

# Scaling Retrieved Documents in Ensemble LLM Reasoning Under Adaptive Adversarial Attacks

Assignee Research

June 9, 2026

## Abstract

This report synthesises findings from 12 peer-reviewed papers addressing the following research question: Does scaling the number of retrieved documents in ensemble configurations improve or degrade LLM reasoning accuracy under adaptive adversarial attacks in AmbiEval benchmarks. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.2/10. This report is a machine-generated literature synthesis and does not constitute original research.

## 1 Introduction

This paper examines: AdversarialCoT: Single-Document Retrieval Poisoning for LLM Reasoning. Research question: Does scaling the number of retrieved documents in ensemble configurations improve or degrade LLM reasoning accuracy under adaptive adversarial attacks in AmbiEval benchmarks?.

## 2 Methodology

Systematic literature search across multiple databases yielded 12 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.2/10.

## 3 Results

12 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 3.2/10.

## 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/2604.12201v1>
- <http://arxiv.org/abs/2307.02055v1>
- <http://arxiv.org/abs/2006.16545v1>