

# Joint Past-Token and Next-Token Prediction for Long-Context Retrieval Accuracy

Assignee Research

June 9, 2026

## Abstract

This report synthesises findings from 11 peer-reviewed papers addressing the following research question: How does integrating past-token prediction objectives alongside next-token prediction affect long-context retrieval accuracy on the Needle-in-a-Haystack benchmark under varying noise levels. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.0/10. This report is a machine-generated literature synthesis and does not constitute original research.

## 1 Introduction

This paper examines: Learning Long-Context Diffusion Policies via Past-Token Prediction. Research question: How does integrating past-token prediction objectives alongside next-token prediction affect long-context retrieval accuracy on the Needle-in-a-Haystack benchmark under varying noise levels?.

## 2 Methodology

Systematic literature search across multiple databases yielded 11 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.0/10.

## 3 Results

11 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 3.0/10.

## 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/2406.11230v2>
- <http://arxiv.org/abs/2408.10151v1>
- <http://arxiv.org/abs/2505.09561v2>