

# FlowKV vs. Traditional Cache Eviction in Long-Context Multi-Turn Conversations

Assignee Research

June 7, 2026

## Abstract

This report synthesises findings from 8 peer-reviewed papers addressing the following research question: How does FlowKV's isolated KV cache management compare to other cache eviction strategies (e.g., FIFO, LRU) in terms of LongBench accuracy and throughput for multi-turn conversations exceeding 64K. 13 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

## 1 Introduction

This paper examines: FlowKV: Enhancing Multi-Turn Conversational Coherence in LLMs via Isolated Key-Value Cache Management. Research question: How does FlowKV's isolated KV cache management compare to other cache eviction strategies (e.g., FIFO, LRU) in terms of LongBench accuracy and throughput for multi-turn conversations exceeding 64K tokens?.

## 2 Methodology

Systematic literature search across multiple databases yielded 8 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.5/10.

## 3 Results

8 papers retrieved. 13 claims extracted; 0 independently verified. Quality review score: 3.5/10.

## 4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

## 5 Extracted Claims

Claim	Verified	Confidence
On the PrefEval benchmark, the Full KV method achieves a score of 75.4.	×	0.04
On the PrefEval benchmark, the FlowKV method achieves a score of 58.7.	×	0.03
On the PrefEval benchmark, the Baseline method achieves a score of 10.6.	×	0.03
In Turn 1 responses, attention is heavily focused on the Turn 1 Query (T1Q) and the local window.	×	0.03
In Turn 2 responses, attention is heavily focused on T1Q, T1R, and the local window.	×	0.03
In Turn 3 responses, attention is heavily focused on T1Q, T1R, T2Q, T2R, and the local window.	×	0.02
Queries in Turns 2 and 3 show increased attention to previous queries and the system prompt.	×	0.02
For the LLaMA model using the SKV strategy, the Baseline Turn 2 performance is 37.08%.	×	0.03
For the LLaMA model using the SKV strategy, FlowKV improves Turn 2 performance by 24.85 percentage points compared to th	×	0.03
For the LLaMA model using the CKV strategy, FlowKV improves Turn 2 performance by 40.27 percentage points compared to th	×	0.04
FlowKV achieves an average performance improvement of over 20% in subsequent conversation turns compared to the baseline	×	0.04
During the initial turn of conversation, FlowKV’s core isolation mechanism is not engaged due to the absence of prior co	×	0.11
For the Qwen model using the SKV strategy, FlowKV achieves a Turn 2 score of 56.72%.	×	0.03

## References

- <http://arxiv.org/abs/2605.07234v1>
- <http://arxiv.org/abs/2503.08879v1>
- <http://arxiv.org/abs/2505.15347v2>