Continuous Contrastive Learning Enhances Distributional Robustness in Dense Retrieval Models

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 12 peer-reviewed papers addressing the following research question: To what extent does continuous contrastive learning on target corpora improve the distributional robustness of dense retrieval models compared to standard fine-tuning on out-of-domain benchmark. 5 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.8/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Pre-training vs. Fine-tuning: A Reproducibility Study on Dense Retrieval Knowledge Acquisition. Research question: To what extent does continuous contrastive learning on target corpora improve the distributional robustness of dense retrieval models compared to standard fine-tuning on out-of-domain benchmark datasets?.

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 3.8/10.

3 Results

12 papers retrieved. 5 claims extracted; 0 independently verified. Quality review score: 3.8/10.

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.

5 Extracted Claims

Claim	Verified	Confidence
The study investigates the consistency of embedding knowledge between dense retrievers and their untrained backbone coun	×	0.09
Fine-tuned dense retrieval models encode knowledge in a manner consistent with their pretraining initialization.	×	0.15
Linear probing involves training a simple linear classifier on embeddings extracted from different layers of a model to	×	0.03
The task considered by Reichman and Heck is that of distinguishing relevant and irrelevant documents with respect to a q	×	0.02
If a linear classifier achieves high accuracy, it suggests that the embeddings contain structure.	×	0.03

References

- http://arxiv.org/abs/2205.02303v1
- http://arxiv.org/abs/2205.02870v2
- http://arxiv.org/abs/2505.07166v1