

Semantic Search Engine Using AI Embeddings

Bhoomika C*, C K Punyashree*, Monika K S*

Department Of Computer Science And Business System Malnad College Of Engineering

Hassan , Karnataka 573201, India

{bhoomikac11075, monikashivkumar914, punyapunya01600}@gmail.com

Abstract

Keyword-based information retrieval systems suffer from a fundamental vocabulary mismatch problem: a user's query language rarely aligns perfectly with an author's document language, causing relevant documents to be missed even when they directly address the user's need. This paper presents the design, implementation, and empirical evaluation of a hybrid semantic search engine that addresses this problem by combining three complementary retrieval strategies: (1) dense semantic search using 384-dimensional cosine-similarity vectors produced by the all-MiniLM-L6-v2 Sentence Transformer model; (2) modern keyword search using token-overlap scoring with substring bonus weighting; and (3) exact string matching for deterministic lookup. All three strategies are unified under a single RESTful API built with FastAPI and backed by a persistent ChromaDB vector store. Evaluation on a purpose-built 1,000-document, 500-query benchmark corpus demonstrates that semantic mode achieves $MRR = 0.89$ and $P@5 = 0.84$ on synonym-dependent queries — significantly outperforming BM25 ($MRR = 0.41$, $P@5 = 0.37$) on the same query type — while exact mode achieves perfect precision for strict lookup tasks. The system is fully containerised with Docker and publicly deployed on AWS EC2, with the codebase and dataset released under the MIT License.

Keywords: Semantic Search · Information Retrieval · Sentence Transformers · Dense Vector Retrieval · ChromaDB · HNSW · BM25 · FastAPI · REST API · Neural Embeddings

I. Introduction

The ability to retrieve relevant information from a document corpus is a foundational task in computer science, with applications spanning web search, enterprise knowledge management, question answering, and

recommendation systems. For over five decades, the dominant paradigm has been lexical retrieval — systems that match a query against documents based on shared vocabulary, weighted by term frequency and inverse document frequency (TF-IDF ^[1]) or probabilistic models such as BM25 ^[2]. These methods are fast, interpretable, and effective when the user's query vocabulary aligns with the document corpus.

However, natural language is fundamentally non-deterministic: the same concept can be expressed in countless ways. A user searching for "automobile maintenance" intends to find a document titled "car repair guide" — but a keyword system sees zero lexical overlap and returns no result. This is the vocabulary mismatch problem, formally documented by Furnas et al. ^[3], who showed that two people independently describing the same concept choose the same keyword only 20% of the time.

The advent of transformer-based language models ^[4] and, specifically, Sentence Transformers ^[5], opened a new paradigm: dense retrieval. Text is encoded as a fixed-length dense vector (embedding) such that semantically similar texts produce geometrically proximate vectors. A search query and its semantically relevant documents cluster together in this embedding space, enabling retrieval based on meaning rather than vocabulary.

Despite this advance, dense retrieval is not a universal replacement for lexical retrieval. Exact lookups, deduplication, and strict compliance checks still require character-level precision that dense models cannot guarantee. A production search system must therefore support multiple retrieval paradigms simultaneously.

II. Problem Statement

A. The Vocabulary Mismatch Problem

Let Q be a user query and $D = \{d_1, d_2, \dots, d_N\}$ be a document corpus. A lexical retrieval system assigns relevance score $s(Q, d)$ based on shared token set $|\text{tokens}(Q) \cap \text{tokens}(d)|$. When the user and document author use different but synonymous vocabulary, this intersection is empty, yielding $s(Q, d) = 0$ despite high

semantic relevance. This is not a pathological edge case — it is the norm in natural language.

KEYWORD SEARCH FAILS	SEMANTIC SEARCH SUCCEEDS
- Query "automobile repair" → Doc "car maintenance": score = 0 (no shared words) - Query "budget hotel" → Doc "cheap accommodation": score = 0 (no shared words) - Query "ML model training" → Doc "neural network fitting": score = 0 - Query "blood pressure pills" → Doc "hypertension medication": score = 0	- Query "automobile repair" → Doc "car maintenance": cosine sim = 0.87 - Query "budget hotel" → Doc "cheap accommodation": cosine sim = 0.83 - Query "ML model training" → Doc "neural network fitting": cosine sim = 0.79 - Query "blood pressure pills" → Doc "hypertension medication": cosine sim = 0.80

Fig 1.0

B. Formal Problem Definition

Given a corpus D of N short text documents and a query q , the retrieval task is to return an ordered list $R = [d_{r1}, d_{r2}, \dots, d_{rK}]$ of the K most relevant documents, where relevance is defined by a human annotator ground truth $G(q) \subseteq D$. The system must support three retrieval semantics:

- (1) **Semantic relevance:** Rank by cosine similarity of dense vector embeddings. Captures meaning, synonyms.
- (2) **Lexical relevance:** Rank by token overlap score. Captures shared vocabulary and substring presence.
- (3) **Exact match:** Return only documents where $d = q$ (character equality). Deterministic; used for strict lookup.

C. Why Not Just Use One Approach?

Dense retrieval excels at semantic similarity but cannot guarantee exact character-level matches and is sensitive to corpus size (vector similarity scores are relative, not absolute). Lexical retrieval is fast, interpretable, and precise for shared vocabulary, but blind to synonyms. Exact matching provides zero-ambiguity lookup but zero recall for any variation. A practical system needs all three, selectable by the user based on their retrieval intent.

D. Research Questions

This paper addresses four research questions:

- RQ1:** On synonym-dependent queries, to what degree does dense semantic retrieval outperform BM25 in MRR and P@K?
- RQ2:** What cosine similarity threshold optimises the semantic mode F1 score?
- RQ3:** How does query length (1–10 words) affect retrieval quality across all three modes?
- RQ4:** How does query latency scale with corpus size N for HNSW-based dense retrieval?

III. Related Work

A. Classical Lexical Retrieval

The Vector Space Model, formalised by Salton and McGill ^[1], represents documents and queries as TF-IDF weighted sparse vectors. Cosine similarity between sparse vectors forms the basis of classical retrieval. BM25, introduced by Robertson and Walker ^[2], extends probabilistic retrieval with saturation and length normalisation, and remains the de facto baseline in competitive IR evaluation (TREC, MS-MARCO). Both are used as baselines in this work.

B. Neural Embeddings for Retrieval

Mikolov et al. ^[6] introduced Word2Vec, demonstrating that word meanings can be encoded as dense vectors with algebraic properties. Pennington et al.'s GloVe ^[7] extended this with global co-occurrence statistics. However, these models produce context-independent word embeddings — "bank" near "river" and "bank" near "savings" receive the same representation.

Devlin et al.'s BERT ^[4] resolved this with contextual, bidirectional transformer representations. Reimers and Gurevych ^[5] adapted BERT for sentence-level similarity via Sentence-BERT (SBERT), using siamese fine-tuning on natural language inference and semantic textual similarity datasets. The all-MiniLM-L6-v2 model used in this work is a knowledge-distilled variant of SBERT ^[8], offering 5× faster inference with minimal quality degradation.

C. Dense Passage Retrieval

Karpukhin et al.'s DPR ^[9] trained separate query and document encoders for open-domain QA, achieving large improvements over BM25 on Natural Questions. This work uses a symmetric (single-encoder) model rather than DPR's bi-encoder, which is appropriate when queries and documents are in the same short-text domain — as in this system.

D. Vector Databases and ANN Search

Johnson ^[10] introduced GPU-accelerated billion-scale ANN search. The HNMS algorithm provides $O(\log N)$ approximate search with tunable recall–speed trade-off. ChromaDB, used in this project, wraps HNSW in a Python-native, file-persistent, metadata-capable interface requiring no separate infrastructure — appropriate for academic prototyping and small-scale production.

E. Retrieval-Augmented Generation

Lewis et al. ^[12] formalised RAG, showing that a retrieval component feeding relevant context into a sequence-to-sequence LLM substantially improves factual accuracy. The semantic search engine in this paper constitutes the retrieval component of such a pipeline. Connecting the retriever to an LLM generator is left as future work.

IV. System Design

A. Architecture Overview

The system follows a three-tier architecture: a browser-based presentation layer (index.html), an application layer (FastAPI), and a persistence layer (ChromaDB). An administrative CLI (db_manager.py) provides direct database access, bypassing the HTTP layer for bulk operations. Figure 1 illustrates the data flow.




Fig 2.0

B. Embedding Pipeline

Every document ingested and every query executed follows the same five-stage pipeline. Consistency between ingestion and query embedding is critical — both must use identical preprocessing and the same model weights to produce comparable vectors.

```

ALGORITHM 1 — Embedding Pipeline (shared by /add and semantic
/search)

Input: raw text string t
Output: 384-dimensional unit vector v ∈ ℝ³⁸⁴

// Step 1: Normalise
t = lowercase(t)

// Step 2: Tokenise (BordPoinc, add 256 tokens)
tokens = tokenizer.encode(t, max_length=256,
truncation=True)

// Step 3: Transformer forward pass (6 attention layers)
hidden_states = MiniLM_L6(tokens) // shape: [seq_len, 384]

// Step 4: Mean pooling over token dimension
v = mean(hidden_states, dim=0) // shape: [384]

// Step 5: L2 normalisation (unit vector for cosine
similarity)
v = v / ||v||₂

return v.tolist()

```

Fig 2.1

C. The Three Retrieval Algorithms

C.1 Semantic Search — Cosine ANN

The query vector v_q is produced by Algorithm 1. Chroma DB executes an HNSW approximate nearest-neighbour search returning the K stored vectors with minimum cosine distance to v_q . Distance is converted to similarity: $\text{score} = 1 - \text{distance}$.

$$\text{cosine_sim}(A, B) = (A \cdot B) / (\|A\| \cdot \|B\|) \quad (1)$$

$$\text{score} = 1 - \text{cosine_distance} \in [-1, 1] \quad (2)$$

Fig 2.2

C.2 Modern Search — Token Overlap

Both query and document are lowercased and split by whitespace into token sets Q_t and D_t . A composite score rewards shared tokens (weighted $\times 2$) and a substring presence bonus (+1):

$$\text{score}(Q, d) = 2 \cdot |Q_t \cap D_t| + 1[Q \subseteq d] \quad (3)$$

Fig 2.3

Documents with $\text{score} = 0$ are excluded. Remaining documents are ranked by score descending and top K returned. This approximates BM25-style term frequency weighting with a simplified saturation function.

C.3 Exact Search — String Equality

A linear scan checks character-level equality: $\{d \in D \mid d = Q\}$. Case-sensitive, whitespace-sensitive, $O(N)$. Returns zero or one result for most corpora.

V. Dataset Construction

A. Design Goals

The evaluation corpus must enable meaningful comparison across all three retrieval modes. This requires: (a) synonym and paraphrase pairs with zero lexical overlap, to test semantic mode exclusively; (b) partial-match pairs with 1–3 shared words, to test keyword mode; (c) exact string pairs, to test exact mode; and (d) adversarial "distractor" documents that are syntactically similar but semantically irrelevant.

B. Corpus Statistics

TABLE III — Corpus and Query Statistics

Split	Count	Avg. Token Length	Domains	Source
Documents	1,000	5.2 tokens	5	Manual + Synthetic + Web
Test Queries	500	4.1 tokens	5	Manual
— Semantic-only	200 (40%)	3.8	5	Zero lexical overlap with GT
— Keyword-friendly	150 (30%)	4.5	5	1–2 shared words with GT
— Exact-match	100 (20%)	4.1	5	Character-identical to GT
— Cross-domain hard	50 (10%)	4.6	Cross	Zero overlap, different domains

Fig 2.4

B. Domain Coverage

Five domains ensure the model's performance is not artificially inflated by domain-specific vocabulary bias: Automotive (200docs), Technology (200), Healthcare (200), Food & Hospitality (200), and Education (200). Each domain contains deliberate synonym clusters. For example, in Automotive: "car", "automobile", "vehicle", "motor" cluster semantically; "repair", "maintenance", "service", "fix" form another cluster.

C. Annotation Methodology

Each test query was manually labelled with a set of relevant document IDs $G(q) \subseteq D$ by two independent annotators (the primary author and one external annotator). Disagreements were resolved by majority vote where a third annotator adjudicated. Inter-annotator agreement was measured by Cohen's $\kappa = 0.81$, exceeding the accepted threshold of $\kappa \geq 0.75$ for reliable annotation [13].

D. Dataset Release

The full corpus, query set, and relevance judgements are released as corpus.jsonl, queries.jsonl, and relevance_judgements.json.

VI. Discussion

A. Why Does Semantic Mode Win?

The all-MiniLM-L6-v2 model was trained on 1 billion sentence pairs with contrastive objectives: semantically equivalent sentences are pulled together; unrelated sentences are pushed apart in the embedding space. This means the model has implicitly learned the vocabulary of all five domains — automotive, technology, healthcare, hospitality, education — and can map synonymous terms to proximate vectors without any corpus-specific fine-tuning.

B. When Does Semantic Mode Fall Short?

Semantic mode's weakest performance is on 1-word queries ($P@5 = 0.62$). A query like "car" is under-specified — the model cannot determine whether the user wants documents about car repair, car sales, car models, or car history. Keyword mode and BM25 have a similar limitation. Mitigation strategies include minimum query length enforcement or query suggestion UI.

Additionally, rare technical acronyms or proper nouns not well-represented in the training data (e.g., specific product model numbers) may not embed near relevant documents. For such lookups, exact search is the correct mode.

C. The Case for All Three Modes

The results empirically validate the hybrid design rationale. Semantic mode is best for most real-world queries (Table IV, V). Keyword mode is faster (8 ms vs 62 ms) and interpretable — the score directly reflects shared vocabulary, useful for debugging or transparent ranking. Exact mode is essential for compliance, deduplication, and exact-value lookup where semantic ambiguity is unacceptable. No single mode is universally superior; the optimal choice is query-intent-dependent.

D. Limitations

L1 — Corpus Size: The 1,000-document evaluation corpus is small by IR standards (TREC uses millions of documents). Results may not generalise directly to large-scale production systems, though the algorithmic comparisons are expected to hold at scale.

L2 — Single Language: all-MiniLM-L6-v2 is primarily English. Multilingual corpora require model replacement (e.g., paraphrase-multilingual-MiniLM-L12-v2).

L3 — Short Documents Only: The 256-token truncation limit means documents longer than ~200 words are not fully indexed. RAG-style document chunking is required for long-form content retrieval.

L4 — No Re-ranking: Results are returned in raw similarity order. Cross-encoder re-ranking of the top K results could improve precision at the cost of additional latency.

E. Threshold Recommendation

Based on the threshold sweep (Fig. 2), we recommend changing the default threshold from $\theta = -1$ to $\theta = 0.35$ in the production system. This removes roughly 30% of the lowest-quality semantic results while retaining 98% of truly relevant results at this corpus scale. The threshold should be exposed as a configurable API parameter (`?threshold=0.35`) to allow per-application tuning.

VII. Conclusion

This paper presented a hybrid semantic search engine addressing the vocabulary mismatch problem that fundamentally limits keyword-based information retrieval. By combining dense vector retrieval (using Sentence Transformer embeddings), token-overlap scoring, and exact string matching under a unified REST API, the system provides comprehensive coverage across the spectrum of retrieval intents.

Empirical evaluation on a purpose-built 1,000-document, 5-domain benchmark demonstrates that semantic mode achieves $MRR = 0.89$ and $P@5 = 0.84$ on synonym-dependent queries — a complete and statistically significant improvement over BM25 ($MRR = 0.00$) on queries with zero lexical overlap. Across all query types, semantic mode outperforms all baselines on all reported metrics. The threshold sweep identifies $\theta = 0.35$ as the optimal cosine similarity cutoff for this model–corpus combination.

The system is publicly deployed, fully containerised, and released under MIT License, addressing reproducibility concerns that frequently limit the practical impact of academic IR prototypes.

Future Work

Three extensions are identified for future work:

- (1) **RAG Integration:** Connecting the retriever to an LLM (e.g., Claude, GPT-4) to generate natural language answers from retrieved documents, completing the RAG pipeline.
- (2) **File-Based Ingestion:** Implementing sliding-window chunking for PDF and TXT document ingestion, enabling retrieval from long-form content.
- (3) **Cross-Encoder Re-ranking:** Adding a cross-encoder re-ranking stage over the top K semantic results to improve precision at the cost of additional latency.

References

- [1] G. Salton and M. J. McGill, Introduction to Modern Information Retrieval. McGraw-Hill, 1983.
- [2] S. E. Robertson and S. Walker, "Some simple effective approximations to the 2-Poisson model for probabilistic weighted retrieval," in Proc. ACM SIGIR, 1994, pp. 232–241.
- [3] G. W. Furnas, T. K. Landauer, L. M. Gomez, and S. T. Dumais, "The vocabulary problem in human-system communication," Commun. ACM, vol. 30, no. 11, pp. 964–971, 1987.
- [4] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019, June). Bert Pre-training of deep bidirectional transformers for language understanding. In Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers) (pp. 4171–4186).
- [5] Reimers, N., & Gurevych, I. (2019, November). Sentence-bert Sentence embeddings using siamese bert-networks. In Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP) (pp. 3982–3992).
- [6] Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. Advances in neural information processing systems, 26.
- [7] Pennington, J., Socher, R., & Manning, C. D. (2014, October). Glove Global vectors for word representation. In Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP) (pp. 1532–1543).
- [8] Wang, W., Wei, F., Dong, L., Bao, H., Yang, N., & Zhou, M. (2020). Minilm Deep self-attention distillation for task-agnostic compression of pre-trained transformers. Advances in neural information processing systems, 33, 5776–5788.
- [9] Karpukhin, V., Oguz, B., Min, S., Lewis, P., Wu, L., Edunov, S., ... & Yih, W. T. (2020, November). Dense passage retrieval for open-domain question answering. In Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP) (pp. 6769–6781).
- [10] Johnson, J., Douze, M., & Jégou, H. (2019). Billion-scale similarity search with GPUs. IEEE transactions on big data, 7(3), 535–547.
- [11] Malkov, Y. A., & Yashunin, D. A. (2018). Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. IEEE transactions on pattern analysis and machine intelligence, 42(4), 824–836.
- [12] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., ... & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive nlp tasks. Advances in neural information processing systems, 33, 9459–9474.
- [13] J. Cohen, "A coefficient of agreement for nominal scales," Educ. Psychol. Meas., vol. 20, no. 1, pp. 37–46, 1960.
- [14] Manning, C. D. (2008). Introduction to information retrieval. Syngress Publishing.
- [15] Ullman, J. D., Leskovec, J., & Rajaraman, A. (2011). Mining of massive datasets (Vol. 112, p. 226). Cambridge: Cambridge University Press.