

# ShardMemo: Scope-Before-Routing for Agentic Memory Retrieval

Yang Zhao<sup>1,\*</sup>,† Chengxiao Dai<sup>2,\*</sup> Mengying Kou<sup>3</sup>

Yue Xiu<sup>4</sup> Dusit Niyato<sup>3</sup>

<sup>1</sup>Independent Researcher <sup>2</sup>The University of Sydney

<sup>3</sup>Nanyang Technological University

<sup>4</sup>University of Electronic Science and Technology of China

\*Equal contribution. †Corresponding author: S180049@e.ntu.edu.sg

## Abstract

Agentic systems accumulate persistent memory across sessions, tools, and tasks, and a later request must retrieve from it under two distinct constraints: which memories it is permitted to access, and which are relevant under a limited search budget. Existing memory systems handle access scope in two flawed ways: applying scope *after* retrieval wastes probe budget on inadmissible memories, while treating scope as a learned ranking feature makes a hard constraint depend on router quality. We present SHARDMEMO, an agentic memory system built on *scope-before-routing*: metadata predicates first identify the admissible shards, and a learned router then selects a small number of them for shard-local approximate nearest neighbor retrieval. Given the supplied scope predicate and metadata, this separates hard admissibility from learned relevance ranking, so inadmissible shards cannot consume shard-probe budget. We evaluate on LoCoMo, HotpotQA, and ToolBench, covering conversational, long-context, and procedural memory. Under matched supervision and fixed budgets, SHARDMEMO improves over a learned router baseline by roughly +3 F1 on LoCoMo; in end-to-end LoCoMo evaluation it improves over the strongest evaluated memory baseline by up to +6.8 F1, with gains on HotpotQA and ToolBench.

## 1 Introduction

Agentic systems built on large language models (LLMs) increasingly rely on persistent memory across sessions, tools, and tasks. These memories include dialogue history, user preferences, tool outputs, task trajectories, and reusable procedures; later requests may depend on information recorded weeks earlier or in another session (Packer et al., 2023; Park et al., 2023; Maharana et al., 2024; Wang et al., 2023; Qian et al., 2026). We refer to retrieval over these persistent stores as *agentic memory retrieval*.

Agentic memory retrieval is not ordinary semantic retrieval. A memory may be relevant to the current task but inadmissible under the request’s user, session, time, tool namespace, or permission state. Relevance can be learned from retrieval signals, whereas admissibility is metadata-defined and must be enforced as a hard constraint.

This creates a retrieval-ordering problem. Memory systems that layer admissibility onto a semantic retriever apply scope only *after* retrieval—as a post-filter on returned evidence, so limited search budget may already have been spent on memories that were never admissible. The alternative, treating scope as one feature among others in a learned ranker, makes a hard access constraint depend on router quality: a routing error can return inadmissible memory. Neither is acceptable when scope is a hard constraint and budget is fixed. Our focus is how an agentic memory system should allocate a small retrieval budget when access scope must be enforced relative to a supplied policy.

We introduce SHARDMEMO, an agentic memory system centered on *scope-before-routing*. Metadata predicates first construct the admissible shard set, and a learned masked router then selects a small number of admissible shards for shard-local approximate nearest neighbor (ANN) retrieval. Unlike predicate-filtered ANN, which enforces constraints inside a single index, SHARDMEMO composes hard predicate masking with a budget-constrained router over heterogeneous shards, so inadmissible shards never consume probe budget while the router still learns which admissible shards are most likely to contain useful evidence.

SHARDMEMO applies this principle to three memory roles common in agentic workflows: active context for interpreting the current request, declarative memory for long-horizon evidence retrieval, and procedural memory for reusable tool-use skills. We evaluate the same scoped retrieval interface on LoCoMo, HotpotQA, and Tool-

Bench, covering long-horizon conversational memory, long-context multi-hop retrieval, and procedural skill reuse (Maharana et al., 2024; Yang et al., 2018; Qin et al., 2024). We test whether enforcing admissibility before routing improves retrieval under fixed budgets, against prior memory systems and controlled routing baselines.

**Contributions.** This paper makes three contributions:

- We formulate agentic memory retrieval as scope-constrained, budgeted shard selection.
- We introduce SHARDMEMO, a scope-before-routing memory system composing hard predicate masking with a budget-constrained router over admissible shards.
- We evaluate the same scoped retrieval interface, showing through controlled routing, noise, scaling, and ablation analyses where scope-before-routing improves the quality-efficiency frontier.

## 2 Background and Related Work

**Agentic memory.** Retrieval-augmented generation grounds language models in external evidence (Lewis et al., 2020; Karpukhin et al., 2020). Agentic memory systems extend this to persistent state across interactions: MemGPT separates working context from external memory (Packer et al., 2023), Generative Agents store profiles and episodic records (Park et al., 2023), and Reflexion reuses prior trajectories through self-reflection (Shinn et al., 2023). Recent systems such as A-Mem, LightMem, Mem0, MemoryOS, and GAM improve how memories are organized, summarized, consolidated, or reused (Xu et al., 2025; Fang et al., 2025; Chhikara et al., 2025; Kang et al., 2025; Yan et al., 2025; Wang et al., 2026; Liu et al., 2026a; Zhang et al., 2026).

**Heterogeneous memory for agents.** Agentic memory is not a single uniform store. Stable user facts, turn-level observations, multi-turn summaries, working context, tool outputs, and reusable tool-use procedures serve different roles and carry different access patterns; tool-using agents in particular benefit from reusing procedures and workflows (Yao et al., 2022; Schick et al., 2023; Wang et al., 2023, 2024; Liu et al., 2026b). SHARDMEMO reflects this heterogeneity with bounded active context, sharded declarative evidence, and a

separate procedural skill library, rather than forcing all memory types into one retrieval layout.

**Constrained and budgeted retrieval.** ANN systems make vector search efficient (Johnson et al., 2021; Guo et al., 2020; Malkov and Yashunin, 2020; Jayaram Subramanya et al., 2019); filtered ANN improves how to search an index under metadata predicates (Gollapudi et al., 2023; Jin et al., 2024), and federated search studies how to select collections under resource limits (Callan, 2002; Shokouhi and Si, 2011). Agentic memory adds a further demand: request-level admissibility over persistent, heterogeneous stores, where a semantically relevant memory may be inaccessible because of user, session, time, tool namespace, or permission state. SHARDMEMO constructs the admissible shard set from metadata first, then performs relevance- and cost-aware routing only within that set, and remains compatible with filtered ANN inside each shard. This addresses a question filtered ANN and federated search do not: how an agent should allocate a small per-request probe budget across admissible memory shards.

## 3 Method

### 3.1 Overview

SHARDMEMO is an agentic memory system for scope-constrained, budgeted retrieval. It implements *scope-before-routing*: metadata-defined admissibility first determines which memory shards a request may access, and learned routing then allocates probe budget only within that admissible set. SHARDMEMO instantiates this principle through Tier A active context, Tier B sharded declarative evidence retrieval, and Tier C reusable skill retrieval with fallback to Tier B.

### 3.2 Problem Formulation

We formalize each memory read as budgeted retrieval under hard scope constraints. Let  $q$  be a request. The system may return Tier A active context  $A(q)$ , Tier B evidence  $R^B(q)$ , and Tier C skills  $\mathcal{U}(q)$ ; let  $\mathcal{P}(q)$  denote the probed Tier B shard set. Each tier  $\tau \in \{A, B, C\}$  defines a metadata predicate

$$\psi^\tau(m, q) \in \{0, 1\},$$

which determines whether metadata  $m$  is admissible for request  $q$ . For a request sequence, we write  $q_t$  for the request at time  $t$  and  $\psi_t^\tau(m) = \psi^\tau(m, q_t)$ .




Figure 1: SHARDMEMO read path. Tier A supplies active context, a tier gate activates Tier B and/or Tier C, and Tier B applies the eligibility mask before routing. Tier C falls back to Tier B when no executable skill applies.

For Tier B, let  $S$  be the number of shards. Each shard  $j \in \{1, \dots, S\}$  has metadata  $m_j$ , inducing the admissible shard set

$$\mathcal{S}_B(q) = \{j : \psi^B(m_j, q) = 1\}.$$

Under budgets  $(M, B_{\text{probe}}, K, R)$ , which cap Tier A context, probed Tier B shards, returned Tier B evidence, and retrieved Tier C skills, the read path must satisfy

$$\begin{aligned} |A(q)| &\leq M, & |P(q)| &\leq B_{\text{probe}}, \\ |R^B(q)| &\leq K, & |U(q)| &\leq R, \\ \mathcal{P}(q) &\subseteq \mathcal{S}_B(q). \end{aligned}$$

The router ranks within the admissible set rather than deciding which shards are allowed, so routing errors can reduce recall within that set but cannot spend probe budget outside the request’s scope.

### 3.3 Tier A: Active Context

Tier A stores bounded active context: session metadata, recent turns, plans, and tool outputs. Its role is query interpretation, not large-scale retrieval. For example, Tier A can resolve temporal or context-dependent references such as “last week” or “that plan” before Tier B searches long-term evidence:

$$A_t = \text{ReadA}(q_t, \psi_t^A), \quad |A_t| \leq M.$$

### 3.4 Tier B: Scope-Before-Routing for Evidence Retrieval

Tier B stores declarative evidence in  $S$  shards, each with a shard-local ANN index  $\text{ANN}_j$ . Shards are

organized into profile shards for stable entity facts, observation shards for turn-level records, and session shards for multi-turn summaries. This layout separates evidence by granularity and access cost.

**Eligibility masking.** For request  $q_t$ , the admissible shard set is

$$\mathcal{S}_t = \mathcal{S}_B(q_t) \subseteq \{1, \dots, S\}.$$

All shards outside  $\mathcal{S}_t$  are masked to logit  $-\infty$  before routing, so they cannot be selected or consume probe budget.

**Budgeted routing and retrieval.** Let  $z_t = \text{Embed}(q_t)$  be the query embedding and  $r_t = [z_t; \varphi_t]$  be the router input, where  $\varphi_t$  contains request-time metadata features such as session, time, tool namespace, permission state, and budget flags. Within  $\mathcal{S}_t$ , the learned shard scorer  $f_\theta$  scores each admissible shard:

$$s_{t,j} = f_\theta(r_t, \sigma_j) - \alpha c_{t,j}, \quad j \in \mathcal{S}_t,$$

where  $\sigma_j$  is a shard summary,  $\alpha$  is the cost-bias coefficient, and  $c_{t,j} = n_j/\bar{n}$  is a latency proxy based on the number of vectors  $n_j$  in shard  $j$  and the mean shard size  $\bar{n}$ . Router probabilities are computed only over admissible shards:

$$p_{t,j} = \text{softmax}(\{s_{t,k}\}_{k \in \mathcal{S}_t})_j.$$

The router selects  $\mathcal{P}_t \subseteq \mathcal{S}_t$  with  $|\mathcal{P}_t| \leq B_{\text{probe}}$  using  $\text{Top-}B_{\text{probe}}$ ; we also evaluate an adaptive  $\text{Top-}P$  efficiency variant capped by  $B_{\text{probe}}$ .

Let  $L$  be the per-shard ANN candidate cap. Candidates from probed shards are merged, checked again at the item level, and globally ranked:

$$\begin{aligned}\tilde{\mathcal{C}}_t &= \bigcup_{j \in \mathcal{P}_t} \text{ANN}_j(z_t, L), \\ \mathcal{C}_t &= \{x \in \tilde{\mathcal{C}}_t : \psi_t^B(m(x)) = 1\}, \\ R_t^B &= \text{TopK}_{x \in \mathcal{C}_t}^K \text{score}(z_t, x).\end{aligned}$$

**Router training.** When evidence-to-shard supervision is available, each query  $q_t$  is paired with a gold shard set  $G_t \subseteq \mathcal{S}_t$  containing admissible shards that store supporting evidence.

We optimize a multi-positive routing loss:

$$L_{\text{route}}(t) = -\log \sum_{j \in G_t} p_{t,j},$$

which rewards probability mass on any supporting shard. To align training with fixed-budget inference, we also use

$$L_{\text{route}}^{\text{BA}}(t) = -\log \sum_{j \in G_t} w_{t,j} p_{t,j}.$$

Here,  $w_{t,j} \in [0, 1]$  softly emphasizes gold shards near the  $B_{\text{probe}}$  probe window. When evidence-to-shard supervision is unavailable, SHARDMEMO uses no-label routing, such as cosine similarity to shard summaries, as an ablation. Algorithm 1 gives the Tier-B read procedure.

### 3.5 Tier C: Skill Memory

**Read.** Tier C filters by scope, keeps the latest in-scope version of each skill, and retrieves up to  $R$  skills:

$$\begin{aligned}\mathcal{W}_t &= \{s \in \mathcal{W} : \psi_t^C(m(s)) = 1\}, \\ \mathcal{W}_t^{\text{cur}} &= \left\{s \in \mathcal{W}_t : \text{ver}(s) = \max_{\substack{s' \in \mathcal{W}_t \\ \text{id}(s') = \text{id}(s)}} \text{ver}(s')\right\}, \\ \mathcal{U}_t &= \text{TopR}_{s \in \mathcal{W}_t^{\text{cur}}}^R \text{sim}(q_t, s).\end{aligned}$$

The executor orders  $\mathcal{U}_t$  by reliability and similarity, fills slots from  $q_t$  and  $\mathcal{E}_t$ , validates preconditions, executes, and validates the result. If no skill succeeds, ShardMemo falls back to Tier B with the original Tier-B budget.

**Write.** Tier C induces skills from successful tool-use traces. Given a trace, SHARDMEMO canonicalizes the procedure, abstracts constants into typed slots, checks whether the canonicalized procedure matches an existing skill, and either

reuses that skill identifier or creates a new one. It then synthesizes deterministic validation tests and stores the skill only if those tests pass. New successful traces for an existing procedure append a new versioned record rather than overwriting earlier validated behavior. This keeps procedural reuse separate from Tier B evidence retrieval while using the same scope-predicate interface. Appendix B.2 gives the read and write algorithms.

## 4 Experiments

### 4.1 Setup

**Benchmarks.** We evaluate SHARDMEMO on three public benchmarks. LoCoMo (Maharana et al., 2024) tests long-horizon conversational memory across single-hop, multi-hop, temporal, and open-domain questions; we exclude adversarial questions because they primarily test instruction-following rather than memory retrieval. HotpotQA (Yang et al., 2018) tests long-context multi-hop evidence retrieval under the MemAgent protocol (Yu et al., 2025), using 56K/224K/448K-token inputs constructed from gold documents plus distractors. ToolBench (Qin et al., 2024) tests procedural skill reuse with in-domain and family-level out-of-domain splits.

**Scope predicates.** Scope-before-routing matters when the admissible shard or skill set is a request-dependent subset of the memory store. On LoCoMo, the Tier-B predicate  $\psi^B$  admits shards matching the conversation and speaker(s) referenced by the request. On HotpotQA,  $\psi^B$  admits only document shards that appear in the request’s input context, filtering out the global distractor pool. On ToolBench, the Tier-C predicate  $\psi^C$  admits skills whose tool namespace is authorized for the request. These predicates are strict rather than no-ops: Figure 2 reports the fraction of memory units admitted per request and the fraction of queries for which the gold evidence or executable skill remains in scope. ToolBench gold-in-scope falls below 100% because some requests have no executable skill in the authorized namespace, motivating the Tier-C  $\rightarrow$  Tier-B fallback in Section 3.5.

**Metrics.** We report token-level F1 and BLEU-1 on LoCoMo, answer F1 on HotpotQA, and Precision@ $R$  plus StepRed on ToolBench. StepRed is the mean reduction in tool-use steps after a valid skill is reused. For retrieval efficiency, VecScan is the mean number of vectors scanned per query across the probed shards, and p95 la-




Figure 2: Admissible-set statistics under the scope predicates. Filtered unit: Tier-B shards (LoCoMo/HotpotQA), Tier-C skills (ToolBench). (a) Mean admitted % per request with p10-p90 whiskers. (b) Gold-in-scope %.

tency is measured from scope filtering through candidate ranking, excluding generator decoding.  $\text{ShardHit}@B_{\text{probe}}$  measures whether at least one probed shard contains gold evidence; we use it as a routing diagnostic rather than a full multi-hop evidence-recall metric. Full formal definitions of all metrics ( $\text{ShardHit}@B_{\text{probe}}$ ,  $\text{Precision}@R$ ,  $\text{StepRed}$ ,  $\text{VecScan}$ , and p95 latency) are given in Appendix D.

**Baselines.** For end-to-end memory retrieval, we compare against Vanilla LLM, centralized RAG, and five agent-memory systems: A-Mem (Xu et al., 2025), LightMem (Fang et al., 2025), Mem0 (Chhikara et al., 2025), MemoryOS (Kang et al., 2025), and GAM (Yan et al., 2025). For controlled Tier-B routing, we compare cosine-to-prototype routing, an untrained router, a simple learned router trained with the same evidence-to-shard supervision, and component ablations. For Tier-C skill retrieval, we compare recency, BM25, embedding similarity, trace kNN, and a static skill library.

**Comparison protocol.** Within each benchmark, we match prompts, output post-processing, and evidence or skill budgets where applicable. Retrieval baselines use the same evidence items and chunking policy. Centralized RAG uses the same embedding model and ANN configuration to search a single global index over the evidence store. Controlled Tier-B routing baselines additionally use the same shard layout, shard-local indexes, and per-request budgets as SHARDMEMO, isolating the routing policy; methods without scope-aware routing enforce admissibility only by post-filtering returned evidence, which tests retrieval order under fixed shard-probe budgets rather than serving

as a production filtered-vector-search benchmark. The LoCoMo routing comparisons use evidence-to-shard supervision where stated, so they should be read as supervised routing comparisons under matched budgets, not as training-free comparisons.

**Implementation and splits.** We use all-MiniLM-L6-v2 embeddings and FAISS-HNSW indexes. Tier-B retrieval is shard-local; centralized RAG uses a single global index. In SHARDMEMO, final evidence ranking combines BM25, dense similarity, and route score with weights 0.4/0.4/0.2. The router is a two-layer MLP trained with Adam. Unless otherwise stated, the answer generator is GPT-OSS-120B with temperature 0, and all results are mean over three random seeds with seed standard deviation shown in blue subscript. LoCoMo uses an 8/2 conversation-level split (1,226/314 QA pairs); ToolBench uses a 70/30 induction/evaluation split plus a family-level OOD split. We run all experiments on a single server without inter-request concurrency. Full hyperparameters, hardware, software, and dataset statistics are in Appendix C.

## 4.2 Main Results

**LoCoMo.** Table 1 reports the broad benchmark comparison on long-horizon conversational memory. SHARDMEMO achieves the best F1 and BLEU-1 in all four non-adversarial categories. Compared with GAM, the strongest evaluated memory baseline, SHARDMEMO improves F1 by +5.70 on Single-Hop, +5.11 on Multi-Hop, +6.82 on Temporal, and +6.03 on Open Domain. The Temporal gains are largest, matching the design role of Tier A active context and scoped Tier B retrieval.

In addition, LoCoMo gains are not tied to the 120B generator. With GPT-OSS-20B and Qwen3-32B backbones, SHARDMEMO improves over GAM by 6.76 and 6.34 average F1, respectively. Temporal gains are the most consistent across backbones (+11.75 and +11.50 F1), supporting the role of Tier A deictic anchoring and scoped Tier B retrieval. Full per-category backbone results are in Appendix E (Table 9).

**HotpotQA.** Table 2 evaluates the same scope-before-routing interface for long-context multi-hop evidence retrieval. SHARDMEMO reaches 64.37/62.65/58.79 F1 at 56K/224K/448K tokens, improving over GAM by +2.27/+1.73/+1.39. These gains are smaller than on LoCoMo, which is expected: as context budgets grow, the marginal

Table 1: Results on LoCoMo with the GPT-OSS-120B backbone.

| LoCoMo           |                                  |                                  |                                  |                                  |                                  |                                  |                                  |                                  |
|------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|
| Method           | Single-Hop                       |                                  | Multi-Hop                        |                                  | Temporal                         |                                  | Open Domain                      |                                  |
|                  | F1                               | BLEU-1                           | F1                               | BLEU-1                           | F1                               | BLEU-1                           | F1                               | BLEU-1                           |
| Vanilla LLM      | 44.72 $\pm$ 0.11                 | 35.28 $\pm$ 0.17                 | 27.10 $\pm$ 0.21                 | 21.45 $\pm$ 0.09                 | 22.86 $\pm$ 0.14                 | 17.02 $\pm$ 0.22                 | 14.25 $\pm$ 0.11                 | 11.86 $\pm$ 0.16                 |
| RAG              | 51.36 $\pm$ 0.28                 | 45.80 $\pm$ 0.31                 | 28.33 $\pm$ 0.17                 | 21.63 $\pm$ 0.26                 | 43.57 $\pm$ 0.23                 | 37.26 $\pm$ 0.34                 | 21.54 $\pm$ 0.19                 | 15.29 $\pm$ 0.21                 |
| A-MEM            | 44.38 $\pm$ 0.37                 | 36.42 $\pm$ 0.42                 | 26.81 $\pm$ 0.24                 | 20.45 $\pm$ 0.33                 | 46.20 $\pm$ 0.29                 | 37.14 $\pm$ 0.46                 | 11.86 $\pm$ 0.31                 | 11.10 $\pm$ 0.27                 |
| LIGHTMEM         | 42.30 $\pm$ 0.26                 | 37.77 $\pm$ 0.39                 | 30.28 $\pm$ 0.35                 | 25.49 $\pm$ 0.28                 | 41.48 $\pm$ 0.41                 | 36.82 $\pm$ 0.24                 | 17.13 $\pm$ 0.22                 | 13.81 $\pm$ 0.36                 |
| MEM0             | 45.35 $\pm$ 0.44                 | 37.95 $\pm$ 0.33                 | 38.60 $\pm$ 0.27                 | 26.54 $\pm$ 0.48                 | 47.64 $\pm$ 0.36                 | 41.27 $\pm$ 0.31                 | 28.32 $\pm$ 0.29                 | 20.68 $\pm$ 0.41                 |
| MEMORYOS         | 47.17 $\pm$ 0.32                 | 42.39 $\pm$ 0.25                 | 34.13 $\pm$ 0.39                 | 26.02 $\pm$ 0.30                 | 42.40 $\pm$ 0.22                 | 32.81 $\pm$ 0.37                 | 19.83 $\pm$ 0.27                 | 14.40 $\pm$ 0.34                 |
| GAM              | 58.38 $\pm$ 0.41                 | 53.00 $\pm$ 0.35                 | 41.17 $\pm$ 0.26                 | 33.06 $\pm$ 0.44                 | 59.52 $\pm$ 0.33                 | 52.70 $\pm$ 0.29                 | 34.10 $\pm$ 0.24                 | 26.85 $\pm$ 0.38                 |
| <b>ShardMemo</b> | <b>64.08<math>\pm</math>0.23</b> | <b>59.70<math>\pm</math>0.30</b> | <b>46.28<math>\pm</math>0.27</b> | <b>39.65<math>\pm</math>0.21</b> | <b>66.34<math>\pm</math>0.19</b> | <b>60.31<math>\pm</math>0.32</b> | <b>40.13<math>\pm</math>0.17</b> | <b>33.82<math>\pm</math>0.25</b> |

Table 2: Results on HotpotQA with the GPT-OSS-120B backbone.

| HotpotQA         |                                  |                                  |                                  |
|------------------|----------------------------------|----------------------------------|----------------------------------|
| Method           | F1                               |                                  |                                  |
|                  | 56K                              | 224K                             | 448K                             |
| Vanilla LLM      | 54.43 $\pm$ 0.14                 | 52.10 $\pm$ 0.19                 | 49.08 $\pm$ 0.13                 |
| RAG              | 52.08 $\pm$ 0.25                 | 50.75 $\pm$ 0.22                 | 51.34 $\pm$ 0.28                 |
| A-MEM            | 34.83 $\pm$ 0.36                 | 31.06 $\pm$ 0.42                 | 29.77 $\pm$ 0.31                 |
| LIGHTMEM         | 37.68 $\pm$ 0.33                 | 35.42 $\pm$ 0.41                 | 29.82 $\pm$ 0.37                 |
| MEM0             | 30.13 $\pm$ 0.39                 | 28.40 $\pm$ 0.45                 | 29.52 $\pm$ 0.34                 |
| MEMORYOS         | 28.82 $\pm$ 0.37                 | 23.56 $\pm$ 0.32                 | 22.18 $\pm$ 0.43                 |
| GAM              | 62.10 $\pm$ 0.34                 | 60.92 $\pm$ 0.29                 | 57.40 $\pm$ 0.38                 |
| <b>ShardMemo</b> | <b>64.37<math>\pm</math>0.21</b> | <b>62.65<math>\pm</math>0.23</b> | <b>58.79<math>\pm</math>0.26</b> |

benefit of precise shard selection decreases. We therefore treat HotpotQA as evidence that scoped routing transfers to long-context evidence retrieval.

**ToolBench.** Table 3 evaluates the procedural-memory path. Tier C achieves the best Precision@3 (0.97) and StepRed (2.01). Compared with the strongest baseline, Tier C improves Precision@3 by +0.09 and StepRed by +0.22. The StepRed gain indicates that retrieved skills are more often executable end-to-end, not merely similar in embedding space.

### 4.3 Quality–Efficiency Trade-off

We test with a controlled routing comparison (Table 6, router-baselines block) and an iso-budget sweep in Figure 3.

**Controlled routing comparison.** The router-baselines block of Table 6 compares routing policies under an identical shard layout and supervision where applicable. A simple learned router improves over cosine routing, showing that evidence-

Table 3: Tier-C skill retrieval on ToolBench ( $R = 3$ ) with the GPT-OSS-120B backbone.

| ToolBench            |                                 |                                 |
|----------------------|---------------------------------|---------------------------------|
| Method               | P@R $\uparrow$                  | StepRed $\uparrow$              |
| Recency              | 0.79 $\pm$ 0.02                 | 1.44 $\pm$ 0.07                 |
| BM25                 | 0.84 $\pm$ 0.02                 | 1.70 $\pm$ 0.08                 |
| Embedding similarity | 0.88 $\pm$ 0.02                 | 1.79 $\pm$ 0.07                 |
| Trace kNN            | 0.69 $\pm$ 0.04                 | 1.65 $\pm$ 0.11                 |
| Static skill library | 0.86 $\pm$ 0.02                 | 1.76 $\pm$ 0.08                 |
| <b>Tier-C Skills</b> | <b>0.97<math>\pm</math>0.01</b> | <b>2.01<math>\pm</math>0.04</b> |

to-shard supervision is useful. The full model further improves F1 while reducing VecScan and p95 latency, indicating that the full routing stack improves the quality–efficiency operating point beyond a simpler learned scorer.

**Iso-budget sweep.** Figure 3 sweeps the shard-probe budget  $B_{\text{probe}} \in \{1, 2, 4, 8\}$  at fixed  $K = 10$ . We use powers-of-two probe budgets to show the trend across small and large retrieval budgets; the main operating point,  $B_{\text{probe}} = 3$ , is reported in Table 6. Across budgets, SHARDMEMO achieves higher F1 at comparable or lower VecScan than similarity and recency routing. Centralized retrieval can recover evidence by searching broadly, but its cost grows sharply.

**Scope before vs. after routing.** Enforcing scope *before* routing, rather than post-filtering, is the source of both gains. Post-filtering returns the same admissible evidence but spends probe budget on inadmissible shards, lowering F1 from 54.21 to 51.68 by reducing within-scope recall and raising VecScan from 414 to 560 (Table 6).

Table 4: Shard-count scaling at  $B_{\text{probe}}=3$ ,  $K=10$ .  $\dagger =$  default.

| Shard-count scaling |             |             |            |           |            |
|---------------------|-------------|-------------|------------|-----------|------------|
| $S$                 | F1          | ShardHit    | VecScan    | p95       | Route(ms)  |
| 5                   | 51.5        | <b>0.93</b> | 580        | 86        | <b>0.2</b> |
| 10                  | 52.8        | 0.89        | 485        | 78        | 0.4        |
| 20                  | 53.7        | 0.85        | 438        | <b>74</b> | 0.7        |
| 40 $^\dagger$       | <b>54.2</b> | 0.82        | 414        | 76        | 1.2        |
| 80                  | 53.8        | 0.76        | <b>388</b> | 81        | 2.2        |




Figure 3: F1 (left,  $\uparrow$ ) and VecScan (right,  $\downarrow$ ) on LoCoMo across  $B_{\text{probe}} \in \{1, 2, 4, 8\}$  at  $K=10$ . Darker is better.

#### 4.4 Component Ablations

Having established the routing mechanism in Section 4.3, we now isolate the contribution of each remaining component on LoCoMo under fixed budgets (Table 6). Heterogeneous shard families matter: collapsing Tier B to session-only shards lowers F1 to 51.12 and raises VecScan to 601. Removing Tier A active context lowers F1 to 50.38, consistent with its role in temporal and deictic grounding. Cost-aware gating and adaptive Top- $P$  mainly affect efficiency: disabling them changes F1 modestly but increases VecScan and p95 latency.

#### 4.5 Robustness and Scaling

All results use LoCoMo with  $B_{\text{probe}} = 3$  and  $K = 10$  unless otherwise stated.

**Shard-count scaling.** Table 4 varies the number of Tier-B shards. F1 remains within 2.7 points across  $S \in \{5, 10, 20, 40, 80\}$  and peaks at  $S = 40$ , our default setting. VecScan does not grow with the total number of shards; it remains controlled by the fixed probe budget and the size of the probed shards. Routing overhead increases with  $S$  but stays below 3 ms at  $S = 80$ . The drop in ShardHit at  $S = 80$  suggests that overly fine shard maps can fragment supporting evidence.

**Router comparison across scales.** Figure 4

Table 5: Label-noise robustness ( $B_{\text{probe}}=3$ ,  $K=10$ ); cosine = no-supervision floor.

| Label-noise robustness |       |          |         |     |
|------------------------|-------|----------|---------|-----|
| Noise $\epsilon$       | F1    | ShardHit | VecScan | p95 |
| 0% (clean)             | 54.21 | 0.82     | 414     | 76  |
| 20%                    | 52.14 | 0.79     | 431     | 79  |
| 40%                    | 49.87 | 0.73     | 462     | 85  |
| Cosine (no labels)     | 47.34 | 0.67     | 521     | 95  |




Figure 4: Router comparison across shard counts ( $B_{\text{probe}} = 3$ ,  $K = 10$ ). (a) F1 and (b) ShardHit for cosine, simple learned, and full routers at  $S \in \{20, 80\}$ .

crosses router type with shard count. The full router maintains its advantage over the simple learned router at both  $S = 20$  and  $S = 80$ , indicating that the family-aware routing stack does not degrade relative to simpler routers as the shard set grows.

**Label-noise robustness.** To test whether routing depends on clean shard labels, we corrupt only the training supervision by replacing each gold shard in  $G_t$  with a uniformly sampled incorrect admissible shard with probability  $\epsilon$ . Test labels remain clean, and the router is retrained from scratch at each noise level. Table 5 shows graceful degradation: even at 40% label noise, the trained router remains above the no-label cosine baseline.

#### 4.6 Case Study

Figure 5 illustrates the two Tier-C control paths on a real ToolBench trace (G1 #10) and a constructed out-of-scope query. **Top:** A customs-lookup request is gated to Tier C; SHARDMEMO retrieves a previously induced `transitaire_lookup` skill, canonicalizes the entity “ACT” via the Tier-A directory cache, validates the filled call against the skill’s stored tests, and executes against the real ToolBench endpoint—folding what would otherwise be a *list + lookup* plan into a single skill call.

Table 6: Ablations on LoCoMo ( $B_{\text{probe}}=3$ ,  $K=10$ ).

| Variant                                                  | F1 $\uparrow$                     | ShardHit $\uparrow$              | VecScan $\downarrow$         | p95 $\downarrow$            |
|----------------------------------------------------------|-----------------------------------|----------------------------------|------------------------------|-----------------------------|
| <i>Router baselines</i>                                  |                                   |                                  |                              |                             |
| Cosine similarity to shard prototype                     | 47.34 $\pm 0.22$                  | 0.67 $\pm 0.01$                  | 521 $\pm 7$                  | 95 $\pm 2$                  |
| Simple learned router (same supervision)                 | 51.14 $\pm 0.31$                  | 0.75 $\pm 0.01$                  | 472 $\pm 8$                  | 86 $\pm 2$                  |
| Untrained router (no supervision)                        | 39.82 $\pm 0.48$                  | 0.52 $\pm 0.02$                  | 627 $\pm 16$                 | 116 $\pm 5$                 |
| <i>Scope-first, no learned router</i>                    |                                   |                                  |                              |                             |
| Scoped Global ANN                                        | 53.74 $\pm 0.19$                  | —                                | 782 $\pm 27$                 | 113 $\pm 6$                 |
| Scoped All-Admissible (equal $K$ , probe all admissible) | 54.58 $\pm 0.17$                  | —                                | 1065 $\pm 38$                | 142 $\pm 8$                 |
| <i>Component ablations</i>                               |                                   |                                  |                              |                             |
| w/o cost-aware gating ( $\alpha=0$ )                     | 53.47 $\pm 0.29$                  | 0.80 $\pm 0.01$                  | 478 $\pm 12$                 | 86 $\pm 3$                  |
| w/o eligibility masking (post-filter only)               | 51.68 $\pm 0.37$                  | 0.73 $\pm 0.02$                  | 560 $\pm 21$                 | 100 $\pm 5$                 |
| w/o adaptive Top- $P$ (fixed Top- $B_{\text{probe}}$ )   | 53.85 $\pm 0.24$                  | 0.81 $\pm 0.01$                  | 465 $\pm 11$                 | 85 $\pm 3$                  |
| w/o heterogeneous shard families (session only)          | 51.12 $\pm 0.41$                  | 0.77 $\pm 0.02$                  | 601 $\pm 24$                 | 106 $\pm 6$                 |
| w/o Tier A active context                                | 50.38 $\pm 0.34$                  | 0.73 $\pm 0.02$                  | 509 $\pm 17$                 | 93 $\pm 4$                  |
| w/o budget-aware surrogate ( $L_{\text{route}}$ only)    | 53.07 $\pm 0.32$                  | 0.78 $\pm 0.01$                  | 428 $\pm 10$                 | 78 $\pm 3$                  |
| <b>ShardMemo (full)</b>                                  | <b>54.21<math>\pm 0.24</math></b> | <b>0.82<math>\pm 0.01</math></b> | <b>414<math>\pm 9</math></b> | <b>76<math>\pm 2</math></b> |

Note: ShardHit is ShardHit@ $B_{\text{probe}}$ ; it is undefined (—) for scope-first baselines that do not select a bounded shard set. VecScan is the mean number of vectors scanned per query. p95 is the 95th-percentile retrieval latency (ms).

#### Case Study: Tier-C Skill Retrieval with C $\rightarrow$ B Fallback

**Task:** Get contact details of the ACT customs agency in New Caledonia (ToolBench G1 #10).  
**Tier gate:** C (procedural)

**Tier C retrieval** ( $R=3$ , scope: transitaires):  
[skill:transitaire\_lookup\_v2]  
transitaire\_for\_transitaires(is\_id={id})  
(0.91 sim.)  
Slot filling: id="ACT\_AGENCE\_CALEDONIENNE\_DE\_TRANSIT" (canonicalized via Tier-A directory cache).

**Execution:** returns {codePostal: "98800", email: "sales@act.nc", phone: "+687 27.55.48"}.  
✓ Success — StepRed: 2 steps saved (list + lookup folded into one skill call).

**Task:** List freight forwarders serving Wallis-and-Futuna.

**Tier gate:** C (procedural)

**Tier C retrieval** ( $R=3$ , scope: transitaires):  
Max sim. 0.34 <  $\tau_C=0.5$  (territory mismatch).

**Execution: Failed (inapplicable).**

**C $\rightarrow$ B fallback:** Tier-B activated under same budgets ( $B_{\text{probe}}=3$ ,  $K=10$ ).

[doc:transitaires/spec] "Covers customs agencies in New Caledonia (country code NC) only."  
(API-spec shard)

**Prediction:** The transitaires API does not cover Wallis-and-Futuna; no freight forwarders are available via this namespace.

✓ Correct via fallback

**Bottom:** For a territory outside the API’s coverage, no skill clears  $\tau_C$ , so SHARDMEMO activates Tier B under the *same* budget and recovers the answer from the tool’s specification shard; a Tier-C miss never degrades below the Tier-B baseline. This discipline is what lets Tier C dominate Table 3 without sacrificing recall on out-of-scope requests. Appendix G gives additional Tier-B case studies.

## 5 Conclusion

We presented SHARDMEMO, a memory architecture for scope-constrained, budgeted agentic retrieval. The central idea is scope-before-routing: metadata predicates first mask inadmissible shards, and a learned router then allocates a small shard-probe budget only among admissible shards. Across conversational QA, long-context retrieval, and tool-use benchmarks, SHARDMEMO improves the quality–efficiency trade-off of memory retrieval under fixed budgets. The controlled ablations show that eligibility masking is central: replacing scope-before-routing with post-filtering lowers F1 and increases VecScan. Separate tiers for active context, declarative evidence, and procedural skills make the retrieval layer usable across different agent-memory roles without forcing all memory types into a single layout. Future work should evaluate dynamic permission changes, shard-map adaptation, concurrent multi-agent access, and distributed deployment.

Figure 5: Two Tier-C paths under shared budget ( $B_{\text{probe}}=3$ ,  $K=10$ ,  $R=3$ ): direct skill execution (top) and C $\rightarrow$ B fallback when no skill applies (bottom).

## Limitations

First, our evaluation is on English-only benchmarks; effectiveness under multilingual or cross-lingual workloads is untested. Second, Tier B router training requires evidence→shard annotations; without such supervision, the untrained router is significantly weaker (39.82 vs. 54.21 F1). The no-label cosine variant (47.34 F1) recovers roughly half this gap and offers a practical fallback. Third, experiments use a fixed shard map; workload drift is described architecturally but not evaluated empirically. Fourth, all experiments are conducted on a single server with no inter-request concurrency; multi-node deployment and concurrent multi-agent workloads are out of scope. Fifth, ToolBench evaluation uses held-out tool families but does not simulate continuous schema evolution; Tier C gains may depend on relative schema stability within the evaluation period. Sixth, our shard-count study covers  $S \in \{5, 10, 20, 40, 80\}$ ; distributed sharding at production scale ( $S \gtrsim 10K$ ) is not evaluated. Seventh, scope enforcement is conditional on externally supplied predicates and metadata. False exclusions cap recall; false inclusions waste budget and, when item metadata are also incorrect, can violate the intended policy despite item-level rechecking. Revocation, policy-conflict resolution, propagation, and audit logging remain external to SHARDMEMO.

## Ethics Statement

All experiments use publicly released benchmarks (LoCoMo, HotpotQA, ToolBench); LoCoMo contains synthetic personas rather than data about real individuals. SHARDMEMO itself does not collect or disclose personal data; its scope predicates are designed to compose with whatever access-control policies a deployment specifies. Operators should audit shard supervision and admissibility under their own data governance before production use.

## Use of AI assistants.

AI assistants were used to support language editing, LaTeX polishing, checklist preparation, and clarity checks. They were not used to generate experimental results, fabricate data, run evaluations, or make final scientific claims. The authors reviewed and verified all manuscript content.

## References

- Jamie Callan. 2002. Distributed information retrieval. In *Advances in information retrieval: Recent research from the Center for Intelligent Information Retrieval*, pages 127–150. Springer.
- Prateek Chhikara, Dev Khant, Saket Aryan, Taranjeet Singh, and Deshraj Yadav. 2025. Mem0: Building production-ready ai agents with scalable long-term memory. *arXiv preprint arXiv:2504.19413*.
- Jizhan Fang, Xinle Deng, Haoming Xu, Ziyang Jiang, Yuqi Tang, Ziwen Xu, Shumin Deng, Yunzhi Yao, Mengru Wang, Shuofei Qiao, Huajun Chen, and Ningyu Zhang. 2025. Lightmem: Lightweight and efficient memory-augmented generation. *arXiv preprint arXiv:2510.18866*.
- Siddharth Gollapudi, Neel Karia, Varun Sivashankar, Ravishankar Krishnaswamy, Nikit Begwani, Swapnil Raz, Yiyong Lin, Yin Zhang, Neelam Mahapatro, Premkumar Srinivasan, Amit Singh, and Harsha Vardhan Simhadri. 2023. Filtered-DiskANN: Graph algorithms for approximate nearest neighbor search with filters. In *Proceedings of the ACM Web Conference 2023*, pages 3406–3416.
- Ruiqi Guo, Philip Sun, Erik Lindgren, Quan Geng, David Simcha, Felix Chern, and Sanjiv Kumar. 2020. Accelerating large-scale inference with anisotropic vector quantization. In *International Conference on Machine Learning*, pages 3887–3896. PMLR.
- Suhas Jayaram Subramanya, Fnu Devvrit, Harsha Vardhan Simhadri, Ravishankar Krishnaswamy, and Rohan Kadekodi. 2019. DiskANN: Fast accurate billion-point nearest neighbor search on a single node. *Advances in Neural Information Processing Systems*, 32.
- Yicheng Jin, Yongji Wu, Wenjun Hu, Bruce M Maggs, Xiao Zhang, and Danyang Zhuo. 2024. Curator: Efficient indexing for multi-tenant vector databases. *arXiv preprint arXiv:2401.07119*.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2021. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*, 7(3):535–547.
- Jiazhen Kang, Mingming Ji, Zhe Zhao, and Ting Bai. 2025. Memory OS of AI agent. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 25961–25970, Suzhou, China. Association for Computational Linguistics.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP)*, pages 6769–6781.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020.

- Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Bingchen Liu, Yuanyuan Fang, Lei Liu, Guangyuan Dong, Xing Fu, Yuanyuan Gao, Shuyue Wei, Xin Li, and Xiangtian Meng. 2026a. [ConMem: Contribution-aware memory for long-horizon manufacturing inspection logs](#). *arXiv preprint arXiv:2607.28126*.
- Yuanzhe Liu, Jingyuan Zhu, Yuchen Mo, Gen Li, Xu Cao, Jin Jin, Yifan Shen, Zhengyuan Li, Tianjiao Yu, Wenzhen Yuan, Fangqiang Ding, and Ismini Lourentzou. 2026b. PALM: Progress-aware policy learning via affordance reasoning for long-horizon robotic manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 28096–28110.
- Adyasha Maharana, Dong-Ho Lee, Sergey Tulyakov, Mohit Bansal, Francesco Barbieri, and Yuwei Fang. 2024. Evaluating very long-term conversational memory of LLM agents. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13851–13870.
- Yu A. Malkov and D. A. Yashunin. 2020. [Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(4):824–836.
- Charles Packer, Sarah Wooders, Kevin Lin, Vivian Fang, Shishir G Patil, Ion Stoica, and Joseph E Gonzalez. 2023. MemGPT: Towards LLMs as Operating Systems. *arXiv preprint arXiv:2310.08560*.
- Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22.
- Pin Qian, Su Wang, Yihang Chen, Qiaolin Yu, Xiaoyuan Wang, Zhitong Guo, Zhicheng Wang, and Junxian You. 2026. [Toward user-conditioned evaluation of personal LLM agents under temporal interventions](#). *arXiv preprint arXiv:2607.21635*.
- Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2024. [ToolLLM: Facilitating large language models to master 16000+ real-world APIs](#). In *The Twelfth International Conference on Learning Representations*.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. Toolformer: Language models can teach themselves to use tools. *Advances in Neural Information Processing Systems*, 36:68539–68551.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36:8634–8652.
- Milad Shokouhi and Luo Si. 2011. Federated search. *Foundations and Trends in Information Retrieval*, 5(1):1–102.
- Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2023. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*.
- Haoyu Wang, Guangyuan Dong, He Liang, Zijing Zhang, Jiachen Luo, Chuang Liu, Chao Xue, and Hao Tang. 2026. [MemGuard: Persisting verifier signals for LLM-agent memory governance](#). *arXiv preprint arXiv:2608.21867*.
- Zora Zhiruo Wang, Jiayuan Mao, Daniel Fried, and Graham Neubig. 2024. Agent workflow memory. *arXiv preprint arXiv:2409.07429*.
- Wujiang Xu, Zujie Liang, Kai Mei, Hang Gao, Juntao Tan, and Yongfeng Zhang. 2025. [A-Mem: Agentic memory for LLM agents](#). In *Advances in Neural Information Processing Systems*, volume 38, pages 17577–17604. Curran Associates, Inc.
- B.Y. Yan, Chaofan Li, Hongjin Qian, Shuqi Lu, and Zheng Liu. 2025. General agentic memory via deep research. *arXiv preprint arXiv:2511.18423*.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 conference on empirical methods in natural language processing*, pages 2369–2380.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*.
- Hongli Yu, Tinghong Chen, Jiangtao Feng, Jiangjie Chen, Weinan Dai, Qiyang Yu, Ya-Qin Zhang, Wei-Ying Ma, Jingjing Liu, Mingxuan Wang, and Hao Zhou. 2025. MemAgent: Reshaping long-context LLM with multi-conv rl-based memory agent. *arXiv preprint arXiv:2507.02259*.
- Haobo Zhang, Xutao Mao, Guangyuan Dong, Ziwei Li, Xuanbo Su, Kaijie Chen, Jing Yang, and Zheng Lin. 2026. [MemMark: State-evolution attribution watermarking for agent long-term memory systems](#). *arXiv preprint arXiv:2605.25002*.

## A Notation Table

| Symbol                                                   | Meaning                                                           |
|----------------------------------------------------------|-------------------------------------------------------------------|
| $q, q_t$                                                 | Request / query (at time $t$ ).                                   |
| $x, s$                                                   | Tier B evidence item; Tier C skill artifact.                      |
| $m(\cdot)$                                               | Metadata (scope keys, permissions, provenance).                   |
| $\psi^\tau(m, q)$                                        | Tier- $\tau$ scope predicate.                                     |
| $M, B_{\text{probe}}$                                    | Tier A and probe budgets.                                         |
| $K, R$                                                   | Tier B evidence and Tier C skill budgets.                         |
| $A_t$                                                    | Tier A active context.                                            |
| $R_t^B, \mathcal{U}_t$                                   | Tier B evidence; Tier C retrieved skills.                         |
| $\mathcal{W}, \mathcal{W}_t, \mathcal{W}_t^{\text{cur}}$ | Skill store; admissible skills; latest admissible skill versions. |
| $S, j$                                                   | Number of shards; shard index.                                    |
| $\text{ANN}_j$                                           | Shard-local ANN index.                                            |
| $\mathcal{S}_t$                                          | Eligible/admissible shard set.                                    |
| $\mathcal{P}_t$                                          | Probed shard set.                                                 |
| $\mathbf{z}_t, \mathbf{r}_t$                             | Query embedding; router input.                                    |
| $\varphi_t$                                              | Request-time metadata features used by the router.                |
| $\sigma_j$                                               | Shard summary.                                                    |
| $f_\theta$                                               | Router scorer.                                                    |
| $s_{t,j}, p_{t,j}$                                       | Shard logit; routing probability.                                 |
| $c_{t,j}, \alpha$                                        | Cost estimate; cost-bias weight.                                  |
| $\gamma, P_{\min}, P_{\max}$                             | Adaptive Top- $P$ parameters.                                     |
| $\tau_t$                                                 | Adaptive Top- $P$ threshold.                                      |
| $G_t$                                                    | Gold shard set.                                                   |
| $L_{\text{route}}$                                       | Set-likelihood routing loss.                                      |
| $L_{\text{route}}^{\text{BA}}$                           | Budget-aware routing loss.                                        |
| $w_{t,j}$                                                | Budget-aware gold weight.                                         |
| VecScan                                                  | Mean vectors scanned per query.                                   |
| $L$                                                      | Per-shard ANN candidate cap.                                      |
| $\mathcal{E}_t$                                          | Available context for Tier-C slot filling.                        |
| $\rho(s)$                                                | Tier-C skill reliability score.                                   |
| $\text{id}(s), \text{ver}(s)$                            | Skill identifier and version of skill record $s$ .                |

Table 7: Notation summary.

## B Algorithms

Let  $[S] = \{1, \dots, S\}$  denote the set of Tier-B shard indices. For request  $q_t$ , write  $\psi_t^\tau(m) = \psi^\tau(m, q_t)$ . The admissible Tier-B shard set is

$$\mathcal{S}_t = \mathcal{S}_B(q_t) = \{j \in [S] : \psi_t^B(m_j) = 1\},$$

where  $m_j$  is the metadata of shard  $j$ . The probed shard set  $\mathcal{P}_t$  must satisfy

$$\mathcal{P}_t \subseteq \mathcal{S}_t, \quad |\mathcal{P}_t| \leq B_{\text{probe}}.$$

$\text{ANN}_j(\mathbf{z}_t, L)$  returns at most  $L$  candidates from shard  $j$ .  $\text{TopK}^K$  and  $\text{TopR}^R$  return at most  $K$

### Algorithm 1 Tier B read: scope-before-routing and global Top- $K$

---

**Require:**  $\mathbf{z}_t, \mathbf{r}_t, \psi_t^B, (B_{\text{probe}}, K, L), \{\sigma_j, m_j, c_{t,j}, \text{ANN}_j\}_{j=1}^S$

**Ensure:** Evidence  $R_t^B$ , probed shards  $\mathcal{P}_t$

- 1:  $\mathcal{S}_t \leftarrow \{j \in [S] : \psi_t^B(m_j) = 1\}$
- 2: **if**  $\mathcal{S}_t = \emptyset$  **then**
- 3:   **return**  $(\emptyset, \emptyset)$
- 4: **end if**
- 5: **for all**  $j \in \mathcal{S}_t$  **do**
- 6:    $s_{t,j} \leftarrow f_\theta(\mathbf{r}_t, \sigma_j) - \alpha c_{t,j}$
- 7: **end for**
- 8:  $(j_1, \dots, j_{|\mathcal{S}_t|}) \leftarrow \text{argsort}_{j \in \mathcal{S}_t}^\downarrow s_{t,j}$
- 9: **if** USEADAPTIVETOP **then**
- 10:    $\{p_{t,j}\}_{j \in \mathcal{S}_t} \leftarrow \text{softmax}(\{s_{t,j}\}_{j \in \mathcal{S}_t})$
- 11:    $\tilde{\tau}_t \leftarrow P_{\min} + \gamma(1 - \max_{j \in \mathcal{S}_t} p_{t,j})$
- 12:    $\tau_t \leftarrow \text{clip}(\tilde{\tau}_t, P_{\min}, P_{\max})$
- 13:    $\hat{b}_t \leftarrow \min\{b : \sum_{\ell=1}^b p_{t,j_\ell} \geq \tau_t\}$
- 14:    $b_t \leftarrow \min(\hat{b}_t, B_{\text{probe}}, |\mathcal{S}_t|)$
- 15: **else**
- 16:    $b_t \leftarrow \min(B_{\text{probe}}, |\mathcal{S}_t|)$
- 17: **end if**
- 18:  $\mathcal{P}_t \leftarrow \{j_1, \dots, j_{b_t}\}$
- 19: **for all**  $j \in \mathcal{P}_t$  **in parallel do**
- 20:    $\mathcal{C}_{t,j} \leftarrow \text{ANN}_j(\mathbf{z}_t, L)$
- 21: **end for**
- 22:  $\mathcal{C}_t \leftarrow \{x \in \bigcup_{j \in \mathcal{P}_t} \mathcal{C}_{t,j} : \psi_t^B(m(x)) = 1\}$
- 23:  $R_t^B \leftarrow \text{TopK}_{x \in \mathcal{C}_t}^K \text{score}(\mathbf{z}_t, x)$
- 24: **return**  $(R_t^B, \mathcal{P}_t)$

---

evidence items and  $R$  skills, respectively. For a skill record  $s$ ,  $\text{id}(s)$ ,  $\text{ver}(s)$ , and  $\rho(s)$  denote its identifier, version, and reliability score.

### B.1 Tier B Read Algorithm

Algorithm 1 implements scope-before-routing in three stages: it first restricts the eligible shard set to those whose metadata satisfy  $\psi_t^B$ , then ranks the eligible shards by the cost-corrected router score  $s_{t,j} = f_\theta(\mathbf{r}_t, \sigma_j) - \alpha c_{t,j}$  (optionally truncating the probe list by adaptive Top- $P$  when the routing distribution is sharp), and finally retrieves shard-local ANN candidates from the probed shards before re-applying admissibility and a global Top- $K$  re-rank. Returning early on  $\mathcal{S}_t = \emptyset$  is what gives scope filtering its compute-saving effect: when no shard is admissible, neither the router nor any ANN index is invoked.

### B.2 Tier C Algorithms

Algorithms 2–3 define the Tier-C read and write operators. Algorithm 2 retrieves and attempts to execute an admissible skill: it filters the skill store to records satisfying  $\psi_t^C$  and keeps only the latest version per skill identifier (Algorithm 3 appends new versions without overwriting old ones, so older versions are retrieved only when they are the newest

version within scope). The current skills are ranked by query similarity to form  $\mathcal{U}_t$ , re-sorted by reliability  $\rho(s)$  with similarity as a tiebreaker, and the procedure walks down this list executing the first skill whose slot fills, preconditions, and post-execution result validation all succeed.

---

**Algorithm 2** Tier C read: retrieve and execute current admissible skills

---

**Require:** Request  $q_t$ ; Tier-C predicate  $\psi_t^C$ ; skill budget  $R$ ; context  $\mathcal{E}_t$ ; skill store  $\mathcal{W}$   
**Ensure:** Skills  $\mathcal{U}_t$ ; flag  $ok$ ; executed skill  $s^*$ ; result  $\hat{y}$

```

1:  $\mathcal{W}_t \leftarrow \{s \in \mathcal{W} : \psi_t^C(m(s)) = 1\}$ 
2: if  $\mathcal{W}_t = \emptyset$  then
3:   return  $(\emptyset, \text{false}, \emptyset, \emptyset)$ 
4: end if
5:  $\mathcal{W}_t^{\text{cur}} \leftarrow \{s \in \mathcal{W}_t : \text{ver}(s) = \max_{\substack{s' \in \mathcal{W}_t \\ \text{id}(s') = \text{id}(s)}} \text{ver}(s')\}$ 
6:  $\mathcal{U}_t \leftarrow \text{TopR}_{s \in \mathcal{W}_t^{\text{cur}}}^R \text{sim}(q_t, s)$ 
7:  $\mathcal{V}_t \leftarrow \text{sort}^\downarrow(\mathcal{U}_t; \rho(s), \text{sim}(q_t, s))$ 
8: for  $s \in \mathcal{V}_t$  do
9:    $(\text{slots}, \text{fillable}) \leftarrow \text{SLOTFILL}(q_t, s, \mathcal{E}_t)$ 
10:  if  $\text{fillable}$  and  $\text{VALIDATEPRECONDITIONS}(s, \text{slots})$  then
11:     $(ok, \hat{y}) \leftarrow \text{EXECUTE}(s, \text{slots})$ 
12:    if  $ok$  and  $\text{VALIDATERESULT}(s, \text{slots}, \hat{y})$  then
13:       $s^* \leftarrow s$ 
14:      return  $(\mathcal{U}_t, \text{true}, s^*, \hat{y})$ 
15:    end if
16:  end if
17: end for
18: return  $(\mathcal{U}_t, \text{false}, \emptyset, \emptyset)$ 

```

---

If Algorithm 2 returns  $ok = \text{false}$  and the tier gate permits declarative retrieval, the controller invokes Algorithm 1 with the original Tier-B budget.

Algorithm 3 induces a new versioned skill from a successful trace  $\tau_t$ : it canonicalizes the trace, abstracts constants into typed slots, deduplicates or mints a skill identifier, synthesizes deterministic regression tests, and (only on passing) appends a fresh version while preserving older ones so that Algorithm 2 can resolve them when no newer version is in scope.

## C Experimental Setup Details

**Hardware and software.** Ubuntu 22.04, Python 3.12, PyTorch 2.8.0, CUDA 12.8; 4× NVIDIA RTX 4090D, 18 vCPUs (AMD EPYC 9754), 128GB RAM, 500GB SSD. The default backbone for headline results is GPT-OSS-120B, accessed via the official API with temperature 0; the GPT-OSS-20B and Qwen3-32B variants used in the backbone sensitivity study (Appendix E) are run locally via Ollama, also at temperature 0.

**Implementation.** all-MiniLM-L6-v2 (384-d, L2-normalized) for embeddings. Shard-

---

**Algorithm 3** Tier C write: induce, validate, and version a skill

---

**Require:** Successful trace  $\tau_t$ ; skill store  $\mathcal{W}$   
**Ensure:** Updated skill store  $\mathcal{W}$

```

1:  $\text{proc} \leftarrow \text{CANONICALIZE}(\tau_t)$ 
2:  $(\text{proc}, \text{slots}) \leftarrow \text{ABSTRACTCONSTANTSINTOSLOTS}(\text{proc})$ 
3:  $\text{id} \leftarrow \text{DEDUPLICATEORCREATEID}(\text{proc}, \mathcal{W})$ 
4:  $\text{tests} \leftarrow \text{SYNTHESIZETESTS}(\tau_t, \text{proc}, \text{slots})$ 
5: if  $\text{VALIDATEINDUCTIONTESTS}(\text{proc}, \text{tests})$  then
6:    $\text{ver} \leftarrow \text{BUMPVERSION}(\text{id}, \mathcal{W})$ 
7:    $s_{\text{new}} \leftarrow \text{SKILLRECORD}(\text{id}, \text{ver}, \text{proc}, \text{slots})$ 
8:    $\text{ATTACH}(s_{\text{new}}, m(\tau_t), \text{tests})$ 
9:    $\mathcal{W} \leftarrow \mathcal{W} \cup \{s_{\text{new}}\}$ 
10: end if
11: return  $\mathcal{W}$ 

```

---

local retrieval uses FAISS HNSW with inner-product similarity; BM25Okapi over tokenized turns for lexical matching. Final ranking:  $0.4 \cdot s_{\text{BM25}} + 0.4 \cdot s_{\text{sem}} + 0.2 \cdot s_{\text{route}}$ . Router: 2-layer MLP (hidden 128, dropout 0.1) over a 768-d input that concatenates the 384-d query embedding with the 384-d shard summary; the shard-family identifier enters the scorer as a separate small learned embedding. Training: Adam, lr=  $10^{-4}$ , weight decay 0.01, cosine annealing, 50 epochs, batch 32. All reported results are the mean over three random seeds, with the seed standard deviation shown as a blue subscript in the result tables.

**Dataset splits.** LoCoMo: 8/2 conversation-level train/test (1,226/314 QA pairs). HotpotQA: 500 questions per context length (56K/224K/448K tokens), drawn from the dev distractor split following the MemAgent protocol; each query’s context is constructed by appending sampled distractor passages to the gold documents until the target token count is reached, and the same 500 questions are used across the three context sizes so that the only varying factor is the context window. ToolBench: in-domain 70/30 induction/evaluation; OOD by tool family holdout (Table 8).

Table 8: Dataset statistics.

| Dataset   | Split   | #Inst.    | Avg. Size |
|-----------|---------|-----------|-----------|
| LoCoMo    | Non-adv | 1,540 QAs | 16.6K tok |
| HotpotQA  | 56K     | 500       | 56K tok   |
|           | 224K    | 500       | 224K tok  |
|           | 448K    | 500       | 448K tok  |
| ToolBench | Train   | 256       | 3.9 steps |
|           | Test    | 92        | 4.1 steps |
|           | OOD     | 181       | 3.9 steps |

## D Evaluation Metrics

**F1 and BLEU-1.** We compute token-level F1 and unigram BLEU between the predicted answer and the gold answer, and report the average over queries.

**ShardHit@ $B_{\text{probe}}$ .** Let  $\mathcal{Q}$  denote the evaluation queries,  $\mathcal{P}_q$  the set of shards probed for query  $q$ , and  $\mathcal{G}_q$  the gold evidence shard set. We define

$$\text{ShardHit}@B_{\text{probe}} = \frac{1}{|\mathcal{Q}|} \sum_{q \in \mathcal{Q}} \mathbb{I}[\mathcal{P}_q \cap \mathcal{G}_q \neq \emptyset].$$

**Precision@ $R$  and StepRed.** Let  $\mathcal{T}$  denote the evaluation targets for skill retrieval,  $\text{Top}_R(t)$  the top- $R$  retrieved skills for target  $t$ , and  $\text{valid}(s, t)$  an indicator that skill  $s$  is valid for  $t$ . We define

$$\text{P@}R = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \frac{|\{s \in \text{Top}_R(t) : \text{valid}(s, t)\}|}{R}.$$

For valid targets  $\mathcal{T}_{\text{valid}}$ , StepRed measures the non-negative reduction in required steps:

$$\text{StepRed} = \frac{1}{|\mathcal{T}_{\text{valid}}|} \sum_{t \in \mathcal{T}_{\text{valid}}} \max\left(0, n_t^{\text{gold}} - n_t^{\text{skill}}\right),$$

where  $n_t^{\text{gold}}$  and  $n_t^{\text{skill}}$  are the number of steps in the gold and skill-assisted solutions, respectively.

**VecScan and latency.** VecScan is the mean number of vectors scanned per query, summed over all probed shards. Latency is measured per query from scope filtering through candidate ranking, excluding generator decoding; we report both the mean and the 95th-percentile latency (p95).

## E Backbone Sensitivity

Table 9 shows that SHARDMEMO consistently outperforms GAM across both generator backbones and all four LoCoMo query types. Macro-averaged over query types, SHARDMEMO improves F1 by 6.76 points with GPT-OSS-20B and 6.34 points with Qwen3-32B. The largest gains occur on temporal queries (+11.75 and +11.50 F1 points, respectively), consistent with the deictic-anchoring role of Tier A described in Section 3. Multi-hop gains are more backbone-sensitive (+5.22 versus +2.49 F1 points), suggesting that final multi-hop answer quality depends on generator capability in addition to retrieval quality.

**Study setup.** Only the answer generator is swapped; the router, embedding model, shard layout, and Tier-B/C pipeline use the default configuration from Section 4. The routing checkpoint trained on GPT-OSS-120B traces is reused without re-fitting, isolating the effect of generator capability. All backbones run at temperature 0 with the same prompts and three random seeds.

**Per-category trends.** Temporal queries depend on Tier A’s session anchor rather than generator capability, so both alternative backbones recover the same +11.5 to +11.8 temporal gain over GAM as GPT-OSS-120B (Table 1). Single-hop and multi-hop gains, in contrast, track backbone reasoning quality: at Qwen3-32B, GAM already reaches 42.81 multi-hop F1 (versus 39.81 at GPT-OSS-20B), compressing the headroom retrieval can recover and narrowing the multi-hop gain to +2.49 from +5.22. Scope-and-route helps most when retrieval is the bottleneck and less when residual error is generation-bound.

**Scope.** The two alternative backbones span model families (Qwen, GPT-OSS) rather than interpolate a single family’s capacity. The study supports cross-family transfer of the architectural benefit; a finer within-family scaling sweep and an extension to closed-API generators are left for future work.

## F Sensitivity Analysis

**Sweep protocol.** We vary one Adaptive Top- $P$  parameter at a time around the default operating point ( $\gamma=0.6$ ,  $P_{\min}=0.5$ ,  $P_{\max}=0.95$ ), holding the other two fixed. All sweeps use the LoCoMo test split with  $K=10$ ,  $B_{\text{probe}}=3$ ,  $S=40$ , and the GPT-OSS-120B backbone; F1 is the macro average over the four query types. The chosen default is marked with  $^\dagger$  and the column-best entry is in bold.

**Sensitivity of Adaptive Top- $P$  (Tables 10, 11, and 12).** Compute-side trends track the truncation rule directly: larger  $\gamma$  tightens the cutoff on sharp distributions (VecScan 448→378, probes 2.85→2.20); larger  $P_{\min}$  raises the floor and admits more probes on flat distributions (VecScan 372→458); larger  $P_{\max}$  loosens the cap (VecScan 388→435). Across all three sweeps F1 varies by at most 0.53 points while VecScan swings by 47–86 scanned vectors per query, so predictive quality is robust around the default and compute remains tunable. We pick  $P_{\min}=0.5$  over the slightly higher

Table 9: LoCoMo per-category F1 / BLEU-1 for SHARDMEMO vs. GAM under two alternative generator backbones. SHARDMEMO numbers are bold.

| LoCoMo with alternative backbones |           |              |              |              |              |              |              |              |              |
|-----------------------------------|-----------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| Backbone                          | Method    | Single Hop   |              | Multi Hop    |              | Temporal     |              | Open Domain  |              |
|                                   |           | F1           | BLEU-1       | F1           | BLEU-1       | F1           | BLEU-1       | F1           | BLEU-1       |
| GPT-OSS-20B                       | GAM       | 57.88        | 52.60        | 39.81        | 33.20        | 51.97        | 44.89        | 31.07        | 25.21        |
|                                   | SHARDMEMO | <b>61.02</b> | <b>57.34</b> | <b>45.03</b> | <b>38.69</b> | <b>63.72</b> | <b>57.60</b> | <b>37.98</b> | <b>32.04</b> |
| Qwen3-32B                         | GAM       | 59.02        | 53.83        | 42.81        | 34.60        | 52.70        | 44.83        | 31.28        | 25.47        |
|                                   | SHARDMEMO | <b>63.17</b> | <b>58.92</b> | <b>45.30</b> | <b>39.10</b> | <b>64.20</b> | <b>58.49</b> | <b>38.50</b> | <b>32.45</b> |

F1 at  $P_{\min}=0.6$  (54.28 vs. 54.21) because the 0.07-point F1 sacrifice buys  $\sim 26$  fewer scanned vectors per query and a lower p95.

**Scope.** These are one-dimensional sweeps; joint interactions among  $\gamma$ ,  $P_{\min}$ ,  $P_{\max}$  are not characterized here. The narrow F1 range along each axis ( $\leq 0.53$  points) suggests the joint surface is locally flat around the operating point.

**Shard-count scaling, router comparison, and label-noise robustness.** These experiments are reported in Section 4.5 of the main paper; we summarize the take-aways here. F1 is stable across  $S \in \{5, 10, 20, 40, 80\}$  at fixed  $B_{\text{probe}}=3$  because shard-local ANN returns at most  $L$  candidates per probed shard, so per-query VecScan is upper-bounded by  $B_{\text{probe}} \cdot L$  regardless of  $S$ , and routing overhead stays below 3 ms at  $S=80$ . The full router keeps a +2.9/+3.1 F1 lead over the simple learned router at  $S=20/S=80$ . Under training-label noise the trained router degrades gracefully and stays above the no-label cosine baseline even at 40% noise, which is the practically relevant regime since deployed shard labels usually come from heuristic assignment rather than gold annotation. This robustness arises from the set-likelihood training objective, which spreads gradient mass across all candidate gold shards instead of over-committing to a single noisy label.

Table 10: Sensitivity sweep over  $\gamma$  with  $P_{\min}=0.5$  and  $P_{\max}=0.95$ .  $\dagger$  marks the default; column-best in bold.

| Adaptive Top- $P$ sweep: $\gamma$ |              |            |           |             |
|-----------------------------------|--------------|------------|-----------|-------------|
| $\gamma$                          | F1           | VecScan    | p95       | Avg Probed  |
| 0.2                               | 54.02        | 448        | 81        | 2.85        |
| 0.4                               | 54.10        | 432        | 79        | 2.72        |
| 0.6 $\dagger$                     | <b>54.21</b> | 414        | 76        | 2.58        |
| 0.8                               | 54.08        | 395        | 73        | 2.38        |
| 1.0                               | 53.82        | <b>378</b> | <b>71</b> | <b>2.20</b> |

Table 11: Sensitivity sweep over  $P_{\min}$  with  $\gamma=0.6$  and  $P_{\max}=0.95$ .  $\dagger$  marks the default; column-best in bold.

| Adaptive Top- $P$ sweep: $P_{\min}$ |              |            |           |             |
|-------------------------------------|--------------|------------|-----------|-------------|
| $P_{\min}$                          | F1           | VecScan    | p95       | Avg Probed  |
| 0.3                                 | 53.75        | <b>372</b> | <b>70</b> | <b>2.12</b> |
| 0.4                                 | 53.98        | 393        | 73        | 2.35        |
| 0.5 $\dagger$                       | 54.21        | 414        | 76        | 2.58        |
| 0.6                                 | <b>54.28</b> | 440        | 80        | 2.78        |
| 0.7                                 | 54.22        | 458        | 82        | 2.90        |

Table 12: Sensitivity sweep over  $P_{\max}$  with  $\gamma=0.6$  and  $P_{\min}=0.5$ .  $\dagger$  marks the default; column-best in bold.

| Adaptive Top- $P$ sweep: $P_{\max}$ |              |            |           |             |
|-------------------------------------|--------------|------------|-----------|-------------|
| $P_{\max}$                          | F1           | VecScan    | p95       | Avg Probed  |
| 0.80                                | 53.88        | <b>388</b> | <b>72</b> | <b>2.28</b> |
| 0.85                                | 54.00        | 398        | 74        | 2.38        |
| 0.90                                | 54.12        | 406        | 75        | 2.48        |
| 0.95 $\dagger$                      | <b>54.21</b> | 414        | 76        | 2.58        |
| 1.00                                | 54.16        | 435        | 79        | 2.72        |

## G Case Studies

We accompany the main paper’s Tier-C case study (Fig. 5) with three case studies on LoCoMo conversation *conv-26* (Caroline and Melanie) that exercise the Tier-B side of SHARDMEMO: cross-session aggregation (Fig. 6), Tier-A deictic anchoring combined with Tier-B retrieval (Fig. 7), and multi-hop aggregation across observation shards (Fig. 8). All three cases share the same retrieval configuration ( $B_{\text{probe}}=3$ ,  $K=10$ ,  $S=40$ ); the question text, gold answer, evidence shard IDs, and dialogue quotations are taken verbatim from LoCoMo (Maharana et al., 2024). Together they cover the three Tier-B retrieval mechanisms SHARDMEMO is designed to support: single-utterance answer recovery (Case 1), cross-tier deictic anchoring (Case 2), and bounded multi-shard aggregation (Case 3).

### Case Study 1: Cross-session aggregation

**Question:** What kind of art does Caroline make?

**Gold:** abstract art

**Setting:** LoCoMo,  $B_{\text{probe}}=3$ ,  $K=10$ ,  $S=40$ .

**Tier-B:** 3/40 shards, 358 vectors, 42.7 ms. **Shard hit: True.**

[D17:13] Caroline: “I’ve been trying out abstract stuff recently. It’s kinda freeing, just putting my feelings on the canvas without too much of a plan.”

[D11:8] Caroline: “Representing inclusivity and diversity in my art is important to me.”

[D11:12] Caroline: “‘Embracing Identity’ is all about finding comfort and love in being yourself.”

**Prediction:** abstract art ✓ **Correct.**

*The answer-bearing token (“abstract”) sits in a Session-17 utterance while supporting shards on Caroline’s art practice live in Session 11; the router aggregates across distant sessions under  $B_{\text{probe}}=3$ .*

Figure 6: Single-hop fact extraction on LoCoMo conv-26. The router aggregates one answer-bearing shard with two topic-supporting shards from a different session under  $B_{\text{probe}}=3$ .

**Case 1 (Fig. 6): cross-session aggregation.** The gold answer (“abstract art”) is supplied by a single Session-17 utterance in which Caroline says she has been “trying out abstract stuff” (D17:13). Two earlier Session-11 utterances anchor the topic “Caroline’s art practice” (style, themes of identity and inclusivity) and serve as context-supporting shards. The case shows the router aggregating across distant sessions (Session 11 and Session 17) within a single budgeted call of  $B_{\text{probe}}=3$ ; the prediction matches the gold without any additional reasoning chain.

**Case 2 (Fig. 7): temporal reasoning needs both tiers.** The query asks for an absolute year, but the Tier-B shard records the information only as “last year” (D1:14); the absolute year is never written into evidence. Tier-A supplies the Session-1 anchor (8 May 2023) and Tier-B recovers the deictic phrase; resolution happens at composition time, “last year” → 2022. A retrieval-only baseline could still surface D1:14 by semantic similarity to “a sunrise”, but with no Tier-A anchor it has no way to convert “last year” into a year.

**Case 3 (Fig. 8): multi-hop aggregation under matched probe budget.** The gold answer is a three-token set (“thoughtful, authentic, driven”), each token attested in a different LoCoMo-labeled gold shard from a different session (D7:4, D13:16,

### Case Study 2: Temporal reasoning (Tier-A anchoring + Tier-B)

**Question:** When did Melanie paint a sunrise? **Gold:** 2022

**Setting:** LoCoMo,  $B_{\text{probe}}=3$ ,  $K=10$ ,  $S=40$ .

**Tier-A:** Session-1 date = 8 May 2023 (enables deictic resolution). **Tier-B:** 3/40 shards, 394 vectors, 52.3 ms. **Shard hit: True.**

[D1:14] Melanie: “Yeah, I painted that lake sunrise last year! It’s special to me.”

**Prediction:** 2022 ✓ **Correct.**

*Two further Session-1 shards were probed under  $B_{\text{probe}}=3$  but contained no temporal anchor and are omitted for brevity. Tier-A supplies the session anchor (8 May 2023); Tier-B recovers the deictic phrase. Resolution: “last year” → 2022 requires both tiers.*

Figure 7: Temporal reasoning on LoCoMo conv-26. Tier-A deictic anchoring and Tier-B retrieval combine to resolve a relative time reference.

### Case Study 3: Multi-hop aggregation across observation shards

**Question:** What personality traits might Melanie say Caroline has?

**Gold:** thoughtful, authentic, driven

**Setting:** LoCoMo,  $B_{\text{probe}}=3$ ,  $K=10$ ,  $S=40$ .

**Tier-B:** 3/40 shards, 386 vectors, 45.2 ms. **Shard hit: True (3/3 gold shards probed).**

[D7:4] Melanie: “Your drive to help is awesome!”

[D13:16] Melanie: “You really care about being real and helping others.”

[D16:18] Melanie: “... you’re so thoughtful!”

**Prediction:** thoughtful, authentic, driven ✓ **Correct.**

*Each probed shard supplies one answer token from a different session (D7, D13, D16); no single shard suffices, motivating global Top-K aggregation under the bounded probe budget.*

Figure 8: Multi-hop aggregation on LoCoMo conv-26. The router probes exactly the three shards LoCoMo marks as gold evidence, one per answer token, under  $B_{\text{probe}}=3$ .

D16:18). The router probes exactly the three shards that carry the answer tokens; global Top-K ranking over the union of probed shards then surfaces one supporting utterance per token. If  $B_{\text{probe}}$  were smaller than the gold cardinality, only the highest-scoring tokens would be recoverable, which is the budget-coverage tradeoff characterized quantitatively in Section 4.3 (Fig. 3).