

# FlowKV and SnapKV Retrieval Performance on BEIR Under Memory Constraints

Assignee Research

June 8, 2026

## Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: How does the retrieval performance of FlowKV and SnapKV compare when evaluated on the BEIR benchmark for cross-domain tool retrieval tasks with Llama-3-8b under 2GB memory constraints. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.8/10. This report is a machine-generated literature synthesis and does not constitute original research.

## 1 Introduction

This paper examines: Make Each Token Count: Towards Improving Long-Context Performance with KV Cache Eviction. Research question: How does the retrieval performance of FlowKV and SnapKV compare when evaluated on the BEIR benchmark for cross-domain tool retrieval tasks with Llama-3-8b under 2GB memory constraints?.

## 2 Methodology

Systematic literature search across multiple databases yielded 14 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.8/10.

## 3 Results

14 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 3.8/10.

## 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.

## References

- <http://arxiv.org/abs/2605.09649v1>
- <http://arxiv.org/abs/2503.01763v2>
- <http://arxiv.org/abs/2605.07234v1>