

Spatial-Temporal KV Cache Eviction for Long-Context Reasoning in Multi-Hop QA

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: Does the integration of spatial-temporal dynamics in KV cache eviction preserve long-context reasoning accuracy on multi-hop QA benchmarks better than static attention scoring methods in scaled LLM. 11 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: ReST-KV: Robust KV Cache Eviction with Layer-wise Output Reconstruction and Spatial-Temporal Smoothing. Research question: Does the integration of spatial-temporal dynamics in KV cache eviction preserve long-context reasoning accuracy on multi-hop QA benchmarks better than static attention scoring methods in scaled LLM architectures?.

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. 11 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.

5 Extracted Claims

Claim	Verified	Confidence
ReST-KV is evaluated on five open-source LLMs: Llama2-Chat, Gemma-Instruct, Llama3-Instruct, Mistral-Instruct-v0.3, and	×	0.08
ReST-KV is compared with five baseline methods: StreamingLLM, H2O, TOVA, SnapKV, and LaCache.	×	0.05
ReST-KV is evaluated on four benchmarks: LongBench, RULER, Needle-in-a-Haystack, and InniteBench.	×	0.10
ReST-KV achieves the best performance in most cases on the LongBench benchmark.	×	0.05
ReST-KV reduces peak memory usage by approximately 36.0% compared to full cache at a context length of 128k.	×	0.08
ReST-KV achieves an approximate 10.61 \times speedup over the full cache method at a 128K context length.	×	0.10
ReST-KV is compatible with prell sparse attention approaches, yielding a Time-To-First-Token (TTFT) speedup of up to 3.	×	0.04
ReST-KV uses a xed cache budget to limit the number of KV pairs, maintaining high efficiency.	×	0.07
ReST-KV only requires computing attention outputs within a small query window, resulting in a computational complexity c	×	0.05
LLMs typically decode text in an auto-regressive manner, which allows them to generate high-quality, contextually cohere	×	0.03
KV cache reduces redundant computation by storing previously computed keys and values.	×	0.06

References

- <http://arxiv.org/abs/2605.09649v1>
- <http://arxiv.org/abs/2605.08840v1>
- <http://arxiv.org/abs/2605.07234v1>