

Misspelling Robustness Training and Zero-Shot Cross-Domain Performance in Dense Retrieval Models

Assignee Research

June 12, 2026

Abstract

Prevailing research practice today often relies on training dense retrievers on existing large datasets such as MSMARCO and then experimenting with ways to improve zero-shot generalization capabilities to unseen domains. While prior work has tackled this challenge through resource-intensive steps such as data augmentation, architectural modifications, increasing model size, or even further base model pretraining, comparatively little investigation has examined whether the training procedures themselves can be improved to yield better generalization capabilities in the resulting models. In this

1 Introduction

This paper examines: Back to Basics: A Simple Recipe for Improving Out-of-Domain Retrieval in Dense Encoders. Research question: What is the impact of misspelling robustness training on the zero-shot cross-domain performance of dense retrieval models evaluated on diverse benchmark suites?.

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 7.5/10.

3 Results

12 papers retrieved. 8 claims extracted; 6 independently verified. Quality review score: 7.5/10.

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
LoRA leads to better out-of-domain generalization performance compared to full parameter tuning.	✓	0.23
Full parameter tuning outperforms LoRA on in-domain settings.	✓	0.22
Mined hard negatives may hurt out-of-domain retrieval performance unless selected with great care.	✓	0.24
Increasing the number of in-batch negatives is consistently beneficial for out-of-domain performance.	✓	0.22
Under identical GPU configurations, increasing the in-batch size typically yields more robust performance compared to ad	✓	0.28
Common full parameter fine-tuning practices are prone to overfitting large popular datasets like MSMARCO.	✓	0.25
Performance tends to increase when using RoBERTa-large compared to BERT-large.	×	0.05
The benefits of LoRA and the possible detriment of hard negatives hold true for both RoBERTa-large and BERT-large.	×	0.05

References

- <http://arxiv.org/abs/2204.00716v2>
- <http://arxiv.org/abs/2311.09765v1>
- <http://arxiv.org/abs/2205.02303v1>