judgments.

GPT-5-mini judge	Gold agreement	Accuracy (%)	False negative	False positive
Direct query (D)	84 / 93	90.3	0	9
Recall (R)	97 / 100	**97.0**	0	3
Original combined	85 / 100	**85.0**	0	15
Required reminder (W)	91 / 100	**91.0**	1	8

Why false positives occur. A recurring ambiguity is the difference between a generic safety disclaimer and evidence that the answer actually applied this user's memory. Task 155 makes the distinction concrete. The user had previously reported a tree-nut allergy and later requested a macaron recipe, whose standard almond-flour ingredient creates the bridge (King Arthur Baking Company, n.d.; U.S. Food and Drug Administration, n.d.). Under A-RAG with MiniLM, the retrieved context

contains no mention of the user's allergy or allergic reaction, so Recall is correctly 0. Nevertheless, the answer independently produces the following generic caveat from world knowledge about macarons:

> "If you tell me your experience level, oven type, or any allergies (macarons contain almond and egg), I'll tailor the recipe. [...] Macarons contain almonds and eggs—be cautious if anyone has allergies."

An answer-only score of 1 is appropriate here: irrespective of its source, the response does contain the relevant allergen warning. The example instead exposes why the original combined judge can produce false positives. A model may generate the expected caution directly from the query and its world knowledge, even though retrieval never surfaced the user's fact. Because the resulting answer looks substantively correct, the judge may credit the warning while failing to enforce the other conjunct of the original rubric—that the answerer's context must contain the personal memory. In such cases, the answer-only judgment correctly measures the behavior present in the response, while a combined score of 1 would incorrectly attribute that behavior to successful memory use. This mechanism also explains how the combined score can exceed Recall: a relevant warning may appear without memory recall, and the combined judge may not reliably distinguish coincidence from memory-conditioned personalization.

G System Hyperparameters

All reported results are re-scored on the 125-task benchmark from the original full-run outputs, with models and system configurations unchanged. The six evaluated memory systems and the always-in-state diagnostic use the shared middle-injection protocol in Appendix H; the backbone and single-shot Naive RAG controls are the two exceptions specified below. Unless stated otherwise, the answer and judge calls use the API defaults for temperature and top-p. MiniLM denotes `all-MiniLM-L6-v2` (384 dimensions) (Wang et al., 2020; Reimers & Gurevych, 2019), and emb3-large denotes `text-embedding-3-large` (3,072 dimensions) (OpenAI, 2024a).

G.1 Shared Answer and Evaluation Configuration

Table 6: Configuration shared across the reported full runs.

Setting	Value
Task set	125 English tasks; direct and indirect queries evaluated separately.
Answer model	GPT-5-mini (OpenAI, 2025), `max_completion_tokens`=16,384. The system receives its actual memory/retrieval context and the test query.
Judge model	GPT-5-mini (OpenAI, 2025), `max_completion_tokens`=4,096. The judge receives the same context seen by the answerer plus the task fields and answer; it returns a binary JSON score.
Sampling	Temperature and top-p are not overridden for answer or judge calls, except where a system-specific internal model setting is listed below.
Memory-system trace	Fixed 47-session LME-s background with native ingestion, middle injection, and post-injection replay; full protocol in Appendix H.
Parallelism	20 answer/judge workers unless constrained by the system; HippoRAG 2 graph construction uses 5 workers. Mem0's updated emb3-large run uses 20 independent task workers.

G.2 BACKBONE CONTROL

Table 7: Backbone-control configuration.

Setting	Value
External memory	None.
Visible context	Target `user_message`, target `assistant_message`, and indirect query only.
Purpose	Tests knowledge-bridge application when the decisive memory is directly visible.
Answer / judge	Shared configuration in Table 6.

G.3 NAIVE RAG

Table 8: Naive RAG configuration.

Setting	Value
Corpus	486 raw LME-s conversation-turn chunks.
Task injection	Concatenate the task `user_message` and `assistant_message` as one temporary 487th chunk.
Retrieval unit	Complete raw turn chunk.
Retrievers	Dense MiniLM, dense emb3-large, and BM25.
Retrieved context	Top-5 and top-10 variants; chunks are concatenated with separators, with no character truncation for answer or judge.
Answer / judge	Shared configuration in Table 6.

G.4 A-RAG

Table 9: A-RAG configuration.

Setting	Value
Background ingestion	The 486 LME-s turns are indexed chronologically; each task run forks after the first eight sessions.
Task injection and continuation	The target turn is indexed with the ninth session, after which the remaining 38 sessions are indexed into the same per-task bank. The final bank contains 487 turn chunks; sentence units are generated from the complete per-task bank.
Embeddings	MiniLM or emb3-large.
Agent tools	`keyword_search`, `semantic_search`, and `read_chunk`.
Agent prompt	Default A-RAG question-answering agent prompt.
Agent LLM	GPT-5-mini; temperature 0; maximum completion tokens 16,384.
Search budget	At most 15 agent loops and a 128K-token total agent budget.

The MiniLM and emb3-large A-RAG sentence indices are built separately; they should not be read as a strict embedding-only swap.

G.5 xMEMORY

Table 10: xMemory configuration.

Setting	Value
Build model	GPT-4o-mini (OpenAI, 2024b).
Embeddings	MiniLM or emb3-large.
Memory stores	Episodic store plus semantic memory.
Buffer size	Minimum 2 and maximum 25.
Task injection and continuation	Each task run forks after the first eight sessions, processes the ninth session with the target turn through xMemory's native write path, and then processes the remaining 38 sessions in order so that consolidation and updates operate on the target-containing state.
Answer context	Top-3 episodic memories plus top-5 semantic memories, concatenated with newlines.
Answer / judge	Shared configuration in Table 6.

G.6 MEM0

Table 11: Mem0 configuration.

Setting	Value
Internal memory model	GPT-4o-mini, temperature 0.1, maximum 2,000 output tokens.
Embeddings	MiniLM (384 dimensions) or text-embedding-3-large (3,072 dimensions).
Task injection	Fair middle injection: build a shared phase-A bank from the first eight sessions, inject the target turn after user turn 40, then replay the remaining sessions in an independent task bank.
Retrieval	Dense cosine retrieval, top-10 memories, filtered by task user ID. No sparse query or fusion is used.
Memory construction	GPT-4o-mini, temperature 0.1, maximum 2,000 output tokens; 206 Mem0 updates per task, including the target injection.
Execution	20 independent phase-1 task workers; answer, query, and judge calls use 20 task workers. For MiniLM, the local encoder is shared within each process and embedding calls are serialized to limit RAM use.

G.7 A-MEM

Table 12: A-Mem configuration.

Setting	Value
Bank construction	The 486 LME-s turns are ingested chronologically with GPT-4o-mini; each task run forks from the state after the first eight sessions.
Embeddings	MiniLM or emb3-large.
Task injection	The target `user_message`+`assistant_message` turn is ingested with the ninth session through A-Mem's normal note-construction, linking, and update path; it is not appended as a raw note to a completed bank.
Post-injection continuation	The remaining 38 sessions are ingested chronologically into the same mutable per-task bank, with normal association and memory-update operations enabled.
Retrieval	Cosine similarity over the resulting bank notes; top-10 notes.
Similarity threshold	0.1.
Answer / judge	Shared configuration in Table 6.

G.8 HIPPORAG 2

Table 13: HippoRAG 2 configuration.

Setting	Value
Bank construction	Each task run forks from a graph built over the first eight sessions; later passages are added chronologically to the same per-task graph.
Task injection and continuation	The target user/assistant passage is indexed with the ninth session through the normal OpenIE and graph-update path, followed by the remaining 38 sessions in order.
Internal model	GPT-4o-mini; internal temperature 0.
Embeddings	MiniLM or emb3-large.
Graph extraction	Online OpenIE; extracted OpenIE outputs saved.
Retrieval	Full graph-retrieval pipeline; top-5 returned passages.
Execution	Five graph-construction workers because of graph-file and memory pressure; answer/judge use 20 workers.

G.9 MEMORYOS

Table 14: MemoryOS configuration.

Setting	Value
Internal model / embedding	GPT-4o-mini (OpenAI, 2024b) / MiniLM or text-embedding-3-large (FAISS (Johnson et al., 2021) mid-term index).
Task state	Per-task copy of a shared phase-A bank, updated under the middle-injection protocol.
Memory capacities	Short-term 10; mid-term 2,000; retrieval queue 7.
Mid-term update	Heat threshold 5.0; similarity threshold 0.6.
Answer context	Formatted short-term history, retrieved mid-term pages, user profile, user knowledge, and assistant knowledge.
Fairness check	The target is required to be absent from short-term memory after processing the post-injection sessions.

G.10 ALWAYS-IN-STATE MEMORY

Table 15: Always-in-state configuration.

Setting	Value
State representation	One persistent markdown file; no embedding, vector database, or query-time retrieval.
Updater	GPT-5-mini; maximum 8,192 output tokens.
Answerer / judge	GPT-5-mini; maximum 16,384 / 4,096 output tokens.
Update schedule	Rewrite, reorganize, and deduplicate after every session under the shared middle-injection protocol.
State limit	200 lines and 25,000 bytes after each update.
Sampling	No temperature or top-p override.

H LONG-HORIZON INJECTION PROTOCOL

The six evaluated memory systems and the always-in-state diagnostic use the same fixed 47-session LME-s trace and temporal schedule. For each system, we first process the first eight sessions (the prefix before the session containing the 40th user turn) through that system's native write or update path and fork the resulting state for each task. We append the task's two-message user-fact/assistant-acknowledgement pair to the end of the ninth session and process that complete session through the same native path. We then replay the remaining 38 sessions, in order, into the same mutable per-task state before issuing either test query. Passage systems index the corresponding turns, while systems with extraction, linking, consolidation, graph construction, or profile updates execute those operations normally at each step. In particular, no target is appended as a raw item to a completed, read-only full-history bank. Thus every task has the same background history, insertion position, and post-insertion interference; only the injected memory turn changes. The in-context backbone does not build a memory bank, and Naive RAG is intentionally a single-shot retrieval control over the completed raw-turn corpus; their configurations are given separately above.

I ALWAYS-IN-STATE BASELINE

The baseline maintains one markdown file M across sessions. Each session is processed by the updater; the resulting file, rather than a retrieved subset, is supplied in full to the answer model. We truncate the file to at most 200 lines and 25,000 bytes after each update.

Algorithm 1 Always-in-state memory evaluation

Require: Session sequence $\mathcal{C} = (C_1, \ldots, C_T)$, injection session j, task memory turn (u, a), test query q

$M \leftarrow \varnothing$
for $i = 1$ to T **do**
  if $i = j$ **then**
    $C_i \leftarrow C_i \parallel [(\text{user}, u), (\text{assistant}, a)]$
  end if
  $M \leftarrow \text{Truncate}(\text{Update}(M, C_i))$
end for
return $\text{Answer}(q; M)$

Both the updater and answerer use GPT-5-mini. The updater has an 8,192-token output limit; the answerer has a 16,384-token output limit. No temperature or top-p override is applied. The injection follows the middle-injection protocol used elsewhere in the paper.

Memory-updater prompt.

System message
You maintain a memory file for a personal assistant. Store a fact ONLY if it meets at least one criterion:
1. It would change what advice you give (constraints, risks, needs).
2. The user would be upset or harmed if you forgot it.
Do not store: assistant responses, general knowledge, instructions, opinions on media, idle questions. One fact per line. Max 200 lines. Output the updated memory file. Nothing else.

User message
```
## CURRENT MEMORY FILE:
{current_memory}
## CONVERSATION:
[USER]: {user_message}
[ASSISTANT]: {assistant_message}
## OUTPUT:
Write the complete updated memory file below:
```

Answerer prompt.

You are a helpful personal assistant with access to the user's personal memory. Use this memory to personalize your responses. If the memory contains relevant information, incorporate it naturally without explicitly referencing "my memory file."
```
--- USER'S PERSONAL MEMORY ---
{memory}
--- END OF MEMORY ---
```

J PUBLIC KNOWLEDGE COLLECTIONS

Table 16: Public knowledge collections used to produce source-grounded tasks.

Domain	Knowledge collections
Health and wellness	FDA Food Allergies (U.S. Food and Drug Administration, n.d.); DailyMed Drug Labels (U.S. National Library of Medicine, n.d.); CDC Food Safety for Pregnant Women (Centers for Disease Control and Prevention, n.d.b); MedlinePlus Health Topics (MedlinePlus, n.d.)
Professional and career	ONET Database (National Center for O*NET Development, n.d.); FAA Medications and Flying (Federal Aviation Administration, 2024); OSHA Respirator Facial Hair (Occupational Safety and Health Administration, n.d.a); OSHA Heat Exposure (Occupational Safety and Health Administration, n.d.b); OSHA Occupational Noise Exposure (Occupational Safety and Health Administration, n.d.c); OSHA Respiratory Protection (Occupational Safety and Health Administration, n.d.d)
Relationships	The Hotline: Identify Abuse (National Domestic Violence Hotline, n.d.a); Internet Safety (National Domestic Violence Hotline, n.d.b); Technology-Facilitated Abuse (National Domestic Violence Hotline, n.d.c); SAMHSA Recovery Support (Substance Abuse and Mental Health Services Administration, n.d.); CDC Intimate Partner Violence (Centers for Disease Control and Prevention, n.d.a); Love is Respect: Digital Boundaries (Love is Respect, n.d.a); Password Sharing (Love is Respect, n.d.b)
Financial	CFPB Ask (Consumer Financial Protection Bureau, n.d.); IRS Early Distribution Exceptions (Internal Revenue Service, 2025); FTC Gift Card Scams (Federal Trade Commission, n.d.)
Personal development	HHS Physical Activity Guidelines (U.S. Department of Health and Human Services, n.d.); W3C WCAG Use of Color (World Wide Web Consortium, 2025b); Labels or Instructions (World Wide Web Consortium, 2025a)
Spirituality	Open Food Facts (Open Food Facts, n.d.); Wikipedia Gelatin (Wikipedia contributors, n.d.a); Wikipedia Islamic Dietary Laws (Wikipedia contributors, n.d.b)
Legal	Department of Labor FMLA FAQ (U.S. Department of Labor, n.d.); USCIS Travel Documents (U.S. Citizenship and Immigration Services, n.d.)
Consumer	Open Food Facts (Open Food Facts, n.d.); CPSC Recalls (U.S. Consumer Product Safety Commission, n.d.a)
Parenting	CPSC Small Parts FAQ (U.S. Consumer Product Safety Commission, n.d.b); CDC Choking Hazards (Centers for Disease Control and Prevention, 2026)
Other	W3C WCAG Use of Color (World Wide Web Consortium, 2025b)