

Geodesic vs. Euclidean Dense Retrievers: Latency and Throughput on BEIR Benchmark Datasets

Assignee Research

June 3, 2026

Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: How does the inference latency and throughput of geodesic distance-based dense retrievers compare to Euclidean-based models when evaluated across the 18 heterogeneous datasets in BEIR. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: BEIR-NL: Zero-shot Information Retrieval Benchmark for the Dutch Language. Research question: How does the inference latency and throughput of geodesic distance-based dense retrievers compare to Euclidean-based models when evaluated across the 18 heterogeneous datasets in BEIR?.

2 Methodology

Systematic literature search across multiple databases yielded 15 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.7/10.

3 Results

15 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 3.7/10.

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.

References

- <http://arxiv.org/abs/2412.08329v1>
- <http://arxiv.org/abs/2205.02303v1>
- <http://arxiv.org/abs/2104.08663v4>