

# Bidirectional Token Prediction Enhances Robustness in Long-Context Key-Value Retrieval

Assignee Research

June 9, 2026

## Abstract

This report synthesises findings from 7 peer-reviewed papers addressing the following research question: What is the impact of bidirectional token prediction strategies on the robustness of key-value retrieval in long-context models when evaluated on the RULER benchmark with corrupted context windows. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 2.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

## 1 Introduction

This paper examines: RULER: What's the Real Context Size of Your Long-Context Language Models?. Research question: What is the impact of bidirectional token prediction strategies on the robustness of key-value retrieval in long-context models when evaluated on the RULER benchmark with corrupted context windows?.

## 2 Methodology

Systematic literature search across multiple databases yielded 7 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 2.7/10.

## 3 Results

7 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 2.7/10.

## 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/2402.07440v3>
- <http://arxiv.org/abs/2404.06654v3>
- <http://arxiv.org/abs/2505.09561v2>