

Scaling Retrievers in RAG Systems Enhances Resistance to Semantic Perturbation Attacks

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 13 peer-reviewed papers addressing the following research question: Does scaling the number of retrievers in a RAG system improve resistance to semantic perturbation attacks compared to single-retriever baselines on literature-focused subsets of AmbiEval. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 6.8/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: To Memorize or to Retrieve: Scaling Laws for RAG-Considerate Pretraining. Research question: Does scaling the number of retrievers in a RAG system improve resistance to semantic perturbation attacks compared to single-retriever baselines on literature-focused subsets of AmbiEval?.

2 Methodology

Systematic literature search across multiple databases yielded 13 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 6.8/10.

3 Results

13 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 6.8/10.

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.

References

- <http://arxiv.org/abs/2411.18583v1>
- <http://arxiv.org/abs/2210.02627v1>
- <http://arxiv.org/abs/2604.00715v1>