SOVEREIGN: What is the trade-off between inference efficiency and robustness to adversarial query perturbations for spars

SOVEREIGN Research Kernel
Autonomous draft — Owner review required before publication

May 28, 2026

Abstract

Pyserini is a Python toolkit for reproducible information retrieval research with sparse and dense representations. It aims to provide effective, reproducible, and easy-to-use first-stage retrieval in a multi-stage ranking architecture. Our toolkit is self-contained as a standard Python package and comes with queries, relevance judgments, pre-built indexes, and evaluation scripts for many commonly used IR test collections. We aim to support, out of the box, the entire research lifecycle of efforts aimed at improving ranking with modern neural approaches. In particular, Pyserini supports sparse

1 Introduction

Analysis of: Pyserini: A Python Toolkit for Reproducible Information Retrieval Research with Sparse and Dense Representations. Research goal: What is the trade-off between inference efficiency and robustness to adversarial query perturbations for sparse vs. dense retrievers on the TREC Deep Learning track?.

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

12 papers retrieved. 6 claims extracted, 6 verified. Tribunal: 7.7/10 → APPROVE (revision_round=0). Policy: AUTO_APPROVE.

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.

5 Extracted Claims

Claim	Verified	Confidence
Pyserini is a Python toolkit for reproducible information retrieval research with sparse and dense representations	✓	0.46
Pyserini supports sparse retrieval such as BM25 scoring using bag-of-words representations	✓	0.29
Pyserini supports dense retrieval such as nearest-neighbor search on transformer-encoded representations	✓	0.27
Pyserini is self-contained as a standard Python package	✓	0.21
Pyserini comes with queries, relevance judgments, pre-built indexes, and evaluation scripts for commonly used IR test co	✓	0.33
Pyserini supports sparse, dense, and hybrid retrieval approaches	✓	0.21

References

- https://doi.org/10.1145/3768153
- https://doi.org/10.1145/3404835.3463238
- https://doi.org/10.48550/arxiv.2312.10997