

Keep It InMind: Benchmarking the Implicit-Association Blind Spot in Agent Memory

Ruizhe Li^{1,*}, Mingxuan Du^{1,*}, Benfeng Xu², Zhendong Mao^{1,†}

¹University of Science and Technology of China, ²Metastone Technology

*Equal contribution †Corresponding author

Long-term memory systems store what a user says in an external store and retrieve it when a related query arrives. This interface rests on an assumption so natural that it is rarely stated: a memory that is needed will resemble the query that needs it. World knowledge breaks the assumption. A tree-nut allergy should change the answer to a macaron request through their almond-flour ingredient, yet the two texts share no cue a retriever can see. We call this failure mode the *implicit-association blind spot* and introduce InMind, a 125-task, expert-verified benchmark spanning ten life domains, with 113 tasks grounded in citable public sources. Its paired controls separate three explanations that existing evaluations conflate: the fact was never stored, the model lacks the bridging knowledge, or the fact was stored and never surfaced. The verdict is clean. With the decisive memory placed in context, the backbone answers 84.0% of indirect queries; when the same memory must be retrieved, six vector, graph, and agentic memory systems reach at most 14.4%, even though they recall the same facts on demand at up to 100%. An embedding with eight times the dimensionality raises answer-blind target recall for every system yet leaves the gap essentially intact. A minimal diagnostic probe that keeps memory visible before the query arrives recovers most of the gap, locating the failure in the query-conditioned interface itself and pointing to *routing*—deciding which facts must stay visible—as the open problem InMind is built to score.

Project: [Project Website](#) [GitHub Repository](#)

Padra

1 Introduction

Ask a memory-augmented agent what its user is allergic to, and it answers correctly: tree nuts. Ask it a moment later for a macaron recipe, and it responds with enthusiasm and almond flour. Figure 1 illustrates the exchange; Appendix D records an evaluated memory system reproducing it verbatim. Nothing was forgotten—the system recalled the allergy on demand seconds earlier. What failed is subtler and, we will argue, structural: the memory was never brought to bear at the one moment it mattered.



WARNING: Traditional macarons are made with almond flour, which is a tree nut.

This is not fiction. This is how state-of-the-art memory systems—including HippoRAG 2, A-Mem, xMemory, ... —actually behave.

Figure 1 Direct recall does not guarantee memory use: the system recalls the allergy when asked, then fails to apply it to a macaron request. Standard macarons use almond flour, and almond is a major tree-nut allergen (King Arthur Baking Company, n.d.; U.S. Food and Drug Administration, n.d.). Appendix D records this failure in an evaluated memory system.

Language agents are increasingly expected to act as persistent assistants rather than stateless chatbots (Zhang et al., 2024; Packer et al., 2023; Pan et al., 2025): a user states a fact about themselves once and expects it to keep mattering days later, in another conversation, on a topic they never connected to it themselves. The dominant way to meet this expectation is retrieval-based memory. A writer distills what the user said into records; when a new query arrives, the system embeds it, retrieves the nearest records, and pastes them into the prompt (Lewis et al., 2020; Karpukhin et al., 2020; Borgeaud et al., 2022; Asai et al., 2023). The design has real virtues—it scales past the context window, updates modularly, and inherits well-tested machinery from retrieval-augmented generation—and it works on direct questions, because “what am I allergic to?” is itself a retrieval cue.

It fails when relevance is created by knowledge that lives in neither text: the allergy note and the macaron request share no token, topic, or embedding-space neighborhood, and what connects them is a fact about French patisserie stated in neither turn. Underneath every retrieval-based design sits what we call the *retrieval hypothesis*, formalized in Section 2.1: any memory needed to answer a query lies close to that query under some computable relevance score. Explicit recall satisfies the hypothesis. Implicit associations—drug–food interactions, cross-allergies, dietary law, occupational restrictions, household hazards—violate it, because retrieval commits to a relevance judgment before the model ever sees the memory, and the model is the only component of the system that could recognize the connection. Figure 1 is not a hypothetical: HippoRAG 2 (Gutierrez et al., 2025), A-Mem (Xu et al., 2025), xMemory (Hu et al., 2026), and A-RAG (Du et al., 2026) all fail exactly this task in the benchmark we introduce below.

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**, a 125-task benchmark built to separate them. Every task pairs a personal fact with an indirect query whose correct answer changes because of that fact; 113 bridges are grounded in citable public sources and the remaining 12 are expert-authored, with all 125 expert-verified. Two paired controls and one answer-blind measurement isolate the three explanations: a direct *naïve query* tests storage, an *in-context control* tests the backbone’s bridging knowledge, and *target recall* tests whether retrieval delivered the fact to the model’s context. The verdict is unambiguous. Six memory systems spanning vector, graph, and agentic designs recall the injected facts on demand at up to 100% and hold the knowledge to bridge them (the in-context backbone scores 84.0%), yet apply them under indirect queries at no more than 14.4%, with no query-time configuration exceeding 16.0%. A final experiment shows what it takes to close the gap: a deliberately minimal *always-in-state* probe, which keeps memory visible before the query arrives instead of retrieving after it, recovers most of it (68.8%).

Our contributions are as follows:

- We define the *implicit-association blind spot* and make explicit the retrieval hypothesis that current memory architectures assume and never test.
- We introduce InMind, a 125-task expert-verified benchmark with 113 source-grounded tasks, whose paired controls separate three failure modes that prior evaluations conflate: storage failure, missing model knowledge, and retrieval failure.
- We evaluate six memory systems spanning vector, graph, and agentic designs. All recall the injected facts on demand and none applies them reliably under indirect queries; an embedding with eight times the dimensionality raises answer-blind target recall for every system without closing the gap, placing the failure in the retrieval interface itself.
- We show what it takes to close the gap: a deliberately minimal always-in-state probe—nothing more than a diligently maintained profile—recovers most of it, while hybrid systems that already keep profile-style state alongside retrieval do not. We propose no system; the contrast isolates *routing*, whether decision-critical facts actually reach and survive in the visible state, as the open problem.

2 The Blind Spot in Query-Conditioned Memory

2.1 Retrieval-Based Memory and Its Hidden Hypothesis

Let $\mathcal{M} = \{m_1, \dots, m_n\}$ 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:

$$\hat{\mathcal{M}} = \text{Retrieve}(\mathcal{M}; q, \theta), \quad a = \text{LLM}(q, \hat{\mathcal{M}}), \quad (1)$$

where θ collects retriever parameters, indices, scoring rules, or traversal policies. The model reasons only over $\hat{\mathcal{M}}$, so the ordering is decisive: 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: $m \in \text{Retrieve}(\mathcal{M}; q, \theta)$ for a suitable choice of θ .*

The quantifier over θ deserves a caveat. The hypothesis is trivially satisfiable by letting 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, 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, explicit recall 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




Figure 2 Representative query-time memory architectures. Blue, lavender, and orange denote the user query, retrieval query, and retrieved memory. Every route to the decisive memory crosses at least one semantic-similarity (dashed) link; the solid machinery each design adds sits upstream of that hop and cannot substitute for it.

this pattern—hard tasks would also defeat the in-context model, forgetting would also defeat direct recall, and a weak model would fail everywhere. InMind is built to test for exactly this signature.

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




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




Figure 4 Query-memory similarity under BGE-small-en-v1.5, an encoder not used for task filtering. Orange denotes target memories and blue the background pool. 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 and Zaragoza, 2009) and MiniLM (Wang et al., 2020; Reimers and 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—with the task’s direct *naive 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 complementary measurements on the indirect query. *Target recall* is an answer-blind, binary post-hoc judgment of whether the answerer’s actual context contains the target personal fact; the judge receives the memory, context, query, and bridge explanation, but never the generated answer. *Indirect application* is a binary, context-aware answer judgment that assesses whether the response follows the source-grounded bridge (Appendix F). Separating them identifies whether an error arises before the model can see the fact or after it has seen it.

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 and Gurevych, 2019; OpenAI, 2024a). We add a single-shot Naive RAG control over raw LME-s turn chunks under MiniLM, emb3-large, or BM25 (Robertson and 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: naive recall reaches up to 100.0%, indirect application never exceeds 16.0%. The naive column runs high almost everywhere, so the memory-write and direct-retrieval pipelines work and nothing here is forgetting. The application column inverts that picture. Every retrieval configuration—vector stores, graph traversal, an agentic searcher with a fifteen-step budget—lands between 3.2% and 16.0%, and the six memory systems themselves top out at 14.4%, below even the Naive RAG control at 16.0%. The spread among them is smaller than the distance separating all of them from the 84.0% backbone, which makes the blind spot a property of the paradigm rather than of any one implementation. The remainder of this section works through the three candidate explanations in turn.

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

Table 1 Results (%) on the 125-task InMind benchmark: direct *Naive* recall, answer-blind *Target Recall* in the indirect-query context, and end-to-end *Application*. The Δ columns report percentage-point differences (Recall–Naive, Application–Recall, Application–Naive); muted coral and cyan-blue indicate negative and positive changes, with darker shading for larger magnitudes. The backbone receives the target memory in context and therefore has no naive-retrieval score.

System	Naive	Target Recall	Application	$\Delta_{\text{Recall-Naive}}$	$\Delta_{\text{App.-Recall}}$	$\Delta_{\text{App.-Naive}}$
<i>Backbone Control</i>						
Backbone (GPT-5-mini)	–	100.0	84.0	–	-16.0	–
<i>Naive RAG</i>						
Dense (MiniLM) (Wang et al., 2020)	92.0	0.8	9.6	-91.2	+8.8	-82.4
Dense (emb3-large) (OpenAI, 2024a)	97.6	6.4	16.0	-91.2	+9.6	-81.6
BM25 (Robertson and Zaragoza, 2009)	53.6	3.2	9.6	-50.4	+6.4	-44.0
<i>MiniLM</i>						
A-RAG (Du et al., 2026)	97.6	5.6	4.8	-92.0	-0.8	-92.8
xMemory (Hu et al., 2026)	84.8	3.2	4.8	-81.6	+1.6	-80.0
Mem0 (Chhikara et al., 2025)	76.0	2.4	3.2	-73.6	+0.8	-72.8
A-Mem (Xu et al., 2025)	99.2	2.4	6.4	-96.8	+4.0	-92.8
HippoRAG 2 (Gutierrez et al., 2025)	93.6	0.8	8.8	-92.8	+8.0	-84.8
MemoryOS (Kang et al., 2025)	87.2	2.4	8.0	-84.8	+5.6	-79.2
<i>text-embedding-3-large</i>						
A-RAG (Du et al., 2026)	93.6	11.2	7.2	-82.4	-4.0	-86.4
xMemory (Hu et al., 2026)	91.2	6.4	6.4	-84.8	0.0	-84.8
Mem0 (Chhikara et al., 2025)	76.8	6.4	6.4	-70.4	0.0	-70.4
A-Mem (Xu et al., 2025)	100.0	12.0	9.6	-88.0	-2.4	-90.4
HippoRAG 2 (Gutierrez et al., 2025)	96.0	2.4	5.6	-93.6	+3.2	-90.4
MemoryOS (Kang et al., 2025)	96.8	7.2	14.4	-89.6	+7.2	-82.4

reaches 84.0% on indirect queries (105/125): the same model, on the same tasks, succeeds five times out of six, and the 95% Wilson interval around the backbone, [76.6, 89.4], sits far above the interval around the best retrieval configuration, [10.6, 23.4]. What separates 84.0% from 16.0% 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 naive 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 target-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 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: target 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). End-to-end application follows less cleanly: five of six improve—MemoryOS 8.0%→14.4%, A-Mem 6.4%→9.6%, A-RAG 4.8%→7.2%, xMemory 4.8%→6.4%, Mem0 3.2%→6.4%—while HippoRAG 2 declines (8.8%→5.6%), and each individual shift amounts to a few tasks out of 125, inside the ± 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

Section 4 eliminated the model, the storage, and the representation, leaving the query-conditioned interface itself. Two questions follow. Can the gap be closed from within the paradigm, by retaining more and searching harder? And if not, what does close it? We take them in order; the answers bracket an open problem: routing.

5.1 Why Searching Harder Cannot Substitute

Retrieval-based memory is plainly necessary: an external store scales past the context window and preserves a fine-grained record. The question is sufficiency, and the strongest case for it comes from General Agentic Memory (GAM), which keeps the complete history and makes search the core runtime operation, with an agentic researcher that plans, searches, integrates, and reflects over a full page store (Yan et al., 2025). Lossless retention and runtime search do solve the failure they target, which is forgetting. But our systems have already cleared that bar—their naive accuracy proves the fact is retained—and they still fail. Retention gives a retriever nothing to act on when no signal tells it to look: as Section 2.3 argues, the query representation is fixed before the decisive fact comes into view, so searching a larger haystack more thoroughly leaves the needle looking nothing like the query. A-RAG makes this concrete: it plans, issues keyword and semantic searches, reads chunks, loops up to fifteen times, and scores 4.8% and 7.2%, among the lowest here. Iteration cannot rescue a search whose every query issues from the same blind representation.

Retrieval alone is therefore insufficient, and the ability an agentic searcher must acquire is more specific than asking better questions. A memory store is not an oracle: it returns what resembles the query it receives, so even the bridge-aware “what about this user could make lilies dangerous?” retrieves nothing—no stored fact resembles that question either. The probe that succeeds is “does the user keep a cat?”, phrased in the vocabulary of the stored fact rather than of the request; to issue it, a searcher must first recall, unprompted, that lilies are toxic to cats, then single cats out among everything lilies might endanger—allergic housemates, birds, curious toddlers—and query each candidate in turn. This is what it means to *hypothesize the bridge*, and every hypothesis is a guess made before anything in the store confirms it: the enumeration is probabilistic exactly where safety demands reliability, which is why we treat it as an open direction rather than a fix.

5.2 Keeping Memory Visible Recovers the Gap

If the query-conditioned interface causes the failure, then removing it—while changing nothing else about the task—should recover most of the gap, even for an unsophisticated method. The claim is falsifiable, so we test it.

An *always-in-state* memory departs from Section 2.1 in one respect: a representation s_t reaches the model before the query arrives, and no query-time retriever selects it:

$$a = \text{LLM}(q, s_t). \quad (2)$$

The state may be an injected profile, a persistent summary, or a latent memory state; persistent visibility defines it, whatever the format. The model reasons over s_t and q jointly, discovering a cat–lily bridge with no mismatched query standing in the way. We emphasize that this design is neither ours nor rare: profile-style

state is a standard component of deployed assistants, and practical memory systems are typically *hybrids* that keep such a state alongside a retrieved store—MemoryOS, evaluated above, already places a user profile and knowledge entries in every answer context (Kang et al., 2025). What has been missing is measurement: no benchmark tells a hybrid whether its persistent slice actually carries the facts that must stay visible.

To isolate the variable, we run the design at its barest: one markdown file, capped at 200 lines, rewritten by a GPT-5-mini updater after each session and prepended whole to the system prompt at answer time—no embeddings, no vector store, no index, no ranking, no query-time retrieval (Appendix I). This is nothing more than a diligently maintained profile, and we propose it for nothing: a 200-line file is no architecture—its state needs rewriting every interaction, and its finite context budget makes facts crowd each other out as volume grows. Table 2 reports the result: 68.8% on indirect queries, against at most 16.0% for any query-time configuration, with direct recall matched at 98.4%.

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

Method	Naive	Indirect	$\Delta_{\text{Indirect-Backbone}}$
Backbone (GPT-5-mini)	–	84.0	0.0
Best retrieval configuration	97.6	16.0	-68.0
Always-in-state (GPT-5-mini)	98.4	68.8	-15.2

The value here is diagnostic, and we state it carefully. The probe is not a controlled ablation of any one system—it differs in state format and in updater as well—but it demonstrates sufficiency: a design whose memory is visible before the query arrives, with nothing more than competent write-time curation behind it, recovers most of what the entire retrieval paradigm loses. What separates it from every row of Table 1 is that no query-time retriever stands between the model and the memory, and that one design change outweighs the accumulated apparatus of vector stores, knowledge graphs, and agentic search. Hypothesis 1 is what fails; the retriever works as designed.

The comparison with MemoryOS sharpens the point. Carrying an always-visible profile is not by itself enough: MemoryOS posts the strongest application score among the six memory systems (14.4% with emb3-large) and still sits more than fifty points below the probe, and its answer-blind target recall (7.2%) shows the decisive fact rarely reached even its persistent slice. Taken together, the probe and MemoryOS pin down the real variable: not persistent visibility as such, but *what is visible when the query arrives*. The question a hybrid must answer is therefore not whether to keep a profile—deployed systems already do—but what earns a place in it, which is the routing problem Section 5.3 takes up.

5.3 Routing as the Open Problem

The two designs are complementary necessities—retrieval preserves a lossless record at scale, always-visible state lets distant knowledge-connected facts be reasoned over—and hybrids combining them already exist. What they lack is a criterion. Every hybrid embeds a routing decision: which memories earn persistent state at write time, and when a query justifies an expensive search over the fine-grained record at read time. Today that decision falls to proxies—recency, frequency, MemoryOS’s heat-thresholded promotion—none of which can see the distant bridge that will one day make a fact decision-critical; and the right criterion is vertical-specific besides, since a medical assistant, a financial advisor, and a coding agent price the same forgotten fact very differently. This is the setting InMind is built to serve: to our knowledge it is the first benchmark that scores the routing decision itself, because a hybrid whose router leaves a decision-critical fact out of the visible state fails its indirect query.

Routing defers the problem rather than solving it. Facts too unimportant to earn persistent state can still turn decision-critical later through a distant bridge, and they are exposed the moment the system falls back on ordinary retrieval. Closing that gap needs a relevance function of a different kind (Section 8).

Table 3 Comparison of memory benchmarks. ✓: supported as a paired, per-task property; ✕: absent; −: not applicable to the benchmark’s construct (the main task is itself direct recall, or the design deliberately excludes external knowledge). The control column marks a paired diagnostic supplied by the benchmark, not merely the presence of recall questions.

Benchmark	Memory type	Long context	Knowledge bridge	Source grounded	Direct-recall control
LoCoMo / LongMemEval	Factual	✓	✕	✕	−
ImplicitMemBench	Behavioral implicit	✕	−	✕	✕
LoCoMo-Plus	Cognitive	✓	✕	✕	✕
InMind	Knowledge-mediated	✓	✓	✓	✓

6 Related Work

Memory for language agents. Early systems retrieve observations to sustain simulated behavior or accumulate skills (Park et al., 2023; Wang et al., 2023), and MemGPT separates a limited context from external storage (Packer et al., 2023). Later work extends the pipeline: vector stores distill compact reusable facts (Zhong et al., 2024; Chhikara et al., 2025); temporal and self-evolving memories revise records over time (Kynoch et al., 2023; Pan et al., 2025; Liang et al., 2024; Salama et al., 2025); graph systems expose relational structure for multi-hop retrieval (Gutierrez et al., 2024, 2025; Rasmussen et al., 2025); heterogeneous stores separate memory types before aggregation (Xu et al., 2025; Hu et al., 2026). Across this line of work, the query-then-retrieve interface inherited from retrieval-augmented generation remains the default (Lewis et al., 2020; Karpukhin et al., 2020; Borgeaud et al., 2022; Asai et al., 2023; Yan et al., 2024; Jeong et al., 2024; Sarthi et al., 2024; Edge et al., 2024). However these designs organize memory, they share the inference pattern of Figure 2: retrieval fires from a representation built before the model can see the decisive memory. These efforts improve what an agent stores; application is the gap we isolate.

Memory benchmarks. Long-term-memory benchmarks ask whether an agent uses information from earlier interactions. Standard formulations test explicit recall, where query and target share a direct cue; state aggregation, where temporally ordered memories combine into a current state; or entity-based multi-hop reasoning, where an intermediate memory links query to answer (Wu et al., 2024; Maharana et al., 2024; Hu et al., 2025); Appendix B illustrates these settings. Recent work reaches past factual recall. Implicit-MemBench evaluates behavioral implicit memory—procedural learning, priming, conditioning—in a short Learning–Interfere–Test protocol (Qin et al., 2026). LoCoMo-Plus evaluates *cognitive* memory: whether latent conversational constraints such as a user’s state, goal, or value survive cue–trigger semantic disconnect (Li et al., 2026). InMind targets knowledge-mediated application of an *explicit* user fact, where relevance rests on an external, verifiable bridge that the conversation never states.

Construction provenance separates it further. Where a benchmark’s associations are authored by an LLM—as in the simulated dialogues and generated tasks common in this space (Maharana et al., 2024; Li et al., 2026)—each association is by definition one models find natural, and the same prior that proposed it at construction time will tend to re-derive it at test time, so such benchmarks preferentially sample the connections models already make unaided. InMind’s source-grounded bridges come instead from public documents: a task enters the pool because a citable source asserts the connection, whether or not any model finds it salient, and the extractor is barred from supplying knowledge beyond its source chunk (Appendix E). The association distribution is thereby anchored to expert knowledge rather than to model priors. Table 3 adds the properties that distinguish it: the knowledge bridge, source grounding, and a paired direct-recall control that separates a stored-but-unsurfaced fact from one never stored. In-context upper bounds we leave untabulated because looser analogues are common—LoCoMo and LongMemEval report full-history baselines, ImplicitMemBench an oracle-injection analysis, LoCoMo-Plus full-context backbone runs—though these bundle long-context search together with knowledge application, where InMind’s paired control isolates application by presenting the target memory alone. In the table, “knowledge bridge” means external, verifiable knowledge absent from the conversation, and “source-grounded” that tasks are anchored to citable external documents (113 of InMind’s 125 tasks; the remainder are expert-authored, and every bridge is expert-verified either way).

7 Limitations

The benchmark is diagnostic, and its 125 tasks lean toward health, wellness, and safety, where failure is easiest to judge, leaving humor, etiquette, long-term goals, and institutional policy uncovered. At $n = 125$ a few percentage points are sampling noise; our conclusions rest on sixty-to-seventy-point effects. GPT-5-mini serves as both answerer and judge, a self-preference risk that could inflate absolute scores. The bias applies to the backbone control and the retrieval systems alike and cannot manufacture the gap between them, and the expert audit in Appendix F.2 bounds the error of each judge; still, an independent judge model remains the most significant methodological gap.

We also assume each bridge is factual and uncontested, where real ones can be probabilistic, jurisdiction-dependent, or disputed; and an agent optimized for InMind could over-warn without our rubric seeing it, since applying the bridge earns credit and over-eagerness costs nothing. Measuring that needs negative controls, left to future work.

8 Conclusion

A memory can be essential to a query without resembling it, its relevance created by world knowledge the retriever lacks. On InMind this blind spot costs the leading memory systems the difference between 84.0% and 16.0%, while those same systems recall those same facts on demand at up to 100%. The uncomfortable reading is that the field has spent its effort where the problem already worked: storage, indexing, graph construction, and iterative search all improve what an agent finds once it knows what to look for—and that is the whole difficulty. Query-conditioned retrieval asks a similarity function for a judgment that requires a world model, then consults the world model afterwards. Until a relevance function conditioned on knowledge exists, an agent’s memory stays reliable where the user could have told it again, and unreliable where they assumed it was kept in mind.

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. Both judge prompts appear in Appendix F; the binary rubric is 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.

References

- Alphabet Inc. Google code of conduct, 2024. <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. <https://www.asPCA.org/news/which-lilies-are-toxic-pets>. Accessed July 2026.
- Anthropic. How people ask claud for personal guidance. Anthropic Research, 2026a. <https://www.anthropic.com/research/claude-personal-guidance>.
- Anthropic. Claude opus 4.8 system card, 2026b. <https://www.anthropic.com/system-cards>.
- Anthropic. Claude sonnet 4.6 system card, 2026c. <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*, pages 2206–2240, 2022. <https://proceedings.mlr.press/v162/borgeaud22a.html>.
- Centers for Disease Control and Prevention. Choking hazards, 2026. <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. <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. <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. <https://www.consumerfinance.gov/ask-cfpb/>. Accessed July 2026.
- DeepSeek-AI. DeepSeek-V4 preview release, 2026. <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. <https://arxiv.org/abs/2602.03442>.
- Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitan, 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. https://www.faa.gov/pilots/medical_certification/medications. Accessed July 2026.
- Federal Trade Commission. Avoiding and reporting gift card scams, n.d. <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. <https://www.irs.gov/retirement-plans/plan-participant-employee/retirement-topics-exceptions-to-tax-on-early-distributions>. Accessed July 2026.
- Soyeong Jeong, Jinheon Baek, Sukmin Cho, Sung Ju Hwang, and Jong C. Park. Adaptive-rag: Learning to adapt retrieval-augmented large language models through question complexity. *arXiv preprint arXiv:2403.14403*, 2024.
- 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. <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. <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. <https://www.loveisrespect.org/resources/digital-boundaries/>. Accessed July 2026.
- Love is Respect. Should we share passwords?, n.d.b. <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. <https://medlineplus.gov/healthtopics.html>. Accessed July 2026.
- National Center for O*NET Development. O*net database, n.d. <https://www.onetcenter.org/database.html>. Accessed July 2026.
- National Domestic Violence Hotline. Identify abuse, n.d.a. <https://www.thehotline.org/identify-abuse>. Accessed July 2026.
- National Domestic Violence Hotline. Internet safety for survivors, n.d.b. <https://www.thehotline.org/plan-for-safety/internet-safety/>. Accessed July 2026.
- National Domestic Violence Hotline. Technology-facilitated abuse, n.d.c. <https://www.thehotline.org/resources/technology-facilitated-abuse/>. Accessed July 2026.
- Novartis Pharmaceuticals Corporation. Tegretol (carbamazepine) prescribing information, n.d. <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. <https://www.osha.gov/node/34013>. Accessed July 2026.
- Occupational Safety and Health Administration. Heat exposure, n.d.b. <https://www.osha.gov/heat-exposure>. Accessed July 2026.

- Occupational Safety and Health Administration. Occupational noise exposure, n.d.c. <https://www.osha.gov/noise>. Accessed July 2026.
- Occupational Safety and Health Administration. Respiratory protection, n.d.d. <https://www.osha.gov/respiratory-protection>. Accessed July 2026.
- Open Food Facts. Open food facts product database, n.d. <https://world.openfoodfacts.org/data>. Accessed July 2026.
- OpenAI. New embedding models and API updates, 2024a. <https://openai.com/index/new-embedding-models-and-api-updates/>.
- OpenAI. GPT-4o system card, 2024b. <https://openai.com/index/gpt-4o-system-card/>.
- OpenAI. GPT-5 system card, 2025. <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*, pages 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*, pages 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: 10.1561/15000000019.
- 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.
- Parth Sarthi, Salman Abdullah, Aditi Tuli, Shubh Khanna, Anna Goldie, and Christopher D. Manning. Raptor: Recursive abstractive processing for tree-organized retrieval. *arXiv preprint arXiv:2401.18059*, 2024.
- Substance Abuse and Mental Health Services Administration. Recovery and recovery support, n.d. <https://www.samhsa.gov/find-help/recovery>. Accessed July 2026.
- Tesla. Supercharging, n.d. <https://www.tesla.com/support/charging/supercharging>. Accessed July 2026.
- U.S. Citizenship and Immigration Services. Travel documents, n.d. <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. <https://www.saferproducts.gov/RestWebServices/Recall?RecallID=10745&format=json>. Accessed July 2026.
- U.S. Consumer Product Safety Commission. Recalls, n.d.a. <https://www.cpsc.gov/Recalls>. Accessed July 2026.
- U.S. Consumer Product Safety Commission. Small parts and choking hazard labeling faqs, n.d.b. <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. <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. <https://www.dol.gov/agencies/whd/fmla/faq>. Accessed July 2026.
- U.S. Food and Drug Administration. Food allergies, n.d. <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. <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, pages 5776–5788, 2020.
- Wikipedia contributors. Gelatin, n.d.a. <https://en.wikipedia.org/wiki/Gelatin>. Accessed July 2026.
- Wikipedia contributors. Islamic dietary laws, n.d.b. 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. <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. <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. <https://arxiv.org/abs/2511.18423>.
- Shi-Qi Yan, Jia-Chen Gu, Yun Zhu, and Zhen-Hua Ling. Corrective retrieval augmented generation. *arXiv preprint arXiv:2401.15884*, 2024.
- 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.

Appendix

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

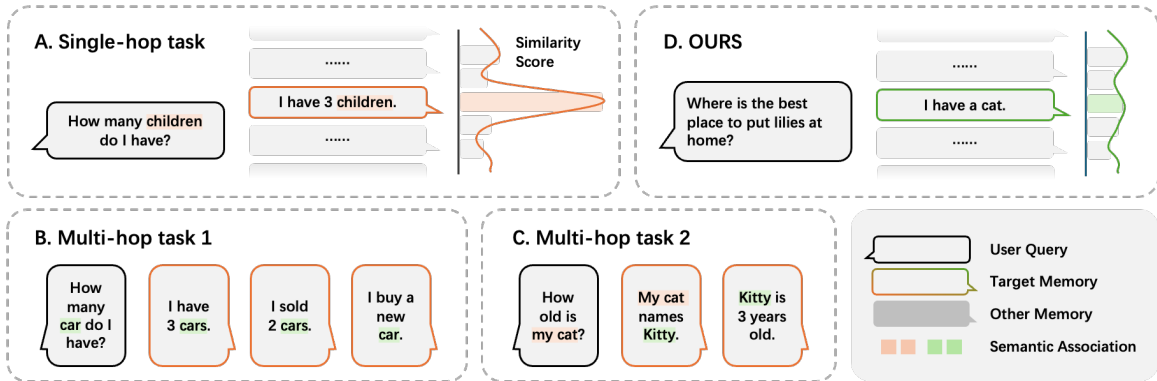


Figure 5 Memory-evaluation settings: explicit recall, state aggregation, entity-based multi-hop reasoning, and InMind’s knowledge-mediated implicit application.

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 naive 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 1 1/2 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). **Naive 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": "...",
    "naive_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 object 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. **naive_query** is a direct-recall question whose answer is the personal fact and which 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** $\rightarrow$ **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 $\rightarrow$ soft-boiled eggs, chronic kidney disease $\rightarrow$ ibuprofen, a job requiring a tight-fitting respirator $\rightarrow$ growing a beard, and electrician $\rightarrow$ electrical-troubleshooting training. Each example demonstrates a personal memory injection, a direct recall 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 a binary, context-aware LLM judge for both evaluation turns. The judge receives the labeled input fields shown below. In the naive condition, the fields are **user_message**, **context**, **query**, and **answer**. In the indirect condition, the same fields are provided together with **explanation**, the source-grounded behavior required for a correct answer.

Naive-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.

Indirect-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**: 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.

Target-recall judge prompt. We separately judge whether retrieval made the target fact available, without exposing the generated answer to the judge. This makes target 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 Post-Hoc Evaluation

As a diagnostic check, we re-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 does not score whether the memory was present in the answerer’s context. Table 4 reports the unmodified binary-judge results over the 125 tasks for every configuration in the main table.

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 answer-only judgments on the 125-task evaluation set.

Configuration	Correct	Accuracy (%)
Backbone (GPT-5-mini)	107 / 125	85.6
Naive RAG, MiniLM	27 / 125	21.6
Naive RAG, emb3-large	26 / 125	20.8
Naive RAG, BM25	23 / 125	18.4
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

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 naive-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 all 16 main-table configurations and all ten domains, and includes both positive and negative model verdicts. Because the backbone control has no naive-retrieval evaluation, the Naive 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. Target Recall is the most reliable judgment at 97.0% accuracy. The original context-aware Application 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
Naive	84 / 93	90.3	0	9
Target Recall	97 / 100	97.0	0	3
Original Application	85 / 100	85.0	0	15
Answer-only	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 Target 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 a context-aware Application 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, Answer-only correctly measures the behavior present in the response, while an end-to-end Application score of 1 would incorrectly attribute that behavior to successful memory use. This mechanism also explains how measured Application can exceed Target 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. Stateful memory systems use the shared middle-injection protocol in Appendix H; the single-shot Naive RAG control instead uses a temporary target chunk, as 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 and 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; naive and indirect query evaluated separately.
Answer model	GPT-5-mini (OpenAI, 2025), <code>max_completion_tokens=16,384</code> . The system receives its actual memory/retrieval context and the test query.
Judge model	GPT-5-mini (OpenAI, 2025), <code>max_completion_tokens=4,096</code> . 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.
Long-horizon trace	Fixed 47-session LME-s background and middle injection; 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 <code>emb3-large</code> run uses 20 independent task workers.

G.2 Backbone Control

Table 7 Backbone-control configuration.

Setting	Value
External memory	None.
Visible context	Target <code>user_message</code> , target <code>assistant_message</code> , 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 <code>user_message</code> and <code>assistant_message</code> as one temporary 487th chunk.
Retrieval unit	Complete raw turn chunk.
Retrievers	Dense MiniLM, dense emb3-large, and BM25.
Retrieved context	Top-5 chunks, concatenated with separators; 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
Memory bank	486 LME-s turns (486 chunks; 4,994 sentence units), with a per-task target-memory insertion.
Embeddings	MiniLM or emb3-large.
Agent tools	<code>keyword_search</code> , <code>semantic_search</code> , and <code>read_chunk</code> .
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.
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	486-note prebuilt bank, constructed with GPT-4o-mini.
Embeddings	MiniLM or emb3-large.
Task injection	Raw <code>user_message+assistant_message</code> note appended to a read-only bank.
Retrieval	Cosine similarity over bank notes and the injected note; 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	Copied prebuilt memory bank; raw target user/assistant passage indexed per task.
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

All memory systems use the same fixed 47-session LME-s trace as their background conversation. The injection point is fixed at the middle position: we first process the first eight sessions (the prefix before the session containing the 40th user turn), append each task’s two-message user-fact/assistant-acknowledgement pair to the end of the ninth session, and process that complete session. We then replay the remaining 38 sessions in their original order. Thus every task has the same background history, insertion position, and post-insertion interference; only the injected memory turn changes.

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 $\mathcal{C} = (C_1, \dots, C_T)$, injection session j , task memory turn (u, a) , test query q

```
1:  $M \leftarrow \emptyset$ 
2: for  $i = 1$  to  $T$  do
3:   if  $i = j$  then
4:      $C_i \leftarrow C_i \parallel [(\text{user}, u), (\text{assistant}, a)]$ 
5:   end if
6:    $M \leftarrow \text{Truncate}(\text{Update}(M, C_i))$ 
7: end for
8: 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)