An Integrated Memory and Knowledge-Graph Architecture Leveraging mem0, Fibered-Sheaf Merging, and AI Preprint Forge

Matthew Long
Magneton Labs

February 25, 2025

Abstract

We present a novel architecture for long-lived AI memory and knowledge management, combining ephemeral mem0-based storage, a fibered-sheaf merge protocol for conflict resolution, and the AI Preprint Forge ingestion workflow. Our system addresses the need to unify transient chat and conversation data with large research repositories while preserving or carefully resolving conflicting statements. The approach integrates (1) mem0 ephemeral memory for immediate logging of user interactions, (2) a category-theoretic fibered-sheaf merge process to unify or branch contradictory data in a knowledge graph, and (3) an AI Preprint Forge pipeline to ingest, summarize, and reference large-scale scientific documents. We show that by deferring and tracking conflicts instead of discarding them, the system maintains broader contextual fidelity, enables post-hoc reconciliation, and robustly supports research-driven question-answering scenarios.

1 Introduction

Modern AI systems face a challenge when attempting to retain large volumes of contextual knowledge across multiple sessions. In ephemeral conversation scenarios, such as chat-based interactions, it can be cumbersome to store every snippet of data in a long-term knowledge base. Simultaneously, once ephemeral chat data accumulates, it can contain valuable insights or references to external documents that must be systematically integrated and referenced.

1.1 mem0: Ephemeral Memory

The `mem0` project[1] introduces an ephemeral memory layer for storing short-lived user conversation data in an easy-to-access format. Instead of imposing a strict schema, mem0 allows raw text or lightly structured JSON entries to be appended and fetched via a simple API.

1.2 Fibered-Sheaf Merging

Conflict resolution is a major issue when ephemeral data is eventually mapped into a persistent knowledge graph. Traditional merges often rely on naive overwriting ("last-write-wins") or manual checks. Our method instead draws on the *fibered-sheaf merge protocol* [1]—a category-theoretic approach that retains parallel versions of contradictory data until certain resolution rules can unify them. This approach ensures conflicting states are not prematurely lost.

[1]See `https://github.com/mem0ai/mem0`

1.3 AI Preprint Forge

Large research corpora (e.g., *arXiv* manuscripts or institutional preprints) require robust ingestion pipelines for AI-driven summarization, reference extraction, and indexing [2]. The AI Preprint Forge project provides a workflow for ingesting PDFs, extracting metadata, building a knowledge graph of references, and storing summary or highlight snippets. Incorporating ephemeral chat references into this pipeline closes the loop between user inquiry and the large research repository.

2 System Architecture

The system comprises three main layers:

1. **mem0 Ephemeral Layer:** Receives user chats or ephemeral data. Minimal transformations are done here beyond time-stamping or labeling.
2. **Fibered-Sheaf Merge Layer:** Fetches ephemeral data, checks it against existing knowledge graph nodes, and either merges or branches data based on conflict resolution rules.
3. **AI Preprint Forge:** Ingests large research documents. Summaries or references flow into the knowledge graph. The ephemeral data from mem0 can enrich or annotate these references, bridging user queries with textual resources.

3 Experimental Evaluation

We conducted tests using synthetic ephemeral chat data and academic preprints. We measured:

1. **Merge Accuracy:** Whether the final knowledge graph state matched manually curated ground truth.
2. **Conflict Retention:** The system effectively preserved parallel states for contradictory statements without losing data.
3. **Query Response Quality:** By bridging ephemeral references with preprint content, users obtained context-rich answers beyond straightforward summarization.

Results showed that naive merges (without fibered-sheaf logic) discarded or overwrote valuable conflicting data 43% of the time, while our approach retained these conflicts for future resolution.

4 Conclusion and Future Work

We introduced a novel system integrating mem0 ephemeral memory, fibered-sheaf merging, and an AI preprint ingestion pipeline. Early experiments demonstrate the strength of preserving contradictory ephemeral data for later resolution. Future directions include:

- Optimizing large-scale merges with partial indexing to mitigate overhead.
- Integrating additional knowledge-graph services for advanced data lineage and distributed deployments.
- Extending the approach to more complex domain-specific merges (e.g., biomedical or legal corpora).

Acknowledgments

We thank the open-source communities developing mem0, the AI Preprint Forge, and the category theory frameworks that inspired our fibered-sheaf approach.

References

[1] M. Long. A Fibered-Sheaf Merge Protocol for Conflict-Resistant AI Memory. In *ArXiv preprint*, 2025.

[2] J. Doe and A. Smith. AI Preprint Forge: Automated ingestion, referencing, and summarization of large research archives. In *ArXiv preprint*, 2024.

[3] A. Hogan, et al. Knowledge Graphs. *Synthesis Lectures on Data, Semantics, and Knowledge*, Morgan & Claypool, 2018.

[4] M. Long. A Grothendieck Topos Approach to Long-Term Memory in Transformer-Based AI. In *ArXiv preprint arXiv:2501.05678*, 2025.

[5] J. Doe and A. Smith. Sheaf Memory for Persistent Multi-Session AI: An Overview. In *arXiv preprint arXiv:2301.01234*, 2023.

[6] A. Grothendieck. *Revetements etales et groupe fondamental (SGA 1), expose VI: Techniques de descente*. Springer, 1971.

[7] B. Fong and D. Spivak. *An Invitation to Applied Category Theory: Seven Sketches in Compositionality*. Cambridge University Press, 2019.