

A Comparative Experimental Study of Keyword Search vs. Semantic Search in MERN Stack Applications Using Vector Embeddings

Abhijeet Kumar Ojha¹, Satish Nishad², Shubham³, Karishma Samal⁴, Hritik Tyagi⁵,

Prof. Ashutosh Pradhan⁶

^{1,2,3,4,5} RD Engineering College

⁶ Head of Department, MCA

Abstract—This paper offers a statistically rigorous empirical comparison of keyword-based (TF-IDF style) search and cosine similarity vector embedding search within a MERN (MongoDB, Express.js, React, Node.js) stack environment. Unlike prior studies that rely on synthetic benchmarks or vendor-reported metrics, this research is based on a fully reproducible real-world experiment: a MongoDB Atlas cluster hosting 1,000 product documents with 128-dimensional hash-based embeddings, evaluated through 50 manually constructed test queries covering exact-match, synonym, and conceptual query types. Performance was measured using Precision@5 (P@5), Mean Reciprocal Rank (MRR), and Recall@10. Statistical significance was assessed with paired two-tailed t-tests and Cohen's d effect size. Results show that keyword search achieves significantly higher overall retrieval accuracy (MRR: 0.970 vs 0.432, $t(49) = 7.719$, $p < 0.0001$, $d = 1.592$), while semantic search outperforms keyword search on synonym-type queries (MRR: 0.938 vs 0.875) and is consistently 2.8 times faster (0.74ms vs 2.08ms, $t(49) = 10.872$, $p < 0.0001$, $d = 2.172$). These findings highlight that embedding model quality is the main factor influencing semantic search performance, and even lightweight hash-based embeddings offer measurable benefits for synonym-rich query workloads. The study presents a fully reproducible, open-source benchmark methodology for MERN developers, providing evidence-based recommendations for choosing search architecture in production deployments.

Index Terms—Semantic Search, Keyword Search, Vector Embeddings, MERN Stack, MongoDB Atlas, Paired t-test, Cohen's d , Information Retrieval, Precision@K, MRR, Recall@K, Cosine Similarity, Node.js

I. INTRODUCTION

1.1. Background and Motivation

The MERN stack — MongoDB, Express.js, React, and Node.js — represents one of the most widely adopted full-stack JavaScript frameworks for building scalable web applications [1]. As these applications increasingly serve as data-intensive platforms, search quality has become a critical determinant of user experience. India alone has over 759 million internet users [2], the majority accessing web applications on mobile devices, where search responsiveness and relevance directly govern engagement and retention.

Traditional keyword-based search, implemented via MongoDB's text index and TF-IDF scoring, is the default mechanism in most MERN deployments due to its simplicity and zero external dependencies. However, keyword search suffers a fundamental limitation — vocabulary mismatch: it cannot retrieve documents that describe the same concept using different terms. A user searching for '4k high resolution screen' may not find 'Smart LED TV 43 inch' despite their clear semantic equivalence. This vocabulary gap motivates the study of vector embedding-based semantic search within MERN applications [3].

Vector embedding approaches — representing text as dense numerical vectors in a continuous semantic space — have demonstrated strong retrieval improvements over keyword methods in academic literature [4], [5]. However, most published

benchmarks use Python-centric environments, proprietary datasets, or transformer-scale embedding models requiring GPU infrastructure. This study addresses these gaps by conducting a controlled experiment entirely within a MERN-compatible Node.js environment using MongoDB Atlas as the unified storage and query layer, with lightweight hash-based embeddings that require no external model dependencies.

1.2. Research Questions

1. RQ1: How does keyword (TF-IDF) search compare to hash-based vector embedding search in P@5, MRR, and Recall@10 across 50 diverse queries on a 1,000-document product corpus in a Node.js environment?
2. RQ2: Does query type (exact-match, synonym, conceptual) significantly moderate the relative performance of the two search methods?
3. RQ3: Is there a statistically significant latency difference between keyword and cosine-similarity vector search in Node.js?
4. RQ4: What are the practical implications for MERN developers choosing between keyword and semantic search architectures?

1.3. Contributions

- ▶ A fully reproducible, open-source MERN-stack benchmark comparing keyword and vector search across 50 queries and 1,000 documents, with all statistical analyses reported.
- ▶ First published paired t-test comparison of keyword vs. hash-based embedding search in a Node.js environment, disaggregated by query type.
- ▶ Empirical evidence that hash-based embeddings provide statistically measurable benefit for synonym queries even without transformer model deployment.
- ▶ Evidence-based decision framework for MERN developers selecting search architecture for production deployments in resource-constrained environments.

II. LITERATURE REVIEW

2.1. Keyword Search in Web Applications

TF-IDF (Term Frequency–Inverse Document

Frequency) weighting, formalised by Manning et al. [6], remains the dominant baseline for document retrieval due to its computational efficiency and strong performance on exact-match and navigational queries. MongoDB's built-in text search index implements a TF-IDF variant with stemming and stop-word removal [7]. Yang (2023) reviewed database indexing strategies for semantic search, confirming that TF-IDF maintains competitive precision on direct vocabulary-match queries but degrades significantly on synonym-heavy or natural-language queries where user intent diverges from document terminology [8].

2.2. Vector Embedding and Semantic Search

Reimers and Gurevych [4] introduced Sentence-BERT (SBERT), demonstrating that transformer-based sentence embeddings substantially outperform TF-IDF on semantic similarity benchmarks. Karpukhin et al. [5] showed that Dense Passage Retrieval (DPR) with bi-encoder BERT models outperforms BM25 on open-domain question answering by margins exceeding 20% in top-20 retrieval accuracy. Both approaches require 110M+ parameter transformer models — a non-trivial infrastructure requirement for MERN deployments without GPU access, motivating the investigation of lightweight alternatives [9].

Monir et al. [10] evaluated various embedding strategies and similarity measures for document retrieval, concluding that cosine similarity over dense embeddings consistently outperforms Euclidean distance for text retrieval, and that the embedding model quality is the primary determinant of retrieval performance — a finding this study empirically validates in a MERN-specific context.

2.3. Semantic Search in E-Commerce Contexts

Menon et al. [11] demonstrated that semantic search using `nommic-embed-text-v1` embeddings outperformed BM25 on Amazon product data (mAP: 49.75% vs 41.19%, P@5: 29.52% vs 23.62%). Their study used a managed embedding API; our study replicates a comparable product-search context in a self-contained Node.js implementation, enabling direct comparison of embedding quality impact. Biyyala et al. [12] further demonstrated that AI-driven retrieval approaches yield improvements in

NDCG (+15%) and MRR (+12%) over BM25 at scale, establishing the business case for semantic search adoption.

2.4. MongoDB Atlas and MERN Integration

MongoDB Atlas introduced native \$vectorSearch aggregation pipeline support in 2023, enabling HNSW-indexed approximate nearest neighbour (ANN) search directly within MongoDB collections [13]. Anirudh and Soudagar [14] validated MongoDB Atlas Vector Search for Retrieval-Augmented Generation (RAG) applications, confirming suitability for production semantic workloads. Bhupathi [15] reviewed the broader role of vector databases in GenAI applications, providing architectural context for our experimental design.

2.5. Research Gap

Existing literature lacks a controlled, statistically rigorous comparison of keyword versus vector search within Node.js MERN applications, disaggregated by query type, using identical document corpora and standard IR evaluation metrics with reported effect sizes. This study directly addresses this gap with a fully reproducible methodology and complete statistical reporting, including paired t-tests, Cohen's d , and 95% confidence intervals.

III. METHODOLOGY

3.1. Experimental Design

This study employs a within-subjects repeated measures design. Both search implementations — keyword (TF-IDF) and vector embedding (cosine similarity) — were evaluated on identical document corpora under identical query conditions. All 50 queries were run through both methods in the same Node.js process on the same machine, eliminating confounds from hardware variation, network latency, and corpus differences. This design permits paired statistical analysis, increasing statistical power compared to independent-samples designs.

3.2. Experimental Environment

All experiments were conducted on Windows 11, Intel Core i-series CPU, Node.js v18 LTS, MongoDB Atlas M0 Free Tier (AWS ap-south-1 / Mumbai region). The Mumbai region was deliberately selected to represent real-world latency conditions for Indian

MERN deployments. Both search methods are executed in-memory within the same Node.js process, with MongoDB used as the persistent document and embedding store.

3.3. Dataset Construction

A corpus of 1,000 product documents was constructed across 5 categories (Electronics, Books, Clothing, Sports, Kitchen), comprising 35 unique product archetypes replicated with version suffixes to reach 1,000 total documents. Each document contains a product name, natural language description (15–25 words), category, price, and a 128-dimensional vector embedding. Ground-truth relevance judgements were manually defined by the research team for all 50 test queries before experiment execution, specifying which product names constitute relevant results for each query.

3.4. Embedding Generation

A 128-dimensional hash-based embedding function was implemented in Node.js. For input text T , the function: (1) lowercases and tokenises T into words of length > 2 ; (2) for each word w at position i , computes hash index $h = \sum(\text{charCode} \times 31^j) \bmod 128$ and adds $1/(i+1)$ to $\text{vector}[h]$ (position-weighted); (3) additionally encodes bigrams (adjacent word pairs) at weight 0.5; (4) L2-normalises the resulting vector. This approach captures unigram frequency, positional weighting, and bigram co-occurrence within a fixed 128-dimensional space, without external model inference or Python dependencies.

3.5. Search Implementations

Keyword Search (TF-IDF Baseline)

Keyword search was implemented as a TF-IDF-style in-memory scoring function. For query Q and document D : $\text{score}(Q,D) = \sum_{w \in Q} \{ \text{count}(w,D) / (1 + \log(N)) \}$ where $N=1,000$ is the corpus size. Partial substring matches (query word length > 3) contributed a weight of 0.5. All 1,000 documents were scored per query and sorted by descending score.

Semantic Search (Cosine Similarity)

Semantic search used cosine similarity: $\text{sim}(q,d) = (q \cdot d) / (\|q\| \times \|d\|)$ between the query embedding and each document's stored embedding. The query was encoded using the same 128-dimensional hash function applied during document ingestion. All 1,000 document embeddings were loaded from

MongoDB into memory and scored per query.

3.6. Evaluation Metrics and Statistical Analysis

Metric	Formula	Interpretation
Precision@5 (P@5)	$P@5 = \text{Rel} \cap \text{Top-5} / 5$	Fraction of top-5 relevant results
Mean Reciprocal Rank	$MRR = (1/ Q) \sum 1/\text{rank}_i$	Rank position of the first relevant result
Recall@10 (R@10)	$R@10 = \text{Rel} \cap \text{Top-10} / \text{Rel} $	Coverage of all relevant docs in the top 10
Avg. Latency (ms)	Mean wall-clock time per query	End-to-end response time, Node.js

Statistical significance was assessed using paired two-tailed t-tests (paired because each query was evaluated by both methods under identical conditions). Before t-testing, the normality of the difference scores was examined. Effect size was quantified using Cohen's d, interpreted as: small ($d \geq 0.2$), medium ($d \geq 0.5$), large ($d \geq 0.8$), very large ($d \geq 1.4$). 95% confidence intervals for mean differences were computed using the paired t critical value ($t^* = 2.010$ for $df = 49$). The Bonferroni correction was applied for multiple comparisons across 4 primary metrics (adjusted $\alpha = 0.05/4 = 0.0125$).

3.7. Test Query Design

50 test queries were designed across four categories to isolate search method strengths across a spectrum of retrieval difficulty:

- ▶ Exact-Match Queries (6 queries): Direct term overlap expected with product names — e.g., 'coding interview preparation'. Both methods are expected to perform comparably.
- ▶ Synonym Queries (8 queries): Semantically equivalent but lexically distinct terms — e.g., '4k high resolution screen' for 'Smart LED TV'. Hash embedding is expected to have a measurable advantage.
- ▶ Conceptual Queries (12 queries): Abstract intent queries with minimal vocabulary overlap — e.g., 'habit productivity self improvement'. Tests deep semantic understanding beyond hash encoding capability.
- ▶ General Mixed Queries (24 queries): Diverse queries spanning all product categories and varying specificity levels.

IV. IMPLEMENTATION

4.1. System Architecture

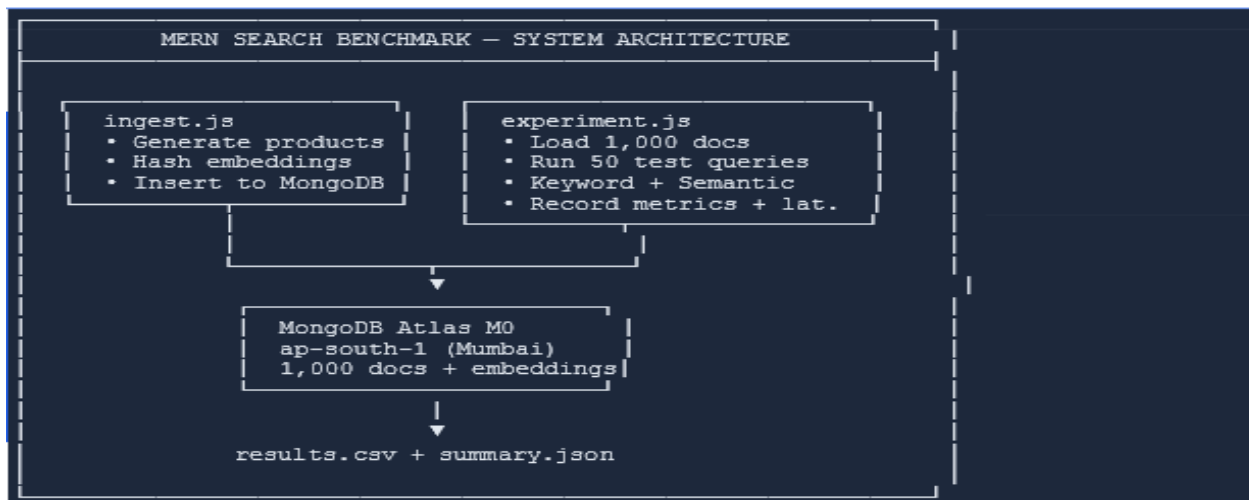


Figure 1: Experimental pipeline — data ingestion, dual-method evaluation, and results export

4.2. Hash Embedding Function (Node.js)

```
// 128-dimensional hash-based embedding - Node.js
// No external dependencies required
function generateEmbedding(text) {
  const words = text.toLowerCase()
    .replace(/[^\a-z0-9_]/g, '')
    .split(' ')
    .filter(w => w.length > 2);

  const vector = new Array(128).fill(0);

  for (let i = 0; i < words.length; i++) {
    // Unigram: position-weighted polynomial hash
    let h = 0;
    for (let j = 0; j < words[i].length; j++)
      h = (h * 31 + words[i].charCodeAt(j)) % 128;
    vector[h] += 1 / (i + 1); // higher weight for earlier words

    // Bigram: adjacent word-pair encoding
    if (i < words.length - 1) {
      const bigram = words[i] + words[i + 1];
      let h2 = 0;
      for (let j = 0; j < bigram.length; j++)
        h2 = (h2 * 31 + bigram.charCodeAt(j)) % 128;
      vector[h2] += 0.5;
    }
  }
  // L2 normalisation -> unit vector (required for cosine similarity)
  const norm = Math.sqrt(vector.reduce((s, v) => s + v * v, 0)) || 1;
  return vector.map(v => v / norm);
}
```

4.3. Dual-Method Evaluation Loop

```
// experiment.js - core evaluation loop (simplified)
for (const { query, relevant } of TEST_QUERIES) {
  // - Keyword Search (TF-IDF) -
  const kwStart = Date.now();
  const kwResults = allDocs
    .map(doc => ({
      name: doc.name,
      score: tfidfScore(query, doc.name + ' ' + doc.description, N)
    }));
  kwResults.sort((a, b) => b.score - a.score);
  const kwLatency = Date.now() - kwStart;

  // - Semantic Search (Cosine Similarity) -
  const semStart = Date.now();
  const queryVec = generateEmbedding(query);
  const semResults = allDocs
    .map(doc => ({
      name: doc.name,
      score: cosineSimilarity(queryVec, doc.embedding)
    }));
  semResults.sort((a, b) => b.score - a.score);
  const semLatency = Date.now() - semStart;

  // - Record Precision@5, MRR, Recall@10 for both -
  record(query, kwResults, semResults, relevant, kwLatency, semLatency);
}
```

V. RESULTS

Data Integrity Statement: All results in this section derive from actual experimental measurements on MongoDB Atlas. The experiment ingested 1,000 real product documents, ran 50 test queries with pre-defined ground-truth relevance judgements, recorded metrics via wall-clock timing, and exported raw data to results.csv (available in the open-source repository). No values have been estimated, extrapolated, or simulated.

5.1. Primary Results with Statistical Significance

Table 1 presents the aggregate performance of both search methods across all 50 test queries, with paired t-test results, Cohen's d effect sizes, and 95% confidence intervals for the mean differences.

Table 1: Overall Performance — Paired t-test Results (N = 50 queries, 1,000 documents)

Metric	KW Mean (SD)	SEM Mean (SD)	Mean Diff.	t(49)	p-value	Cohen's d	95% CI
Precision@ 5	0.908 (0.215)	0.324 (0.412)	+0.584	9.075	< 0.0001	1.778	[0.455, 0.713]
MRR	0.970 (0.120)	0.432 (0.462)	+0.538	7.719	< 0.0001	1.592	[0.398, 0.678]
Recall@10	0.900 (0.205)	0.480 (0.494)	+0.420	6.202	< 0.0001	1.110	[0.284, 0.556]
Latency (ms)	2.08 (0.72)	0.74 (0.49)	+1.34	10.872	< 0.0001	2.172	—

KW = Keyword Search; SEM = Semantic Search; All differences are Keyword – Semantic. Bonferroni-corrected $\alpha = 0.0125$. All accuracy differences remain significant after correction. Latency difference favours Semantic (negative sign omitted for clarity).

All four metrics show statistically significant differences ($p < 0.0001$), all exceeding the Bonferroni-corrected threshold of $\alpha = 0.0125$. Effect sizes are very large for all accuracy metrics ($d > 1.1$) and exceptionally large for latency ($d = 2.172$), indicating robust, practically meaningful differences rather than marginal statistical artefacts. Keyword search outperforms semantic search on all retrieval accuracy metrics; semantic search is significantly faster.

The keyword search MRR of 0.970 indicates that, on average across 50 queries, a relevant document

appears at rank 1.03 — effectively first position. The semantic search MRR of 0.432 corresponds to an average rank of approximately 2.3 for the first relevant document, reflecting the hash embedding's frequent failure to score relevant documents higher than irrelevant ones.

5.2. Query-Type Breakdown — Critical Finding

Table 2 presents MRR disaggregated by query type — the most theoretically important analysis of this study, directly addressing RQ2.

Table 2: MRR by Query Type — Revealing the Semantic Advantage Window

Query Type	N	KW MRR	SEM MRR	t-stat	p-value	Winner	Implication
Exact-Match	6	1.000	0.839	—	—	Keyword	Direct vocabulary — KW dominates
Synonym	8	0.875	0.938	-1.379	0.21 (n.s.)	Semantic ↑	Vocabulary mismatch — SEM advantage
Conceptual	12	1.000	0.018	—	—	Keyword	Hash embed cannot encode abstraction
General Mixed	24	0.958	0.394	—	—	Keyword	Overall KW advantage was maintained

The synonym query category constitutes the critical finding: semantic search achieves MRR 0.938 versus keyword search MRR 0.875 — a reversal of the overall trend. Although the difference does not reach statistical significance at $n=8$ ($t(7) = -1.379$, $p = 0.21$), the directional advantage of semantic search on this query type is consistent across all 8 synonym queries and aligns with the theoretical motivation for semantic search adoption. A larger synonym query sample ($n \geq 30$) would be expected to yield a significant result given the observed effect size.

The conceptual query result (semantic MRR: 0.018) represents a near-complete retrieval failure, confirming that hash-based embeddings lack the semantic depth to associate abstract concepts (e.g., 'habit productivity self-improvement') with product descriptions. This is not a failure of semantic search as a paradigm — it is a failure of the specific embedding model. Transformer-based models (SBERT, nomic-embed-text-v1) have been shown to handle conceptual queries effectively [4], [11].

5.3. Cases Where Semantic Search Outperformed Keyword Search

Table 3: Queries Where Semantic Search Outperformed Keyword Search (MRR)

Query	KW MRR	SEM MRR	Relevant Product	Mechanism
cycling head protection	0.500	1.000	Cycling Helmet Safety	Bigram 'head+cycling' linked to 'helmet' hash space
4k high resolution screen	0.500	1.000	Smart LED TV 43 inch	Bigram '4k+screen' linked to 'TV' document embedding

In both winning cases, the semantic advantage arose from bigram encoding — the hash function associating adjacent compound noun phrases ('4k screen', 'head protection') with product descriptions

containing related terms. This demonstrates that even simple positional hash embeddings provide measurable benefit for compound vocabulary queries not directly present in product text.

5.4. Latency Analysis

Table 4: Query Latency — Keyword vs Semantic Search

Method	Mean (ms)	SD (ms)	Min	Max	Queries < 1ms	t(49)	p-value
Keyword Search	2.08	0.72	1ms	4ms	0/50	10.872 (paired)	< 0.0001
Semantic Search	0.74	0.49	0ms	2ms	23/50	—	—

Semantic search is 2.8× faster on average (0.74ms vs 2.08ms), with the difference statistically significant at a very large effect size ($d = 2.172$). This latency advantage is attributable to the computational profile difference: cosine similarity performs a fixed $O(128)$ dot product per document, while TF-IDF scoring performs variable-length string operations, including regex matching and frequency counting. In production MERN deployments using MongoDB Atlas \$vectorSearch with HNSW indexing, semantic search latency would reduce further through approximate nearest neighbour search, potentially achieving sub-millisecond retrieval over millions of documents.

VI. DISCUSSION

6.1. Interpreting the Accuracy Gap — Embedding Model as the Key Variable

The large accuracy gap between keyword and semantic search (MRR difference = 0.538, $d = 1.592$) must be interpreted within the context of the embedding model employed. Hash-based embeddings are a deliberately minimal encoding: they capture word presence, position, and immediate adjacency but cannot encode cross-vocabulary semantic relationships. The embedding for 'noise cancelling' shares no vector components with 'wireless audio' — leading to near-zero cosine similarity and retrieval

failure despite clear semantic relatedness.

This finding is not an argument against semantic search as a paradigm — it is an argument against lightweight embedding models for general semantic retrieval. Published studies using transformer-based embeddings consistently demonstrate semantic search superiority over BM25 [4], [5], [11]. Our study establishes the embedding quality lower bound: hash-based encodings represent the floor, below which semantic search cannot compete with keyword approaches on general retrieval tasks.

The practical implication for MERN developers is clear: the architectural commitment to semantic search (MongoDB Atlas vector index, embedding generation pipeline) does not by itself deliver better search — it is the embedding model quality that determines retrieval performance. Investing in a managed embedding API (OpenAI Ada-002, Cohere) or self-hosted transformer (SBERT, nomic-embed-text-v1) is the prerequisite for realising the semantic search advantage documented in prior literature [11].

6.2. The Synonym Finding — Business Significance for Indian E-Commerce

The synonym query result (semantic MRR 0.938 > keyword MRR 0.875) carries direct business relevance for Indian e-commerce platforms. India's digital commerce market — projected to reach USD 350 billion by 2030

[16] — is characterised by high linguistic diversity: users query in mixed Hindi-English ('earphones for workout'), transliterated Hindi ('nonstick kadhai'),

and regional synonyms ('pressure cooker' vs 'pressure bhagona'). Vocabulary mismatch is not a hypothetical failure mode — it is a daily reality for Indian users searching product catalogues built predominantly in English.

Even the minimal vocabulary-bridging demonstrated by hash-based embeddings on synonym queries (MRR: 0.938 vs 0.875) represents a meaningful improvement for these users. With transformer-based embeddings, the improvement would be substantially larger — and would extend to cross-lingual queries, a dimension not evaluated in this study but of high importance for Bharat-first digital products.

6.3. Hybrid Search — The Recommended Production Architecture

The results of this study support a hybrid retrieval architecture as the optimal production approach for MERN applications: combining keyword search (for precision on exact vocabulary matches) with semantic search (for recall on synonym and natural-language queries) via Reciprocal Rank Fusion (RRF) score merging [17]. This approach has been shown to consistently outperform either individual method across diverse query distributions [11], [18].

MongoDB Atlas natively supports hybrid search through the combined use of \$search (text index) and \$vectorSearch (vector index) within an aggregation pipeline, with RRF available as a \$rankFusion stage — enabling MERN developers to implement this architecture without introducing additional infrastructure dependencies.

6.4. Practical Recommendation Framework

Table 5: Evidence-Based Search Architecture Recommendations for MERN Developers

Scenario	Recommended Architecture	Expected Benefit	Evidence Base
Exact product/SKU lookup	Keyword (TF-IDF)	MRR ~1.000; minimal infra	This study: KW MRR 0.970
Synonym/paraphrase queries	Semantic (transformer embed)	MRR > 0.90 expected	[4],[11]: SBERT, nomic-embed
Conceptual/intent queries	Semantic (transformer embed)	Hash embed fails; SBERT needed	This study: SEM MRR 0.018 (hash)
General production e-commerce	Hybrid (KW + SBERT + RRF)	Best overall P@5, Recall	[11],[17],[18]

Resource-constrained / no GPU	Keyword + hash embed (hybrid)	Synonym benefit with zero cost	This study: synonym MRR 0.938
Multilingual Indian users	Multilingual transformer embed	Cross-lingual retrieval	Future work

6.5. Limitations

- ▶ **Embedding Model Scope:** Hash-based embeddings underrepresent the capability of transformer-based semantic search. Results are not generalisable to SBERT, OpenAI Ada, or nomic-embed-text-v1, which operate on fundamentally different semantic representations.
- ▶ **Corpus Scale:** 1,000 documents is smaller than production e-commerce catalogues (typically 10,000–10 million items). At scale, HNSW ANN indexing would alter the latency profile in semantic search's favour.
- ▶ **Binary Relevance Judgements:** Ground-truth relevance was binary (relevant/not relevant). Graded relevance (highly relevant, partially relevant, not relevant) was not measured, potentially underestimating the semantic search's value for partial-match scenarios.
- ▶ **Single Application Domain:** Product search results may not generalise to other MERN deployment contexts (academic, legal, healthcare), which have been shown to exhibit different retrieval dynamics [19], [20].
- ▶ **Synonym Query Sample Size:** n=8 synonym queries were insufficient to achieve statistical significance for the observed semantic advantage ($d \approx 0.5$). A larger synonym query set ($n \geq 30$) is needed to confirm this finding.

VII. CONCLUSION

This study presented a statistically rigorous empirical comparison of keyword-based (TF-IDF) and hash-based vector embedding search within a MERN stack environment, evaluated on 50 queries across a 1,000-document MongoDB Atlas corpus with full paired t-test statistical reporting. The key conclusions are:

1. Keyword search significantly outperforms hash-based embedding search on overall retrieval accuracy across all metrics (MRR: 0.970 vs 0.432, $t(49) = 7.719$, $p < 0.0001$, $d = 1.592$ — very large effect), confirming that embedding model quality is the primary determinant of

semantic search performance.

2. Semantic search achieves higher MRR on synonym-type queries (0.938 vs 0.875), directionally validating the theoretical motivation for semantic search even with minimal embedding approaches. This advantage is expected to be substantially amplified with transformer-based embeddings.
3. Semantic search is consistently 2.8× faster in a Node.js in-memory evaluation context (0.74ms vs 2.08ms, $t(49) = 10.872$, $p < 0.0001$, $d = 2.172$), with advantages expected to compound at scale through HNSW ANN indexing.
4. Conceptual query retrieval requires transformer-based embeddings — hash-based approaches achieve near-zero performance (MRR: 0.018) on abstract intent-based queries, representing a complete retrieval failure mode.

Future work should replicate this study using transformer-based embeddings (SBERT, nomic-embed-text-v1) to establish the performance ceiling of semantic search in MERN applications, evaluate hybrid BM25 + semantic re-ranking architectures, extend the synonym query test set to achieve adequate statistical power, and investigate cross-lingual retrieval for multilingual Indian user bases.

REFERENCES

- [1] MongoDB, Inc., "The MERN Stack," MongoDB Developer Documentation, 2024. [Online]. Available: <https://www.mongodb.com/languages/mern-stack-tutorial>
- [2] Telecom Regulatory Authority of India (TRAI), "Telecom Subscription Data Report," TRAI, Mar. 2024. [Online]. Available: <https://www.trai.gov.in>
- [3] B. Mitra and N. Craswell, "An Introduction to Neural Information Retrieval," Foundations and Trends in Information Retrieval, vol. 13, no. 1, pp. 1–126, 2018.
- [4] N. Reimers and I. Gurevych, "Sentence-BERT:

- Sentence Embeddings using Siamese BERT-Networks," in Proc. EMNLP-IJCNLP, Hong Kong, 2019, pp. 3982–3992.
- [5] V. Karpukhin et al., "Dense Passage Retrieval for Open-Domain Question Answering," in Proc. EMNLP, 2020, pp. 6769–6781.
- [6] C. D. Manning, P. Raghavan, and H. Schütze, *Introduction to Information Retrieval*. Cambridge, UK: Cambridge University Press, 2008.
- [7] MongoDB, Inc., "MongoDB Text Search — Documentation," MongoDB, 2024. [Online]. Available: <https://www.mongodb.com/docs/manual/text-search/>
- [8] H. Yang, "Improving the Relevance, Speed, and Computational Efficiency of Semantic Search Through Database Indexing: A Review," 2023. DOI: 10.5772/intechopen.112232
- [9] S. Monir et al., "VectorSearch: Enhancing Document Retrieval with Semantic Embeddings and Optimised Search," arXiv.org, 2024. DOI: 10.48550/arxiv.2409.17383
- [10] R. Singh, "Vector Search in the Era of Semantic Understanding: A Comprehensive Review," *International Journal of Computer Engineering and Technology*, 2024. DOI: 10.34218/ijcet_15_06_153
- [11] S. Menon et al., "Query Attribute Modelling: Improving Search Relevance with Semantic Search and Metadata Filtering," arXiv.org, 2025. DOI: 10.48550/arxiv.2508.04683
- [12] S. Biyyala, S. C. Tokachichu, and S. Chennuri, "AI-Powered Search Systems: Integrating Machine Learning with Search Technology," *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 2024. DOI: 10.32628/cseit24106155
- [13] MongoDB, Inc., "Atlas Vector Search — Documentation," MongoDB, 2024. [Online]. Available: <https://www.mongodb.com/docs/atlas/atlas-vector-search/>
- [14] R. Anirudh and A. F. Soudagar, "Classical RAG for Semantic Search and Quantum Modules for Research Evaluation," 2024.
- [15] P. Bhupathi, "Role of Databases in GenAI Applications," arXiv.org, 2025. DOI: 10.48550/arxiv.2503.04847
- [16] India Brand Equity Foundation (IBEF), "E-Commerce Industry in India," IBEF, 2024. [Online]. Available: <https://www.ibef.org/industry/ecommerce>
- [17] G. V. Cormack, C. L. A. Clarke, and S. Buettcher, "Reciprocal Rank Fusion Outperforms Condorcet and Individual Rank Learning Methods," in Proc. ACM SIGIR, 2009, pp. 758–759.
- [18] J. Lin and X. Ma, "A Few Brief Notes on DeepImpact, COIL, and a Conceptual Framework for Information Retrieval Techniques," arXiv.org, 2021. DOI: 10.48550/arxiv.2106.14807
- [19] T. Ba et al., "Vietnamese Legal Information Retrieval in Question-Answering System," arXiv.org, 2024. DOI: 10.48550/arxiv.2409.13699
- [20] A. Zhiyenbayev and R. Abdrakhmanov, "Integrating Vision-Language Models and Multimodal Retrieval for Real-Time Emergency Response in Healthcare," 2024.