

ReST-KV Layer-Wise Reconstruction and Perplexity Scaling on LongEval Beyond 64k Tokens

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: How does ReST-KV's layer-wise output reconstruction impact perplexity degradation on the LongEval benchmark compared to SmootherMem when context lengths scale from 64k to 256k tokens. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 7.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: LV-Eval: A Balanced Long-Context Benchmark with 5 Length Levels Up to 256K. Research question: How does ReST-KV's layer-wise output reconstruction impact perplexity degradation on the LongEval benchmark compared to SmootherMem when context lengths scale from 64k to 256k tokens?.

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 7.5/10.

3 Results

14 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 7.5/10.

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.

References

- <http://arxiv.org/abs/2402.05136v3>
- <http://arxiv.org/abs/2311.11509v3>
- <http://arxiv.org/abs/2605.08840v1>