Fine-Tuning Learnable Eviction in ReST-KV for Adversarial Long-Context Robustness

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 8 peer-reviewed papers addressing the following research question: Can the learnable eviction mechanism in ReST-KV be fine-tuned to improve robustness against adversarial long-context inputs (e.g., noisy or irrelevant tokens) while maintaining accuracy on HumanEval. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 6.7/10. 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: Can the learnable eviction mechanism in ReST-KV be fine-tuned to improve robustness against adversarial long-context inputs (e.g., noisy or irrelevant tokens) while maintaining accuracy on HumanEval and MBPP benchmarks?.

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 6.7/10.

3 Results

8 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 6.7/10.

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/2103.15670v3
- http://arxiv.org/abs/2605.08840v1