

Robustness of Retrieval-Augmented 7B Models Across Medical Subdomains

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 12 peer-reviewed papers addressing the following research question: How does the robustness of retrieval-augmented 7B models vary across different medical subdomains (e.g., pharmacology vs. radiology) compared to non-retrieval-augmented baselines, as evaluated by. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 6.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: RadioRAG: Online Retrieval-augmented Generation for Radiology Question Answering. Research question: How does the robustness of retrieval-augmented 7B models vary across different medical subdomains (e.g., pharmacology vs. radiology) compared to non-retrieval-augmented baselines, as evaluated by domain-specific accuracy and hallucination detection scores?.

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 6.7/10.

3 Results

12 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 6.7/10.

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/2407.15621v3>
- <http://arxiv.org/abs/2411.18583v1>
- <http://arxiv.org/abs/2412.01496v2>