

Parallel Context Windows vs Sliding Window Accuracy on Long-Context Needle-in-a-Haystack Benchmarks

Assignee Research

June 4, 2026

Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: How does the Parallel Context Windows method impact accuracy on the Needle In A Haystack benchmark compared to sliding window approaches for context lengths exceeding 100k 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 6.0/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Adaptive Chunking: Optimizing Chunking-Method Selection for RAG. Research question: How does the Parallel Context Windows method impact accuracy on the Needle In A Haystack benchmark compared to sliding window approaches for context lengths exceeding 100k tokens?.

2 Methodology

Systematic literature search across multiple databases yielded 15 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 6.0/10.

3 Results

15 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 6.0/10.

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/2603.25333v1>
- <http://arxiv.org/abs/2212.10947v3>
- <http://arxiv.org/abs/2602.14200v5>