

The Integration Of Domain-Adaptive Fine-Tuning Before Retrieval Augmentation Performance On The Precision-Recall

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 4 peer-reviewed papers addressing the following research question: How does the integration of domain-adaptive fine-tuning before retrieval augmentation affect the precision-recall trade-off in 7B vs. 70B models on the Qur'anQA benchmark. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Domain Fine-Tuning vs. Retrieval-Augmented Generation for Medical Multiple-Choice Question Answering: A Controlled Comparison at the 4B-Parameter Scale. Research question: How does the integration of domain-adaptive fine-tuning before retrieval augmentation affect the precision-recall trade-off in 7B vs. 70B models on the Qur'anQA benchmark?.

2 Methodology

Systematic literature search across multiple databases yielded 4 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.7/10.

3 Results

4 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 3.7/10.

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/2510.22531v1>
- <http://arxiv.org/abs/2604.23801v1>
- <http://arxiv.org/abs/2110.06500v2>