
GENAI-BASED ENTERPRISE KNOWLEDGE MANAGEMENT ASSISTANT

¹ K Sunanda, ² Y Kavya Sree, ³ K Poojitha, ⁴ A Koushik

¹AssistantProfessor, ²³⁴Students

Department of IOT

Siddhartha Institute of Technology & Sciences, Narapally

kondapallisunanda.cse@siddhartha.co.in, 23tq1a6935@siddhartha.co.in, 23tq1a6932@siddhartha.co.in,
23tq1a6953@siddhartha.co.in

Abstract

Organizations generate and maintain large volumes of knowledge through documents, reports, policies, manuals, emails, project records, meeting notes, and internal communications. As the amount of information increases, employees may find it difficult to locate accurate and relevant knowledge quickly. Traditional enterprise knowledge-management systems often depend on keyword-based search and manually organized repositories, which can make information discovery time-consuming and inefficient.

The proposed GenAI-Based Enterprise Knowledge Management Assistant is an intelligent platform designed to help employees access organizational knowledge through natural-language interaction. The system uses Generative Artificial Intelligence to understand user questions and retrieve relevant information from authorized enterprise knowledge sources. Instead of requiring users to search through multiple documents or systems, the assistant can provide contextual answers based on available organizational information.

The platform can use document processing, semantic search, embeddings, retrieval-augmented generation, and large language models to understand and respond to enterprise queries. Documents can be collected from approved knowledge repositories, processed into searchable representations, and indexed for efficient retrieval. When a user submits a question, the system retrieves relevant information and uses the language model to generate a concise and context-aware response.

Security and access control are important components of the proposed system. Users should only receive information that they are authorized to access according to organizational roles and permissions. The system can maintain source references for generated answers, helping users verify information against the original enterprise documents. Conversation history, feedback, and usage analytics can also be used to improve the quality and usefulness of the knowledge assistant.

I. Introduction

Modern enterprises generate significant amounts of digital knowledge through business documents, technical manuals, policies, project files, reports, presentations, emails, and internal communications. This information is valuable because it contains organizational experience, procedures, decisions, and expertise. However, the increasing volume and distribution of information can make it difficult for employees to locate the right information at the right time.

Traditional enterprise knowledge-management systems generally use document repositories, folders, databases, intranets, and keyword-based search engines. Although these systems provide centralized storage, employees may still need to search through multiple documents to find specific information. Keyword search can also produce many irrelevant results when a user's question does not exactly match the terminology used in the stored documents.

Generative AI provides new opportunities for improving enterprise knowledge access. Large language models can understand natural-language questions and generate human-readable responses. However, using a language model alone can introduce risks such as inaccurate responses, outdated information, and unsupported statements. Therefore, enterprise knowledge assistants can benefit from retrieval-augmented generation, where relevant information is retrieved from approved organizational sources before generating a response.

The proposed GenAI-Based Enterprise Knowledge Management Assistant addresses these challenges by combining document ingestion, semantic search, retrieval, large language models, access control, and response generation. Users can ask questions using natural language rather than searching manually through large repositories. The assistant retrieves relevant organizational information and generates responses based on the available knowledge sources.

The main goal of the system is to improve knowledge accessibility while maintaining security, accuracy, and organizational control. By providing source references, role-based access, feedback mechanisms, and centralized knowledge management, the platform can help employees use enterprise information more effectively. The system can become an intelligent interface between employees and the organization's existing knowledge repositories.

II. Literature Survey

Enterprise knowledge management has traditionally focused on collecting, organizing, storing, and sharing organizational information. Knowledge-management systems commonly use document repositories, databases, intranets, content-management systems, and search engines. These approaches provide organizations with centralized access to information, but users may still experience difficulties when knowledge is distributed across multiple repositories or stored in large document collections.

Information-retrieval research has introduced techniques for improving the discovery of relevant documents. Traditional keyword-based approaches depend on matching words between user queries and documents. Semantic retrieval techniques, including vector representations and embeddings, can instead identify relationships between the meaning of a query and stored content. This allows systems to retrieve information even when the user's wording differs from the terminology used in the source document.

Large language models have significantly expanded the capabilities of conversational knowledge systems. These models can understand natural-language questions,

summarize information, and generate fluent responses. However, general-purpose language models may not contain an organization's private knowledge and can potentially generate unsupported information. This makes enterprise-specific grounding an important requirement for reliable knowledge assistants.

Retrieval-Augmented Generation (RAG) has emerged as an approach for combining information retrieval with generative AI. In a RAG architecture, relevant documents or document sections are retrieved from a knowledge repository and provided as context to a language model. The model then generates a response based on the retrieved information. This approach can improve grounding and allow enterprise knowledge bases to be updated without retraining the language model for every document change. Enterprise AI systems also require strong security and access-control mechanisms. Organizational documents may contain confidential business information, employee information, technical details, or proprietary knowledge.

Existing enterprise search and AI-assistant solutions demonstrate the potential of conversational interfaces, semantic retrieval, document analysis, and automated summarization. However, organizations may still face challenges involving knowledge fragmentation, source verification, access control, hallucination management, and knowledge freshness. The proposed system builds on these concepts by combining GenAI, semantic retrieval, RAG, enterprise document processing, role-based access control, source attribution, feedback, analytics, and centralized knowledge management.

III. System Analysis

System analysis for the proposed platform focuses on the requirements for building a secure and useful enterprise knowledge assistant. The system should provide authenticated access to employees and administrators. It should support integration with authorized enterprise knowledge sources such as documents, policies, manuals, and internal repositories. The document-processing module should extract and normalize relevant content from supported sources. Documents should be divided into meaningful sections and converted into searchable semantic representations. A retrieval engine should identify relevant information based on the meaning of the user's question. The system should apply user permissions before retrieving restricted enterprise content. A GenAI module should generate responses using the retrieved information as contextual evidence. Responses should include source references where appropriate to support verification. The platform should maintain conversation and feedback information according to organizational policies. An analytics module should monitor usage, query patterns, and response-quality indicators. Administrators should be able to manage knowledge sources, users, permissions, and system configurations. The system should also provide mechanisms for updating outdated knowledge and handling unsupported or low-confidence responses.

Existing System

Existing enterprise knowledge-management environments commonly use document repositories, shared drives, intranets, databases, enterprise search engines, and

content-management platforms. Employees generally locate information by browsing folders or entering keywords into search interfaces. These systems can provide access to large amounts of organizational information but may require users to know where relevant information is stored. Keyword-based search may return numerous results without understanding the actual meaning of the user's question. Employees may need to open several documents and manually compare information. Knowledge can also become fragmented across departments and different storage platforms. Some systems provide limited conversational interaction with enterprise information. Traditional systems generally do not generate contextual answers automatically. Maintaining document indexes and metadata can also require administrative effort. Access-control configurations may differ between repositories, making centralized knowledge access difficult. These limitations can increase the time required to locate organizational information. A GenAI-powered knowledge assistant can provide a more natural and efficient interface for accessing authorized enterprise knowledge.

Disadvantages of Existing System

- Heavy dependence on keyword-based search.
- Users may need to know where information is stored.
- Information can be distributed across multiple repositories.
- Search results may contain irrelevant documents.
- Employees may need to manually read multiple documents.
- Limited natural-language interaction.
- Limited contextual answer generation.
- Difficult to obtain concise answers from large document collections.

Proposed System

The proposed GenAI-Based Enterprise Knowledge Management Assistant provides a centralized conversational interface for accessing authorized organizational knowledge. The system allows administrators to connect approved enterprise knowledge sources such as documents, policies, manuals, reports, and internal repositories. A document-ingestion module extracts, cleans, and divides source content into meaningful sections. The processed content is converted into embeddings and stored in a vector-based knowledge index for semantic retrieval. When an employee submits a natural-language question, the system analyzes the query and retrieves relevant information from the authorized knowledge base. Access-control mechanisms verify whether the user is permitted to retrieve the selected information. The retrieved content is passed to a GenAI model as contextual information for response generation. The assistant produces a concise answer grounded in the retrieved enterprise knowledge and can provide source references for verification. Low-confidence or unsupported questions can be identified and handled appropriately instead of presenting unsupported information as fact. Conversation feedback can be collected to improve retrieval and response quality. Administrators can monitor knowledge sources, usage statistics, permissions, and system performance. The platform therefore provides an intelligent workflow for enterprise knowledge ingestion, retrieval, response generation, verification, and continuous knowledge improvement.

Advantages of Proposed System

- Provides natural-language access to enterprise knowledge.
- Reduces time spent searching through documents.
- Uses semantic retrieval instead of keyword matching alone.
- Supports Retrieval-Augmented Generation.
- Provides contextual and concise answers.
- Can provide references to original knowledge sources.
- Supports role-based access to enterprise information.
- Integrates multiple authorized knowledge repositories.
- Can be continuously updated as enterprise knowledge changes.


IV. Methodology

The methodology begins with authenticating the employee or administrator using the organization's approved identity system. Authorized enterprise knowledge sources are then connected to the platform according to defined access policies. The document-ingestion module collects supported documents and extracts their textual content. Extracted content is cleaned, normalized, and divided into meaningful chunks for efficient retrieval. Each chunk is converted into a semantic embedding and stored in the enterprise knowledge index. When a user submits a question, the query is processed and converted into a representation suitable for semantic retrieval. The retrieval engine searches the knowledge index for information relevant to the user's question. Access-control policies are applied to ensure that only authorized information is retrieved. Relevant content is passed to the GenAI model as contextual information for grounded response generation. The generated response is presented to the user along with available source references and confidence or limitation indicators. User feedback can be collected to evaluate response usefulness and identify retrieval problems. Administrators can analyze usage, update knowledge sources, review permissions, and maintain the knowledge base. The overall process creates a continuous cycle of knowledge ingestion, secure retrieval, GenAI response generation, verification, feedback, and improvement.

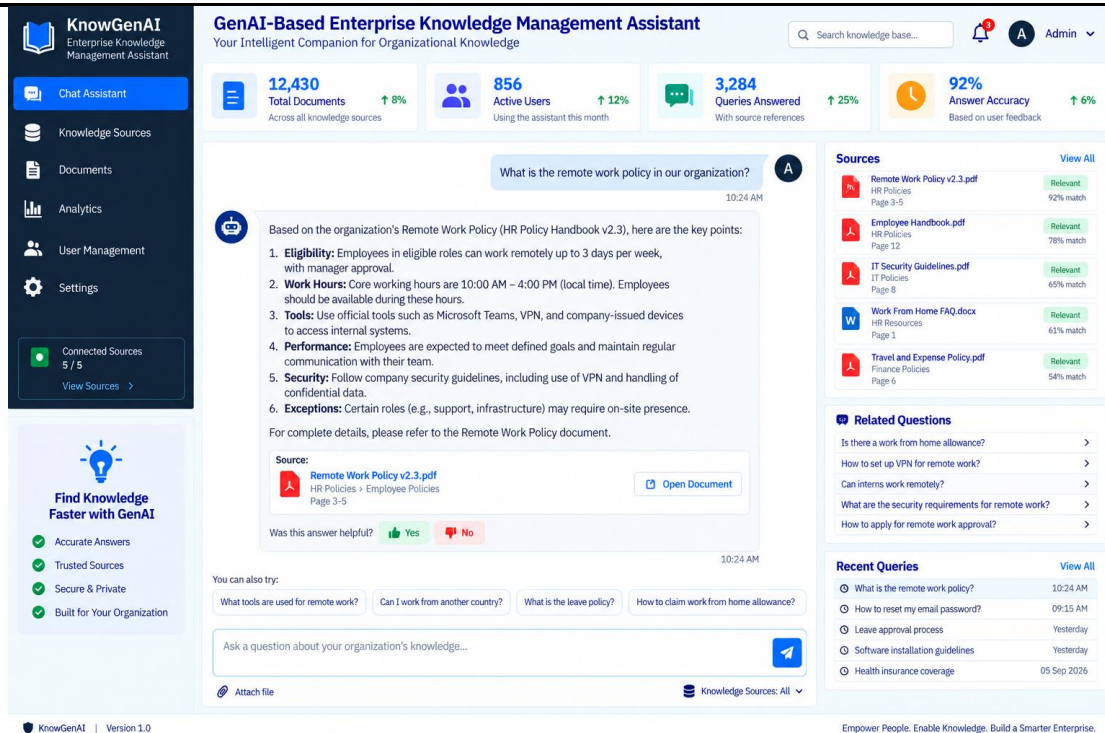
System Architecture

The system architecture consists of multiple interconnected layers designed to provide secure and intelligent enterprise knowledge management. The User Layer contains employees, knowledge managers, administrators, and authorized business users who interact with the assistant according to their organizational permissions. The User Interface Layer provides authentication, conversational chat, search, source viewing, feedback, and knowledge-management features. The Application Layer manages user sessions, roles, query processing, business logic, access control, logging, and communication between system components. The Knowledge Ingestion Layer connects to approved enterprise repositories and processes documents, manuals, policies, reports, and other authorized knowledge sources. The Document Processing Layer extracts text, cleans content, divides documents into meaningful chunks, and prepares metadata. The Embedding and Vector Search Layer converts knowledge segments into semantic representations and stores them in a searchable vector index.

The Query Processing Layer interprets user questions and generates retrieval queries for finding relevant enterprise information. The Authorization Layer verifies user roles and permissions before restricted information is retrieved or included in responses. The Retrieval-Augmented Generation Layer combines retrieved enterprise content with a GenAI model to produce grounded responses. The Source Attribution Layer provides references to relevant documents or knowledge sections where available. The Database Layer stores users, permissions, document metadata, conversation records, feedback, system configurations, and audit information. The Analytics Layer provides usage statistics, query trends, retrieval performance, feedback analysis, and knowledge-base monitoring. The Administration Layer allows authorized administrators to manage users, knowledge sources, permissions, models, and system settings. The architecture also includes security controls, audit logging, data protection, and knowledge-retention policies. The overall workflow follows the sequence Ingest → Process → Index → Retrieve → Generate → Verify → Learn, providing employees with secure and intelligent access to enterprise knowledge.



V. Result and output



VI. Conclusion

The GenAI-Based Enterprise Knowledge Management Assistant provides an intelligent and centralized approach to accessing and managing organizational knowledge. By combining Generative AI, semantic search, document processing, and Retrieval-Augmented Generation, the system enables employees to ask questions in natural language and receive relevant, contextual answers from authorized enterprise knowledge sources.

The proposed system reduces the limitations of traditional knowledge-management systems by minimizing manual document searching and providing faster access to relevant information. Source references, access-control mechanisms, and knowledge management features help improve the reliability, security, and usefulness of generated responses.

The platform also supports knowledge-source management, conversation history, user feedback, analytics, and administrative controls. These features help organizations monitor knowledge usage, identify information gaps, maintain updated content, and continuously improve the quality of the knowledge assistant.

References

- [1] Kumar, R. D., Prudhvairaj, G., Vijay, K., Kumar, P. S., & Plugmann, P. (2024). Exploring COVID-19 through intensive investigation with supervised machine learning algorithm. In Handbook of Artificial Intelligence and Wearables (pp. 145-158). CRC Press.
- [2] Swathi, B., Vijay, K., Sushanth Babu, M., & Dinesh Kumar, R. (2024, November). Machine Learning Techniques in Cloud Based Intrusion Detection. In The

International Conference on Artificial Intelligence and Smart Environment (pp. 557-564). Cham: Springer Nature Switzerland.

[3] Sv satyakrishna, shirisha rangu ,bhargavi nalacheruve.(2024) Prospective investigation on colorectal cancer with SMOTE on machine learning Algorithm

[4] Dr.G.Vishnu Murthy, BhargaviNalacheruve 1Professor, Department of computer Science & engineering, Anurag University, TS, India. 2Student, Department of computer Science & engineering, Anurag University, TS, India.

[5] V. N. S. Manaswini, K. K, C. Nigam, S. S. Ali, R. Niranjana, and Suman, “Real-Time Object Detection in Drone Surveillance Using YOLOv5,” in Proc. 2025 3rd Int. Conf. IoT, Communication and Automation Technology (ICICAT), Gorakhpur, India, 2025, pp. 1–6, doi: 10.1109/ICICAT68430.2025.11414670.

[6] B. Soundarya, V. N. S. Manaswini, M. Ayyakrishnan, R. D. Kumar, “Contextual Analysis of Big Data Analytics in Intelligent Transportation Frameworks,” in Intersection of Artificial Intelligence, Data Science, and Cutting-Edge Technologies: From Concepts to Applications in Smart Environment, Lecture Notes in Networks and Systems, vol. 1353, Cham: Springer, 2025, doi: 10.1007/978-3-031-88304-0_79.

[7] R. D. Kumar, V. N. S. Manaswini, “Applications of blockchain in smart cities: detecting fake documents from land records using blockchain technology,” in Blockchain for Smart Cities, Elsevier, 2021, pp. 105–117, doi: 10.1016/B978-0-12-824446-3.00017-X.

[8] Tejavath Veeramma, Badarla Anil, Guguloth Ravinder, “An advanced movie recommender using collaborative filtering and sentiment analysis,” International Research Journal of Modernization in Engineering Technology and Science, vol. 7, no. 7, July 2025, doi: 10.56726/IRJMETS81618.

[9] Ravi Kumar Banoth, Ramana Murthy B V, “Automatic crop recommendation system using LightGBM and decision tree machine learning models,” Journal of Machine and Computing, vol. 5, no. 1, pp. 343, Jan. 2025, doi: 10.53759/7669/jmc202505026.

[10] Ravi Kumar Banoth, Dr. B.V. Ramana Murthy, “Smart agriculture through IoT and machine learning for analyzing carbon footprints,” in Proc. Int. Conf. Computer Science and Communication Engineering (ICCSCE), Apr. 2025.

[11] Ravi Kumar Banoth, B. V. Ramana Murthy, “Soil image classification using transfer learning approach: MobileNetV2 with CNN,” SN Computer Science, vol. 5, art. no. 199, 2024, doi: 10.1007/s42979-023-02500-x.