

A Comparative Analysis of Keyword, Vector, and Hybrid Search Strategies for Enterprise Document Intelligence Platforms

Tanmay Mangesh Pawar, Aarti Sanjay Gawai

Student, Dept. of Information Technology and Computer Science, D.G Ruparel College of Arts, Science and Commerce,
Mumbai, Maharashtra, India

Assistant Professor, Dept. of Information Technology and Computer Science, D.G Ruparel College of Arts,
Science and Commerce, Mumbai, Maharashtra, India pawartanmay045@gmail.com, aartigawai1611@gmail.com

Abstract - Organizations today store large amounts of information in digital documents such as company policies, technical manuals, reports, contracts, employee records, research documents, and product documentation. As the amount of digital information increases, finding the required information quickly and accurately becomes an important challenge. Search technology plays an important role in solving this problem.

Three important approaches used for document retrieval are keyword-based search, vector-based search, and hybrid search. Keyword search mainly identifies documents by matching words in the user query with words present in documents. Vector search represents text as numerical embeddings and retrieves documents based on semantic similarity. Hybrid search combines keyword and vector retrieval to use both exact matching and semantic understanding.

This study presents a comparative analysis of these three search strategies for enterprise document intelligence platforms. The research uses secondary data obtained from existing studies published in the area of information retrieval, semantic search, enterprise search, and hybrid retrieval. The selected studies are examined according to their retrieval approach, application area, evaluation methods, reported findings, advantages, and limitations.

The analysis shows that keyword search remains useful for exact terms, technical identifiers, and situations where fast retrieval is important. Vector search provides advantages for natural-language queries and situations where the meaning of the query is more important than exact word matching. Hybrid search combines both approaches and can provide a more flexible retrieval mechanism for mixed query types. However, hybrid systems also introduce additional complexity and computational requirements.

The study concludes that there is no single search strategy that is suitable for every enterprise retrieval problem. The appropriate approach depends on the type of documents, query characteristics, accuracy requirements, response-time requirements, and available system resources.

Key Words: Keyword Search, Vector Search, Hybrid Search, Document Retrieval, Semantic Search, Information Retrieval, Enterprise Document Intelligence, BM25.

1. INTRODUCTION

1.1 Background

The amount of digital information produced by organizations has increased significantly in recent years. Businesses store information in many different forms, including reports,

employee guidelines, financial documents, technical manuals, product information, contracts, emails, research documents, and frequently asked questions. Although storing these documents digitally makes information easier to preserve, finding the correct information can become difficult when the number of documents becomes very large.

Enterprise document intelligence platforms are designed to help organizations search, retrieve, and understand information stored in their document collections. These systems can use traditional information retrieval techniques as well as modern artificial intelligence methods. The retrieval component is particularly important because the quality of the final result depends heavily on whether the system can identify the most relevant documents or passages.

Traditional keyword-based search is one of the oldest and most widely used retrieval approaches. It searches for words or phrases appearing in the user's query and attempts to find documents containing those terms. Methods such as BM25 are commonly used to rank documents according to the importance and frequency of query terms.

Keyword search is effective when the user knows the exact terminology used in the document. For example, searching for a specific employee ID, product code, error number, or policy name can benefit from exact term matching. However, keyword search may have difficulty when the user expresses the same concept using different words.

For example, consider the query:

"How can I recover access to my account?"

A document may contain the phrase:

"Procedure for restoring access to a disabled user account."

A simple keyword system may not recognize the relationship between these two expressions as effectively as a semantic retrieval system.

Vector search attempts to solve this problem by representing documents and queries as numerical vectors called embeddings. Documents with similar meanings can be located close to each other in a vector space. This allows vector-based systems to retrieve semantically related information even when the exact words used in the query do not occur in the document.

However, vector search also has limitations. A semantically similar document may not necessarily contain the exact information required by the user. Vector retrieval can also be less suitable for some exact identifiers, codes, names, and highly specialized terms.

Hybrid search combines keyword and vector retrieval. The purpose is to use the strengths of both approaches. A hybrid system can use lexical matching for exact terms while using vector similarity for semantic relationships. Recent research on enterprise document search and hybrid retrieval continues to

investigate this combination, including approaches involving BM25, dense retrieval, rank fusion, and reranking.

Therefore, understanding the differences between keyword, vector, and hybrid search is important when designing enterprise document intelligence platforms.

1.2 Keyword-Based Search

Keyword-based search retrieves documents using terms contained in the user's query. Traditional information retrieval systems commonly use an inverted index to efficiently locate documents containing particular terms.

BM25 is a popular ranking function used in lexical information retrieval. It considers factors such as the frequency of a query term in a document and the frequency of the term across the document collection.

The main advantage of keyword search is its ability to handle exact terms. It is particularly useful when users search for:

- Employee IDs
- Product codes
- Error codes
- Policy numbers
- Technical terms
- Names
- Specific phrases

However, keyword search generally depends on the relationship between the words used in the query and the words stored in the document. When users use synonyms, paraphrases, or different terminology, relevant documents may be missed.

1.3 Vector-Based Search

Vector search uses machine-learning models to convert text into numerical representations called embeddings.

Instead of asking only:

"Does this document contain the same words?"

vector search attempts to determine:

"Does this document have a similar meaning?"

For example, the following two statements use different words:

- "How do I change my password?"
- "What are the steps for resetting a forgotten password?"

A semantic retrieval system may recognize that both statements refer to a similar concept.

Vector search is therefore useful for natural-language questions, semantic search, document discovery, and applications where users may not know the exact terminology used in the document.

However, vector search can sometimes return documents that are related to the topic but do not contain the exact answer required. This is particularly important in enterprise environments where technical codes, product identifiers, policy numbers, and other exact terms may be important.

1.4 Hybrid Search

Hybrid search combines two or more retrieval signals, commonly keyword retrieval and vector retrieval.

A simplified hybrid architecture can be represented as:

User Query → Keyword Search + Vector Search → Result Fusion → Ranked Documents

Keyword retrieval provides lexical information, while vector retrieval provides semantic information.

One common method for combining rankings is Reciprocal Rank Fusion (RRF). Instead of directly comparing scores produced by two different retrieval systems, RRF combines the rankings generated by those systems.

Recent research demonstrates several forms of hybrid retrieval. Prasad et al., for example, describe a document-management framework combining dense vector search, BM25 keyword search, and RRF.

Hybrid retrieval is attractive for enterprise systems because enterprise queries can contain both ordinary natural language and exact technical information.

1.5 Problem Statement

Enterprise organizations maintain large collections of digital documents, making accurate and efficient information retrieval increasingly important. Keyword search can provide strong exact matching but may have difficulty with semantic relationships and different wording. Vector search can identify semantically related documents but may not always provide reliable exact matching for identifiers and specialized terminology. Hybrid search attempts to combine these two approaches.

However, existing studies evaluate retrieval strategies using different datasets, domains, metrics, and system configurations. This makes it difficult to obtain a simple overall understanding of the strengths and limitations of keyword, vector, and hybrid search.

Therefore, this study compares these three retrieval strategies using findings from existing research in order to identify their characteristics, advantages, limitations, and suitable applications in enterprise document intelligence platforms.

1.6 Need for the Study

The study is required for several reasons.

First, organizations increasingly depend on digital documents for their daily operations. Employees need to find information quickly without manually searching through large numbers of files.

Second, different search techniques behave differently for different types of queries. An exact product code may be handled well by keyword search, while a natural-language question may benefit from vector search.

Third, hybrid retrieval has become an important research area because it attempts to combine lexical and semantic retrieval.

Finally, developers and organizations need to understand the trade-offs between these approaches before selecting a retrieval strategy for an enterprise system.

1.7 Objectives

The main objectives of this study are:

1. To study keyword-based search for enterprise document retrieval.
2. To examine vector-based semantic search and its applications.
3. To analyze hybrid search techniques combining keyword and vector retrieval.
4. To compare the advantages and limitations of the three retrieval approaches.
5. To analyze findings reported by existing research studies.

6. To identify suitable retrieval strategies for different enterprise query types.

1.8 Research Questions

RQ1: How does keyword-based search differ from vector-based search for enterprise document retrieval?

RQ2: What are the major advantages and limitations of keyword, vector, and hybrid search?

RQ3: Can hybrid search provide a more balanced retrieval approach than keyword or vector search alone?

RQ4: Which search strategy is more suitable for different types of enterprise queries and documents?

1.9 Scope of the Study

The study focuses on three retrieval strategies:

- Keyword-based retrieval
- Vector-based semantic retrieval
- Hybrid retrieval

The study concentrates on document and information retrieval in enterprise-oriented environments.

The research does not attempt to develop a complete enterprise document intelligence platform. It also does not train a new embedding model or conduct a new large-scale retrieval experiment.

Instead, the study analyzes findings reported in existing research and provides a comparative interpretation.

2. LITERATURE REVIEW

2.1 Overview

Research in information retrieval has traditionally focused on identifying relevant information from large document collections. Keyword-based methods have provided a reliable foundation for exact retrieval, while recent advances in natural-language processing have introduced semantic vector retrieval. More recent research has investigated hybrid systems that combine lexical and semantic methods.

The selected studies for this research cover enterprise search, semantic search, financial document retrieval, academic information retrieval, vector databases, RAG systems, and document intelligence.

2.2 Keyword-Based Retrieval

Keyword retrieval remains important because many enterprise queries contain terms that need to be matched accurately.

For example, an employee searching for a policy number or a developer searching for a specific error code may expect an exact match rather than a semantically similar document.

Traditional lexical retrieval therefore continues to have practical value even with the development of modern vector search systems.

Research comparing classical and hybrid search has also emphasized the importance of measuring retrieval performance rather than assuming that newer techniques are automatically better. Kapanke, Stephan, and Mertens specifically examine the transition from classical to hybrid search using a metric-driven approach.

2.3 Vector-Based Retrieval

Vector retrieval has become increasingly important with the development of embedding models.

An embedding converts text into a numerical representation. A query can then be compared with document embeddings using a similarity measure.

Bhoomika C. presents a semantic search engine based on AI embeddings, illustrating the use of embeddings for semantic document retrieval.

The major benefit of vector retrieval is that it can identify relationships based on meaning. This is useful when the query and document do not use exactly the same vocabulary.

However, semantic similarity alone does not guarantee that the retrieved document contains the exact information required by the user. This limitation is one reason hybrid approaches continue to receive attention.

2.4 Hybrid Retrieval

Hybrid retrieval combines different retrieval methods.

The basic idea is:

Lexical Retrieval + Semantic Retrieval = Hybrid Retrieval

Hybrid retrieval can be implemented in different ways. A system may combine the scores produced by two retrieval models or combine their ranked result lists using a method such as RRF.

Research by Tua and Fitrianah focuses directly on enterprise document retrieval and proposes a lightweight hybrid retrieval and reranking approach designed with privacy and on-premise deployment requirements in mind. Their work demonstrates that enterprise retrieval is not only about relevance but can also involve privacy, deployment, and computational considerations. Kathi and Garg examine enterprise application search using hybrid indexing and reasoning. Their work highlights the difficulty of searching across heterogeneous enterprise artifacts such as source code, configuration files, logs, tickets, and technical documentation.

2.5 Hybrid Retrieval in Document Intelligence

Hybrid retrieval is also being used in document-management and RAG systems.

Prasad, Yegireddi, and Anjaneyulu describe a framework that combines dense vector search using an embedding model with sparse BM25 retrieval and RRF. Their reported results demonstrate the practical use of combining lexical and semantic retrieval within an intelligent document-management system. This is relevant to enterprise document intelligence because modern systems often need to retrieve information before generating summaries, answers, or other AI-generated outputs.

2.6 Hybrid Retrieval in Financial Documents

Financial documents present a different retrieval challenge because they often contain long documents, tables, abbreviations, and technical terminology.

Wu and Lin examine sparse-dense fusion for financial document retrieval using corporate filings. Their study highlights an important methodological issue: the retrieval unit and document segmentation can affect the fairness of comparisons between sparse and dense methods. Their results also report benefits from combining BM25 and dense retrieval.

This suggests that retrieval performance cannot be separated completely from decisions about document segmentation and evaluation design.

2.7 Hybrid Retrieval in Academic Search

Htun and Khin investigate hybrid academic information retrieval and emphasize the trade-off between efficiency and retrieval quality. Their reported evaluation considers precision, recall, and nDCG while examining a system that combines lexical and semantic retrieval.

This demonstrates that retrieval systems should not be evaluated only on relevance. Response time and system complexity are also important.

2.8 Enterprise Natural-Language Search

Ngoc's 2025 study, *From Keywords to Context: An AI Agent for Natural Language Document Lookup in the Enterprise*, focuses on moving from traditional keyword-based lookup toward natural-language document access in enterprise settings. The paper was published in the 2025 RIVF International Conference on Computing and Communication Technologies.

This direction reflects a broader change in enterprise search: users increasingly expect to ask questions in ordinary language rather than constructing precise keyword queries.

2.9 Hybrid Querying and Vector Systems

Zhu et al. examine hybrid queries that combine vector search with structured filters such as attribute and range filtering. Their study evaluates multiple algorithms using large-scale datasets and different system conditions. The work demonstrates that "hybrid" retrieval can also refer to combining vector similarity with structured constraints, rather than only combining BM25 with embeddings.

This distinction is important because hybrid retrieval is a broad concept rather than a single fixed algorithm.

3. COMPARATIVE ANALYSIS OF EXISTING STUDIES

3.1 Summary of Selected Studies

No.	Study	Year	Main Focus
1	Tua & Fitrianah	2026	Privacy-preserving hybrid enterprise retrieval
2	Kathi & Garg	2026	Hybrid indexing and reasoning for enterprise search
3	Wu & Lin	2026	Sparse-dense fusion for financial document retrieval
4	Kapanke, Stephan & Mertens	2026	Classical vs hybrid search
5	C.	2026	Semantic search using AI embeddings
6	Prasad, Yegireddi & Anjaneyulu	2026	Hybrid RAG and document management
7	Ngoc	2025	Natural-language enterprise document lookup
8	Ruparel & Bussari	2025	Benchmarking retrieval methods for enterprise RAG

9	Zhu et al.	2025	Experimental evaluation of hybrid vector queries
10	Htun & Khin	2025	Efficiency and design choices in hybrid academic retrieval

3.2 Comparison of Search Strategies

Factor	Keyword Search	Vector Search	Hybrid Search
Exact matching	Excellent	Moderate	Excellent
Semantic understanding	Low	Excellent	Excellent
Synonym handling	Limited	Strong	Strong
Natural-language queries	Moderate	Strong	Strong
IDs and codes	Strong	Can be weaker	Strong
Technical terminology	Strong	Variable	Strong
Implementation complexity	Low	Medium	High
Computational requirements	Low	Medium/High	Higher
Flexibility	Moderate	High	Very High
Enterprise suitability	High for exact retrieval	High for semantic retrieval	High for mixed retrieval

4. DISCUSSION

4.1 Keyword Search vs Vector Search

Keyword and vector retrieval solve related problems in different ways.

Keyword search focuses primarily on the words contained in the query. This makes it highly useful when exact terms are important.

Vector search focuses on semantic similarity. It can retrieve documents that express a concept using different terminology. Consider the following query:

"How can an employee take medical leave?"

A keyword system may search for the words "employee," "medical," and "leave."

A vector system may also identify a document titled:

"Employee Sick Leave Procedure"

even if the exact phrase "medical leave" does not appear.

Therefore, vector search can provide an advantage when users express the same concept using different words.

However, consider a query such as:

"Find policy HR-2026-017."

Here, exact matching is important. Keyword retrieval can be more appropriate.

This demonstrates why comparing the two approaches as universally better or worse is misleading.

4.2 Keyword Search vs Hybrid Search

Hybrid search extends keyword retrieval by adding another source of evidence.

Keyword retrieval can provide exact matches, while semantic retrieval can identify conceptually related documents.

For enterprise applications, this combination is useful because document collections often contain both ordinary language and structured technical information.

For example:

"Show me documents explaining error code ERR-451."

A keyword system can identify "ERR-451."

A vector system can identify documents explaining the underlying issue.

A hybrid system can use both signals.

4.3 Vector Search vs Hybrid Search

Vector search is particularly useful when meaning is important. However, enterprise documents often contain information that should not be treated only as semantic meaning.

Examples include:

- Product numbers
- Employee IDs
- Version numbers
- Error codes
- Policy numbers
- Technical abbreviations

Hybrid search can preserve the semantic benefits of vector retrieval while adding lexical evidence.

This is one reason recent enterprise retrieval studies investigate combinations of sparse and dense retrieval rather than relying on only one signal.

4.4 Query-Type Analysis

The appropriate search method can depend strongly on the type of query.

Query Type	Keyword	Vector	Hybrid
Exact name	High suitability	Medium	High
Product code	High	Low/Medium	High
Natural-language question	Medium	High	High
Synonym-based query	Low/Medium	High	High
Technical query	High	Medium	High
Broad information search	Medium	High	High
Mixed query	Medium	High	Very High

This comparison suggests that hybrid retrieval is particularly useful when the query type cannot be predicted in advance.

4.5 Retrieval Effectiveness

Retrieval effectiveness can be measured using metrics such as:

- Precision
- Recall
- F1-score
- Mean Reciprocal Rank
- nDCG
- Hit@K

Precision measures how many retrieved results are relevant.

Recall measures how many of the relevant documents were successfully retrieved.

F1-score provides a combined measure based on precision and recall.

nDCG considers both relevance and the position of retrieved results.

The selected studies use different metrics, meaning that numerical results from different papers should not simply be added together or treated as directly equivalent.

This is an important limitation of secondary comparative research.

4.6 Efficiency and Latency

Retrieval quality is not the only concern.

An enterprise search system also needs to return results within a reasonable amount of time.

Keyword search can be relatively efficient because mature inverted-index technology is highly optimized.

Vector search requires embedding generation and similarity search.

Hybrid retrieval performs multiple retrieval operations and therefore may introduce additional processing.

Htun and Khin's work illustrates this trade-off by considering both retrieval effectiveness and efficiency in hybrid academic information retrieval.

Similarly, enterprise-focused research has considered deployment and privacy constraints in addition to retrieval quality.

4.7 Complexity

A simple keyword system can be easier to design and maintain.

A vector system requires additional components such as:

- Embedding model
- Vector index/database
- Similarity calculation
- Embedding storage

A hybrid system requires both lexical and semantic components and an additional mechanism for combining their results.

Therefore, organizations should consider whether the improvement in retrieval quality justifies the additional complexity.

5. OVERALL FINDINGS

The comparative analysis produced the following major findings.

Finding 1: Keyword search remains relevant

Keyword retrieval should not be considered outdated. It remains useful for exact terms, identifiers, names, codes, and technical terminology.

Finding 2: Vector search improves semantic retrieval

Vector search provides an advantage when users express information using different words from those used in the documents.

Finding 3: Hybrid retrieval combines complementary strengths

Keyword and vector retrieval have different strengths. Combining them can provide a broader retrieval capability.

Finding 4: Hybrid search is not automatically superior

The effectiveness of hybrid retrieval depends on factors such as:

- Query type
- Dataset
- Document structure
- Embedding model
- Fusion method
- Retrieval configuration
- Computational resources

Finding 5: Evaluation methodology matters

Different studies use different datasets and metrics. Therefore, numerical results from separate studies should not be treated as if they came from one common experiment.

The work of Wu and Lin particularly demonstrates how document segmentation can influence the comparison between sparse and dense retrieval.

Finding 6: Enterprise retrieval has additional requirements

Enterprise search is not only about relevance. Privacy, access control, latency, scalability, deployment requirements, and system complexity can also affect the choice of retrieval strategy.

6. PRACTICAL APPLICATIONS

6.1 Human Resources

Enterprise HR systems can use keyword retrieval for:

- Employee IDs
- Policy numbers
- Leave codes

Vector retrieval can help employees ask natural-language questions.

Hybrid retrieval can support both.

6.2 Technical Support

Technical support systems often contain:

- Error codes
- Product names
- Technical documentation
- Troubleshooting guides

Keyword search is useful for exact error codes, while vector search can identify conceptually similar troubleshooting information.

Hybrid retrieval can combine both.

6.3 Legal and Compliance Documents

Legal documents often contain exact clauses and terminology. Keyword retrieval can locate specific terms, while vector retrieval can help identify documents discussing similar concepts.

Hybrid retrieval can therefore be useful for document discovery.

6.4 Research and Academic Documents

Researchers may search using either exact terms or natural-language descriptions.

For example:

"papers discussing privacy risks in cloud computing"
is more semantic than:

"cloud security privacy 2026."

Vector or hybrid search can therefore be useful for academic document discovery.

7. RESEARCH GAP

Although several studies examine keyword, vector, and hybrid retrieval, the existing literature is spread across different application areas and evaluation settings.

Some studies focus on enterprise retrieval, while others focus on financial documents, academic search, RAG systems, or vector databases.

The datasets, evaluation metrics, document structures, retrieval models, and system environments also vary.

This creates a challenge when attempting to make a direct comparison.

The literature therefore indicates a need for a consolidated comparison of the three retrieval strategies using common factors such as:

- Retrieval effectiveness
- Semantic understanding
- Exact matching
- Query flexibility
- Computational requirements
- Latency
- Complexity
- Enterprise suitability

This study addresses this gap by bringing the findings together into one comparative framework.

8. CONCLUSION

This research presented a comparative analysis of keyword, vector, and hybrid search strategies for enterprise document intelligence platforms.

Keyword search remains an important retrieval technique because it provides strong exact matching and is useful for names, identifiers, codes, technical terms, and other situations where precise lexical matching is required.

Vector search provides a different capability by focusing on semantic similarity. It can retrieve documents that express similar concepts even when the wording of the query and document is different. This makes it particularly useful for natural-language questions and semantic document discovery. Hybrid search combines keyword and vector retrieval. The reviewed research indicates that this approach can be useful when enterprise queries contain both semantic concepts and exact terminology. Hybrid retrieval can therefore provide a flexible solution for systems that must support different types of user queries.

However, hybrid retrieval should not automatically be considered superior in every situation. It introduces additional components, computational requirements, and design decisions. Its effectiveness depends on the dataset, query type, retrieval models, fusion technique, and evaluation method.

The overall analysis therefore suggests that the choice of retrieval strategy should be based on the requirements of the particular enterprise application. Keyword search remains suitable for exact retrieval, vector search is valuable for semantic retrieval, and hybrid search provides a balanced approach when both types of retrieval are required.

A controlled experiment using the same documents, queries, hardware, and evaluation metrics would be a useful next step. Such an experiment could provide a more direct numerical comparison of the three approaches.

REFERENCES

- [1] H. M. Tua and D. Fitriana, "A lightweight privacy-preserving reranker enables customizable hybrid retrieval for enterprise document search," *Discover Computing*, vol. 29, art. no. 89, 2026, doi: 10.1007/s10791-026-09993-z.
- [2] S. R. Kathi and A. Garg, "Enhancing Enterprise Application Search Using Retrieval-Augmented Generation: A Hybrid Indexing and Reasoning Approach," *European Journal of Information Technologies and Computer Science*, vol. 6, no. 1, pp. 13–19, 2026, doi: 10.24018/compute.2026.6.1.70183.
- [3] C. Wu and Y. Lin, "Evidence-Unit Fairness and the Limits of Query-Adaptive Sparse-Dense Fusion in Financial Document Retrieval," 2026.
- [4] L. Kapanke, P. Stephan, and R. Mertens, "Transitioning from Classical to Hybrid Search: A Metric-Driven and Explainable Approach," in *2026 International Conference on Semantic Computing (ICSC)*, pp. 221–225, 2026, doi: 10.1109/ICSC67292.2026.00037.
- [5] B. C., "Semantic Search Engine Using AI Embeddings," *International Journal of Scientific Research in Engineering and Management*, vol. 10, no. 6, 2026, doi: 10.55041/IJSREM64845.
- [6] D. R. Prasad, R. Yegireddi, and P. Anjaneyulu, "A Hybrid RAG Framework for AI Powered Personalized Document Summarization and Contextual Querying," in *2026 International Conference on Data Science and Business Systems*, 2026, doi: 10.1109/ICDSBS69077.2026.11634727.
- [7] H. L. Ngoc, "From Keywords to Context: An AI Agent for Natural Language Document Lookup in the Enterprise," in *2025 RIVF International Conference on Computing and Communication Technologies*, pp. 1151–1156, 2025, doi: 10.1109/RIVF68649.2025.11365175.
- [8] R. Ruparel and S. Bussari, "Survey and Benchmarking of Retrieval Methods for Enterprise Retrieval-Augmented Generation," in *International Conference on Computer and Applications (ICCA) – Proceedings*, 2025, doi: 10.1109/ICCA66035.2025.11430920.
- [9] J. Zhu, J. Yuan, K. Yang, X. Chen, S. Yu, H. Lv, Y. Li, and B. Zheng, "An Experimental Evaluation of Hybrid Querying on Vectors," *Proceedings of the VLDB Endowment*, vol. 19, no. 2, pp. 183–195, 2025, doi: 10.14778/3773749.3773757.
- [10] C. N. Htun and M. P. Khin, "Efficiency Trade-offs and Design Choices in Hybrid Academic Information Retrieval," in *2025 6th International Conference on Advanced Information Technologies (ICAIT)*, 2025, doi: 10.1109/ICAIT68809.2025.11236772.

BIOGRAPHIES



Mr. Tanmay Mangesh Pawar
M. Sc. Information Technology
Student, Dept. of IT/CS,
D. G. Ruparel College.



Ms. Aarti Sanjay Gawai
M. Sc. Information Technology
Asst. Prof., Dept. of IT/CS,
D. G. Ruparel College