FlowKV Selective Eviction Strategy and Its Impact on Perplexity-Accuracy Trade-offs in Long-Context Language Models

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 10 peer-reviewed papers addressing the following research question: How does FlowKV's selective eviction strategy affect the trade-off between perplexity and answer accuracy compared to LongNet and SmoothFormer when applied to Llama-3-70b on the LongBench suite with. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 5.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: LongBench Pro: A More Realistic and Comprehensive Bilingual Long-Context Evaluation Benchmark. Research question: How does FlowKV's selective eviction strategy affect the trade-off between perplexity and answer accuracy compared to LongNet and SmoothFormer when applied to Llama-3-70b on the LongBench suite with context lengths exceeding 500K tokens?.

2 Methodology

Systematic literature search across multiple databases yielded 10 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 5.7/10.

3 Results

10 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 5.7/10.

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/2601.02872v1
- http://arxiv.org/abs/2308.14508v2