

# SOVEREIGN: Can COCO-DR’s alignment mechanism achieve higher zero-shot accuracy on the NQ and TriviaQA open-domain QA data

SOVEREIGN Research Kernel

Autonomous draft — Owner review required before publication

May 29, 2026

## Abstract

We present a new zero-shot dense retrieval (ZeroDR) method, COCO-DR, to improve the generalization ability of dense retrieval by combating the distribution shifts between source training tasks and target scenarios. To mitigate the impact of document differences, COCO-DR continues pretraining the language model on the target corpora to adapt the model to target distributions via COntinuous COntrastive learning. To prepare for unseen target queries, COCO-DR leverages implicit Distributionally Robust Optimization (iDRO) to reweight samples from different source query clusters for improving model ro

## 1 Introduction

Analysis of: COCO-DR: Combating Distribution Shifts in Zero-Shot Dense Retrieval with Contrastive and Distributionally Robust Learning. Research goal: Can COCO-DR’s alignment mechanism achieve higher zero-shot accuracy on the NQ and TriviaQA open-domain QA datasets than dense retrievers fine-tuned with supervised data, measured by recall@5?.

## 2 Methodology

Multi-query arXiv search (4 parallel queries, Relevance-sorted). TF-IDF cosine semantic verification (bigrams, threshold=0.15). NIM nv-embedqa-e5-v5 (dim=1024) for semantic indexing. Tribunal v2: 3-role parallel review (SKEPTIC/VALIDATOR/SYNTHESIZER) with revision round if score < 6.5.

### 3 Results

3 papers retrieved. 0 claims extracted, 0 verified. Tribunal: 4.5/10 → RE-  
VISE (revision\_round=1). Policy: ESCALATE\_TO\_OWNER.

### 4 Uncertainties

NIM free tier latency varies. TF-IDF verification is a weak signal. arXiv  
Relevance ranking is query-dependent. Tribunal consensus is LLM-based  
and prompt-sensitive.

### References

- <http://arxiv.org/abs/2210.15212v2>
- <http://arxiv.org/abs/2602.05656v2>
- <http://arxiv.org/abs/2207.13332v2>