Continuous Contrastive Learning for Zero-Shot Retrieval in Out-of-Domain Benchmarks

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: To what extent does continuous contrastive learning on domain-specific corpora improve zero-shot retrieval accuracy on out-of-domain benchmarks compared to standard fine-tuning of LLaMA-based dense. 16 claims were extracted from source literature; 3 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: COCO-DR: Combating Distribution Shifts in Zero-Shot Dense Retrieval with Contrastive and Distributionally Robust Learning. Research question: To what extent does continuous contrastive learning on domain-specific corpora improve zero-shot retrieval accuracy on out-of-domain benchmarks compared to standard fine-tuning of LLaMA-based dense retrievers?.

2 Methodology

Systematic literature search across multiple databases yielded 14 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.5/10.

3 Results

14 papers retrieved. 16 claims extracted; 3 independently verified. Quality review score: 4.5/10.

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
Dense retrieval methods show stronger advantages over sparse retrieval methods in supervised settings compared to zero-s	×	0.10
Dense retrieval models perform worse than BM25 on BEIR tasks that are less similar to the MS MARCO dataset.	×	0.06
ANCE initialized from coCondenser underperforms BM25 on BEIR tasks where distribution shifts are severe.	×	0.07
COCO-DR introduces two training techniques: Continuous Contrastive pretraining (COCO) and implicit Distributionally Robu	✓	0.19
The COCO component continuously pretrains the language model on target corpora to handle document distribution shifts.	✓	0.19
The iDRO component improves model robustness during fine-tuning to lead to better generalization for unseen target queri	✓	0.15
Sequence Contrastive Learning (SCL) aims to improve the alignment of similar text sequences and the uniformity of unrela	×	0.04
SCL-pretrained models still suffer from distribution shifts in zero-shot settings.	×	0.15
COCO pretraining involves randomly extracting two disjoint sequences from a document to form a positive pair for contras	×	0.05
COCO-DR Large achieves an nDCG@10 score of 0.789 on the TREC-COVID dataset.	×	0.10
COCO-DR Base achieves an nDCG@10 score of 0.789 on the TREC-COVID dataset.	×	0.09
Removing the iDRO component from COCO-DR Base reduces the nDCG@10 score on TREC-COVID from 0.789 to 0.771.	×	0.08
Removing the COCO component from COCO-DR Base reduces the nDCG@10 score on TREC-COVID from 0.789 to 0.763.	×	0.08
COCO-DR Base achieves an nDCG@10 score of 0.867 on the Quora dataset.	×	0.09
COCO-DR Base achieves an nDCG@10 score of 0.355 on the NFCorpus dataset.	×	0.09
COCO-DR Base achieves an nDCG@10 score of 0.505 on the NQ dataset.	×	0.09

References

- http://arxiv.org/abs/2402.15059v1
- http://arxiv.org/abs/2210.15212v2
- http://arxiv.org/abs/2505.07166v1