

MA-DPR Performance on Cross-Domain OOD Queries: A Manifold-Aware Retrieval Benchmark

Assignee Research

June 1, 2026

Abstract

This report synthesises findings from 13 peer-reviewed papers addressing the following research question: How does the performance of MA-DPR on cross-domain OOD queries (e.g., TriviaQA to HotpotQA) compare to other manifold-aware retrieval methods like SimCLR or contrastive learning-based approaches in. Many questions cannot be answered simply; their answers must include numerous nuanced details and additional context. Complex Answer Retrieval (CAR) is the retrieval of answers to such questions. 14 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: Overcoming low-utility facets for complex answer retrieval. Research question: How does the performance of MA-DPR on cross-domain OOD queries (e.g., TriviaQA to HotpotQA) compare to other manifold-aware retrieval methods like SimCLR or contrastive learning-based approaches in terms of exact match accuracy and recall@K?

2 Methodology

Systematic literature search across multiple databases yielded 13 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

13 papers retrieved. 14 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
Manual relevance judgments are graded on a scale from Must be mentioned (3) to Trash (-2).	×	0.01
Manual relevance judgments cover a subset of queries (702 of the 2,125 queries).	×	0.05
Two versions of the knowledge graph are generated: one using hyperlinks as entity mentions and one using entity mentions	×	0.06
Edge labels are limited to the top 1000 most frequently-used labels.	×	0.03
HolE embeddings are trained using the link graph for 5,000 iterations.	×	0.04
Top 2 entity scores are included in the model.	×	0.05
Each model is trained for 80 epochs with samples from train.fold1-2.	×	0.03
Automatic relevance judgments serve as a source of relevant documents.	×	0.02
Top non-relevant BM25 documents are used as negative training examples.	×	0.03
For each positive sample, 6 negative samples are included.	×	0.11
The training iteration that yields the highest R-Precision value on the validation dataset (test200) is selected for evaluation	×	0.03
Top 100 BM25 results for each query in benchmarkY1test are reranked.	×	0.02
Neural models have been shown to be competitive with conventional ranking techniques due to their ability to learn match	×	0.06
Preliminary work has shown neural ranking architectures to be more effective than conventional approaches out-of-the-box	×	0.05

References

- <http://arxiv.org/abs/1811.08772v1>
- <http://arxiv.org/abs/2404.14464v1>
- <http://arxiv.org/abs/2407.20114v3>