TransMem: Transition-Aware Retrieval for Evolving Personal Memory

Shigeng Chen[1*] Linhao Luo[2*] Changlong Shi[3] Zhangchi Qiu[1] Shirui Pan[1†]

[1] Griffith University [2] Monash University [3] The Hong Kong Polytechnic University

Paper | GitHub Code | Let's Connect!

NeurIPS 2026

Background & Motivation

Preference evolution

EARLIER PREFERENCE

Training on my own for a fast first marathon.

NEW CONSTRAINT

Injured: stopped chasing race times; want low-pressure exercise.

Saturday morning plan?

Easy neighborhood walk + coffee

Interleaved sessions

Past → Now

Fitness Routine

Volunteering Challenges

Fitness Routine

Reading & Books

Weekend Activities

Fragmented user context | Interleaved task topics | Dynamic preference drift

- **Memory Construction Challenges**

  Preference evolution is hard to model when evidence is scattered across **interleaved topics and multiple sessions**.

- **Memory Retrieval Challenges**

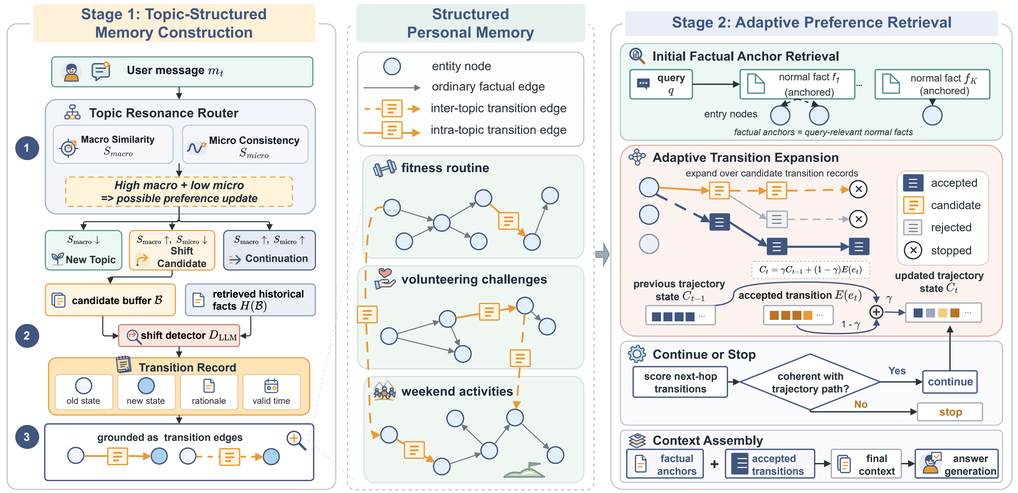
  Retrieving **isolated facts** can miss the preference **evolution trajectory**, especially when key updates do not match the query.

Transition record

Earlier state	Training for a fast marathon
Later state	Stay active without performance pressure
Reason	Injury changed the exercise goal
Time	4 Jan 2024

Record what changed, why it changed, and when it changed.

TransMem Framework



1. Topic-Structured Memory Management

A lightweight *Topic Resonance Router* organizes multi-topic dialogues into topic graphs.

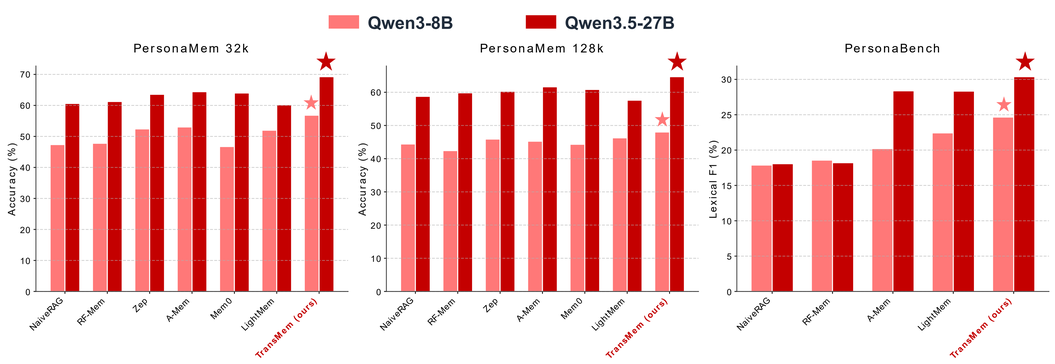
2. Explicit Transition Extraction & Graph Linking

An LLM *Shift Detector* extracts *Transition Records* grounded as intra- and inter-topic edges.

3. Trajectory-Guided Adaptive Retrieval

Expands from *factual anchors*, updating the *trajectory state* to guide expansion and stopping.

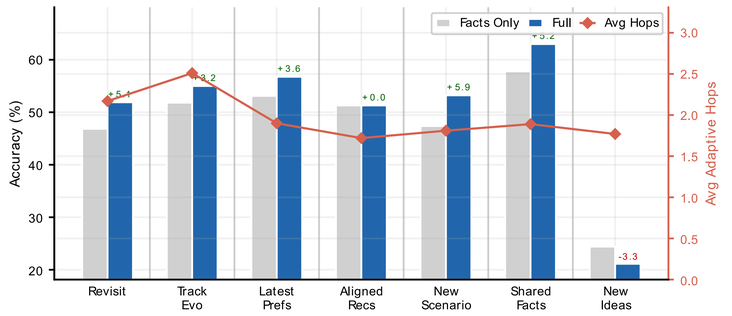
Main Experimental Results



Findings: TransMem achieves the best overall results among the compared methods across both backbones on PersonaMem and PersonaBench.

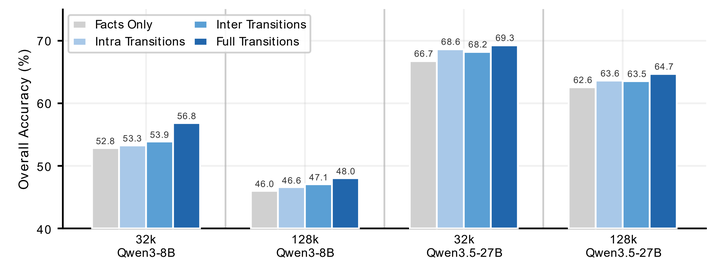
Ablation & Efficiency Analysis

Per-Category: Facts Only vs. Full (128k, Qwen3-8B)



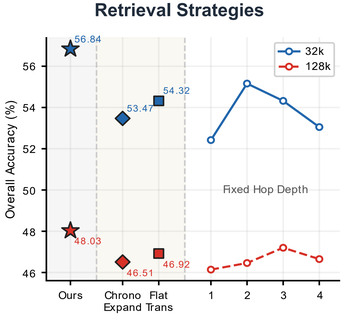
Findings: TransMem improves most categories and dynamically explores deeper paths for evolution-heavy tasks (Track Evo, Revisit).

Intra- vs. Inter-Topic Transitions

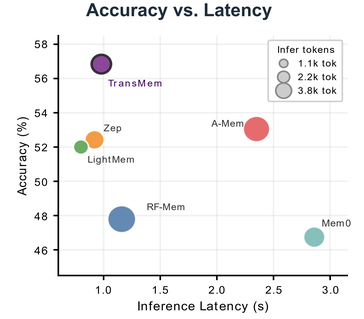


Findings: Both transitions beat Facts Only and perform best when combined.

Retrieval Strategies



Accuracy vs. Latency



Findings: Outperforms chronological, flat, and fixed-depth baselines; the strongest overall balance of accuracy, latency and inference context.