WORLD COMET RESEARCH GROUP

International Journal of Applied Data Science & Modern Computing
Vol. 8, No. 1, 2025
Doi: 10.67228/30713498/IJADSMC-2025PI8M2D
PP. 01-17

Original Article

Large Language Model Integration for Enterprise Knowledge Management Platforms

Farhan Malik[1], Zara Ahmed[2]

[1]IT Manager, Systems Limited, Pakistan.

[2]Business Development Manager, Engro Corporation, Pakistan.

Received: 04-05-2025 Revised: 04-26-2025 Accepted: 05-03-2025 Published: 05-05-2025

ABSTRACT

Enterprise Knowledge Management Platforms (EKMPs) help organizations capture, organize, share, and utilize knowledge to improve decision-making and operational efficiency. Traditional knowledge management systems often face challenges such as data silos, unstructured information, limited contextual understanding, and ineffective search capabilities. The integration of Large Language Models (LLMs) addresses these limitations by enabling intelligent search, semantic understanding, automated content generation, and conversational interfaces. This study proposes an LLM-powered knowledge management framework that combines Retrieval-Augmented Generation (RAG), semantic embeddings, enterprise-specific language models, and vector databases to transform enterprise data into actionable knowledge. The framework includes data ingestion pipelines, knowledge repositories, embedding generation modules, retrieval systems, and conversational AI interfaces while emphasizing security, privacy, governance, scalability, and explainability. It also addresses challenges such as hallucination reduction, domain adaptation, and knowledge freshness. Experimental results demonstrate that LLM-based systems significantly improve retrieval accuracy, response relevance, knowledge reuse, and user satisfaction compared with traditional keyword-based approaches. These advancements enhance employee productivity, decision quality, collaboration, and innovation, positioning LLM-driven knowledge management systems as a key enabler of enterprise digital transformation.

KEYWORDS

Large Language Models (LLMs), Enterprise Knowledge Management, Artificial Intelligence, Retrieval-Augmented Generation, Semantic Search, Knowledge Discovery, Transformer Models, Enterprise AI.

ISSN: 3071-3498
PP: 01-17.

1. INTRODUCTION

1.1. Background and Motivation

In today's digital world, knowledge has emerged as a strategic resource for organizations to improve their innovation, efficiency, and competitiveness. In today's digital world, enterprises face the challenge of handling vast amounts of information generated from business transactions, customer interactions, emails, reports, research activities, and collaboration. Efficiently managing all this knowledge is crucial to enable accurate decision-making and to retain organizational knowledge. Typical Enterprise Knowledge Management Platforms focused on a method of storing, organizing, and retrieving information, but were constrained by keyword-based search capabilities and predefined indexing structures, which cannot effectively index unstructured and large data sources. In recent years, the field of Artificial Intelligence has witnessed significant progress, especially with the advent of Large Language Models (LLMs), which have opened new avenues for intelligent knowledge management. These models can follow context, understand intention, come up with human-like answers, condense complex information and carry out semantic search on huge repositories. The ability to embed the LLM into enterprise platforms opens the door to natural language access to the organization's knowledge base, streamlines information retrieval time, and enhances decision support. As an enterprise transitions to a data-driven business model, it is, therefore, becoming increasingly important to develop intelligent knowledge ecosystems that enable collaboration and knowledge sharing, continuous learning, and organizational transformation.

1.2. Role of Large Language Models in Knowledge Management

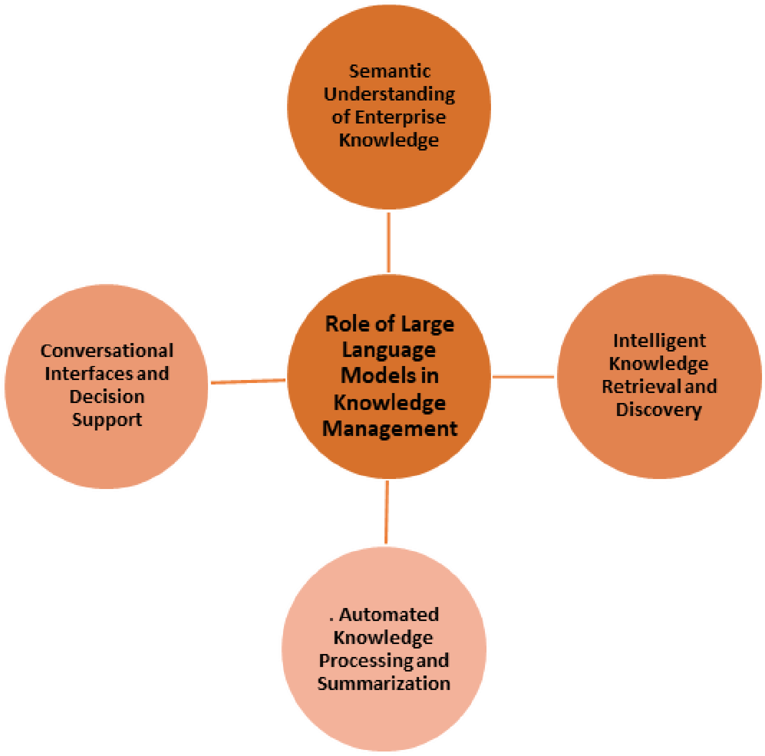

Fig 1 - Role of Large Language Models in Knowledge Management

ISSN: 3071-3498
PP: 01-17.

1.2.1. Semantic Understanding of Enterprise Knowledge

Large Language Models (LLMs) have become an important part of improving enterprise knowledge's semantic understanding. The most common traditional knowledge management systems are keyword matching and metadata-based indexing, but they don't capture the actual meaning and context behind the information. LLMs tackle these problems by understanding how language is structured within documents, the relationships between different sections, and the meaning behind the content. With this feature, companies can pull out information when it matches intent and meaning, not an exact word match. Employees are then able to access knowledge that is relevant to them more efficiently, regardless of the fact that different terminology or expressions are used within departments or between business units.

1.2.2. Intelligent Knowledge Retrieval and Discovery

LLMs have significantly aided intelligent knowledge retrieval and discovery, a vital aspect of knowledge management. LLMs have the capability to process natural language queries and extract relevant information from extensive corporate databases using sophisticated natural language processing capabilities. The Semantic Search mechanisms based on LLMs facilitate the identification of the relationships between the documents, reports, policies and knowledge records. This ensures retrieval accuracy and aids employees in finding the critical information they need easily. Moreover, LLM can help with knowledge discovery, as they can detect trends, patterns and insights that might not be obvious using traditional search techniques.

1.2.3. Automated Knowledge Processing and Summarization

LLMs can automate a range of information processing tasks, contributing to substantial enhancements in knowledge management. In enterprise repositories, there may be huge amounts of unstructured content such as documents, emails, meeting notes, technical manuals, research reports, and so on. LLMs can automatically summarize long content, extract key information, categorize documents, add metadata, and recognize key entities and relations. These abilities decrease manual work to organise and take care of knowledge assets. Automated summarization also allows workers and decision makers to rapidly grasp complex information in a manner that boosts productivity and makes knowledge more easily accessible throughout the organization.

1.2.4. Conversational Interfaces and Decision Support

Large Language Models can be used to create conversational interfaces that change the way enterprise knowledge systems are used. Employees can interact with the system via natural language queries, rather than having to navigate multiple databases or use complex search commands. The LLM understands what the user is trying to do, pulls the necessary details and produces responses that are relevant to the context and in conversational tone. This allows users to easily access the knowledge and improves the user experience. Further, LLMs can be used as a powerful tool for decision support, offering recommendations, answers to business-related inquiries, and actionable insights based on organizational knowledge. These work together to enhance organizations' decision making, collaboration and knowledge management environments.

ISSN: 3071-3498
PP: 01-17.

1.3. Large Language Model Integration for Enterprise Knowledge Management Platforms

The advent of Large Language Models (LLMs) in Enterprise Knowledge Management Platforms (EKMPs) marks a major leap in how organizations manage, capture, organize, retrieve, and use their knowledge assets. The mainstay of traditional knowledge management systems is the keyword search, manual categorisation and categorised databases, which are unable to deal with the increasing volume and complexity of enterprise information. Structured and unstructured data from emails, reports, project documents, customer interactions, research records, and collaborative platforms are produced in large quantities by modern organizations. Combining LLMs into enterprise knowledge environments can help these systems to process and understand information at a semantic level, which can overcome many of the limitations of conventional approaches. LLMs have the ability to comprehend natural language, understand context, extract meaning, and produce human-like responses, thereby simplifying the process of accessing knowledge. In the context of an enterprise knowledge management system, LLMs serve as smarter mediators between users and enterprise knowledge repositories. They can include some semantic search features that enable the employees to find information based on meaning and purpose instead of exact matches. This vastly enhances the accuracy of retrieval and decrease the time taken to find information relevant to their search. In addition, the LLM can be used to classify documents, summarize content, extract information, and recommend knowledge, which can help organizations manage large volumes of information from their repositories more effectively. Implementing Retrieval-Augmented Generation (RAG), which leverages enterprise document retrieval and advanced language generation, further improves system performance. This method ensures that the responses are based on the latest organizational information, which helps to minimize hallucinations and enhance factual correctness. In addition, conversational interfaces with natural language interaction with enterprise systems are possible with LLM integration. These interfaces make the knowledge of the organization more accessible, enhance user experience, and promote a deeper engagement with knowledge resources among various departments. The integration also enables the discovery of expertise, collaborative learning, and intelligent decision support by making users aware of the information they need as well as best practices to follow within the organization. LLMs help turn static repositories into dynamic and interactive knowledge ecosystems, improving productivity, speeding up decision-making, and encouraging ongoing innovation. The merger of LLM into enterprise knowledge management systems is projected to be an essential ingredient of knowledge-driven, adaptive, and more intelligent enterprises in the future as the artificial intelligence technologies continue to evolve.

2. LITERARY SURVEY

2.1. Evolution of Enterprise Knowledge Management Systems

The development of Enterprise Knowledge Management Systems (EKMS) has been motivated by the need to effectively capture, organise and leverage organisation knowledge. The first knowledge management systems were mainly concerned with functions such as storing, archiving and retrieving documents from relational databases and based on the metadata contained in them. With these systems, organizations were able to store vast amounts of structured data and access

ISSN: 3071-3498
PP: 01-17.

documents via keyword search. They had a hard time, however, in dealing with unstructured information like e-mails, reports, multimedia information, and informal information of the organization. To address these shortcomings, the solution of ontology-based knowledge representation and Semantic Web has been proposed, which allows the creation of more meaningful relations between information assets, and the facilitation of the knowledge organization. These approaches improved the semantic comprehension, but at the same time, it took a lot of manual work to create, maintain and update the ontologies. Later, the machine learning techniques were integrated into knowledge management platforms to automate document classification, recommendation, clustering, and information extraction processes. These developments decreased reliance on people and enhanced the availability of enterprise information. Recently, with the advent of deep learning and transformer architectures, knowledge management has been revolutionized by the ability to understand the context, do logical reasoning, and engage with organizational knowledge in intelligent ways. Today's EKMS platforms are capable of handling both structured and unstructured data, can handle conversational interfaces and knowledge discovery. The evolution illustrates how storage-based systems have been transformed to become smart, knowledge-based ecosystems that enable data-based decision making and organizational learning.

2.2. Large Language Models for Knowledge Discovery

As a result of their capabilities in understanding and generating natural language, Large Language Models (LLMs) are becoming very powerful tools in enterprise knowledge discovery. These models, like BERT, GPT, T5, and LLaMA, which are based on transformer structures, can process a substantial amount of textual information and extract intricate semantic connections from enterprise data repositories. Whereas traditional keyword search systems rely on queries consisting of specific keywords, LLMs can interpret the context and intent behind the query, making it possible to ask questions in natural language that get meaningful answers. They can be used for various knowledge discovery applications such as document summarization, information extraction, semantic search, content classification, topic modeling, and question-answer systems. The ability of LLMs to create semantic embeddings allows them to determine the similarities in concepts between documents, even if they do not contain the same exact words or phrases, thus enhancing retrieval performance. Furthermore, LLMs can facilitate automated construction of a knowledge graph from unstructured text sources by detecting entities and relationships. They can handle multiple data formats, enabling businesses to gain insights, patterns, and knowledge from massive enterprise data stores. Research studies have shown that LLM-powered knowledge discovery systems can provide better precision, recall and user satisfaction than traditional information retrieval systems. As a result, companies are increasingly looking to cut down employee work time, expedite their decisions, and convert business knowledge into usable information.

2.3. Retrieval-Augmented Generation in Enterprise Systems

Large Language Models can be integrated into enterprise knowledge management systems, and one of the most popular techniques for doing so is called Retrieval-Augmented Generation (RAG). In the past, language models generated responses mainly from what they've been taught, and

ISSN: 3071-3498
PP: 01-17.

as time goes on, that can change or grow incomplete. RAG is designed to overcome this challenge by integrating information retrieval with generative AI. In a RAG approach, the relevant documents are first retrieved from enterprise knowledge bases or databases or document management systems and then fed as contextual information to the language model. This mechanism enables the model to produce responses based on the up-to-date information in the organisation instead of just on the pre-saved data. This means that RAG can create more accurate outputs, lower the risk of hallucination, and provide greater reliability in the factual content of the outputs. Moreover, it allows organizations to keep knowledge up-to-date without the need for expensive model retraining when enterprise data is updated. One of the major benefits of RAG is that it can be explained, for example, by providing a link back to the source document to help users determine the accuracy of the information and build trust in the system. It has been proven that the RAG-based approaches significantly enhance enterprise Q&A, customer service, technical documentation support, and knowledge discovery systems. RAG is thus becoming a key technology for the development of intelligent, trustworthy, and scalable enterprise knowledge management platforms.

2.4. Research Gaps and Opportunities

While there are great strides with LLM and enterprise knowledge management technologies, there are still a number of research challenges and opportunities. A major worry is the protection of data privacy and security, particularly when AI systems are handling sensitive organizational data. The confidence and transparency of enterprises depend on the effective governance mechanisms for accessing knowledge assets, complying with regulations and ensuring security of confidential knowledge assets. One of the significant challenges is controlling hallucinations, where language models might provide incorrect or misleading information which could impact organizational decision-making. Explainability is also a crucial challenge as many deep learning models are black boxes which means that inference is hard for the users to understand. However, having models change with time adds even more complexity, especially if these organisations are creating many dynamic data.

Moreover, the problem of cross-domain, multi-departmental and multi-source data integration, which involves data format, terminology and structures, is difficult. Scalability is another key issue, particularly for large organisations with extensive knowledge bases and many users accessing them simultaneously. Most existing studies have concentrated on a specific application, like a chatbot, document retrieval or recommendation systems, and not offered enterprise-wide integration frameworks. Hence, the potential for building complete architectures that integrate LLM, RAG, knowledge management governance, and scalable infrastructure to build intelligent, secure, and efficient knowledge management solutions to meet real-world business needs.

ISSN: 3071-3498
PP: 01-17.

3. METHODOLOGY

3.1. Proposed LLM-Enhanced Enterprise Knowledge Management Architecture

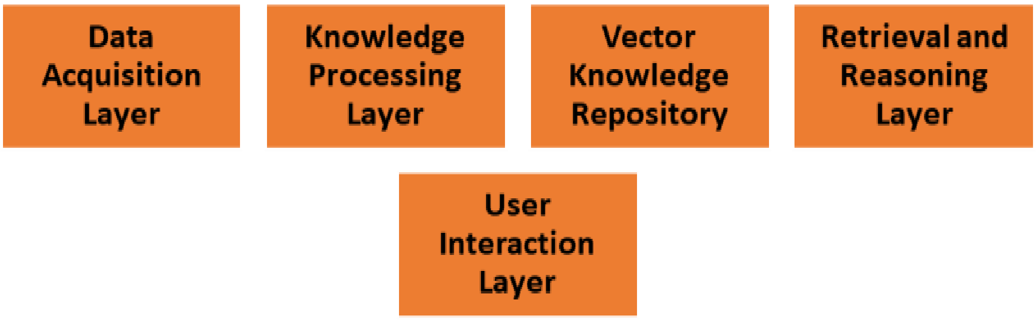

Fig 2 - Proposed LLM-Enhanced Enterprise Knowledge Management Architecture

3.1.1. Data Acquisition Layer

The Data Acquisition Layer is a key part of an enterprise Knowledge Management (KM) architecture that pulls data from multiple sources within an organization. The sources include databases, enterprise applications, document management systems, emails, reports, web portals, cloud storage platforms, and collaborative tools. The main goal of this layer is to collect structured and unstructured data, and to ensure data quality, consistency and completeness. Data ingestion mechanisms continuously pull new data and updates from a variety of repositories, allowing the knowledge management system to keep up to date with organizational knowledge. This layer also does some basic data cleansing and data integration to facilitate further processing.

3.1.2. Knowledge Processing Layer

The Knowledge Processing Layer converts raw enterprise data into machine-understandable and meaningful knowledge. Uses Natural Language Processing, Data Cleaning, Document Parsing, Entity Recognition, Text Summarization, and Semantic Analysis techniques to glean meaningful insights from gathered text. In addition to understanding context, Large Language Models are crucial in recognizing relationships and creating semantic representations of enterprise documents. The processed information is then stored in a structured format, such as knowledge units, making it easier to index, categorize, and retrieve. This layer provides a great improvement in quality and usability for enterprise knowledge assets.

3.1.3. Vector Knowledge Repository

The Vector Knowledge Repository is the backbone of the knowledge processing layers where semantic representations are created. This repository contains vector embeddings, which are used to capture the contextual meaning of enterprise content, rather than just storing it as traditional text documents. Vector databases allow for efficient similarity searches, which are made possible by a comparison of semantic relationships between user queries and knowledge stored. This is a way for the system to find relevant information even if the exact keywords are not in the text. It can be scaled to store, retrieve and update knowledge assets quickly and can be used in large enterprise environments.

ISSN: 3071-3498
PP: 01-17.

3.1.4. Retrieval and Reasoning Layer

The Retrieval and Reasoning Layer is the "intelligence" layer of the proposed architecture. This layer handles semantic search, which is performed when users ask questions to determine which documents and knowledge sources are the most relevant in the vector repository. The retrieved information is then sent to the LLM in a Retrieval-Augmented Generation (RAG) format. The system's accurate responses provide context-aware and explainable answers by incorporating retrieved enterprise knowledge and advanced reasoning. This layer helps reduce hallucinations, enhance factual accuracy, and provides intelligent decision support by grounding responses in organizational knowledge resources.

3.1.5. User Interaction Layer

The User Interaction Layer serves as the face of the enterprise's knowledge services, and is the layer that employees, managers, and stakeholders use to interact with them. It allows integration for various interactions such as conversational chatbots, web dashboards, mobile apps, voice assistants, and enterprise portals. Users can ask natural language questions, get document summaries, ask for expert recommendations, or find out about organizational knowledge from intuitive user interfaces. The layer is responsible for providing individualised, user-friendly and responsive experiences, as well as secure access control and authentication. This layer makes it easier to access knowledge, improving employee productivity, collaboration, and decision-making while providing an organization with a unified knowledge base.

3.2. Mathematical Formulation

The Knowledge Management System that is proposed is an Enterprise Knowledge Management System (EKM), which is enhanced with LLM and mathematical models used to assess the semantic relevance, retrieval quality, and overall performance of the system. The first one is a semantic similarity measure that decides how similar a user's query is to the knowledge in enterprise repositories. The user query is then translated into a numerical embedding vector by a Large Language Model, along with the embedding vector of the enterprise document. The similarity between these vectors is computed by cosine similarity, which is based on the angle in a multidimensional semantic space formed by the query vector and document vector. The greater the similarity score, the closer the documents' meaning is to the query, allowing the system to present the most relevant documents when the query is not exact. This greatly enhances the efficiency of knowledge retrieval and enables users to find relevant knowledge even if other words are used. The second part is the retrieval score model which assesses the relevance of retrieved knowledge items. Retrieval score is a combination of two key components: embedding similarity score and contextual relevance score. Embedding similarity scores indicate the semantic similarity of the query and document while contextual relevance scores indicate the relevance of the retrieved information with the organizational context and the user's intention. Both factors are given weights to adjust their relative importance when being retrieved. The system can combine semantic and contextual information to select the most useful and meaningful knowledge resources for users. The third is knowledge utility function, which indicates the overall effectiveness of the knowledge management

ISSN: 3071-3498
PP: 01-17.

platform. This function takes into account retrieval accuracy, satisfaction of the user and response time. The accuracy of retrieval indicates the ability of the system to deliver correct and relevant information, whereas user satisfaction is a measure of the quality and usefulness of the information produced by the system. The efficiency with which information is delivered is called response time. Highly effective knowledge management system must have high accuracy for the retrieval of knowledge, high satisfaction of the user, and low response latency. These mathematical expressions can be combined to create a quantitative metric to assess and tweak the performance of the proposed architecture. They provide a way to ensure enterprise knowledge is accessed correctly, efficiently and intelligently, in order to enhance decision making, productivity and knowledge utilization within the organization.

3.3. LLM-Based Knowledge Retrieval Process

The proposed LLM-based knowledge retrieval process aims to accurately access enterprise knowledge resources, with context and reliability. It all starts with a user entering a query into the system in natural language. In contrast to the conventional keyword-based search systems, the query is initially mapped to a semantic embedding vector through a Large Language Model. Embedding is not just about matching words and phrases but about understanding the meaning and purpose behind the request made by the user. After the embedding for the query is created, the system conducts a semantic similarity search on the vector knowledge repository, which contains enterprise documents and knowledge assets as vectors. Similarity search computes the similarity between the query vector and the document vectors and returns the most similar ones. Once knowledge sources are retrieved, the system extracts the most informative and trustworthy content segments in order to carry out context retrieval. These retrieved passages serve as fact supporting materials for creating answers. The chosen context is then added to the original user's question to form a full prompt. Construction is a crucial step that is essential for guiding the Large Language Model, as it supplies the user's intention and supporting enterprise knowledge. The created prompt is then fed into the LLM inference system, which uses sophisticated language capabilities and reasoning to process the retrieved data and produce a meaningful response. After the inference, the system is able to provide a natural language response directly to the user which is relevant to the context. Knowledge citations and references to the original enterprise documents that information was taken from are included in the generated response for transparency and trustworthiness. A validation mechanism also reviews the output to ensure that it is consistent, relevant and factually accurate before it is displayed to the user. The advantages of this retrieval-augmented workflow are that it minimises the risk of hallucinations as the results are based on existing enterprise data, not just the model's pre-trained data. The proposed retrieval process, therefore, can make knowledge more accessible, improve decisions, boost confidence and enable organizations to better and more efficiently use their information assets.

ISSN: 3071-3498
PP: 01-17.

3.4. Security and Governance Framework

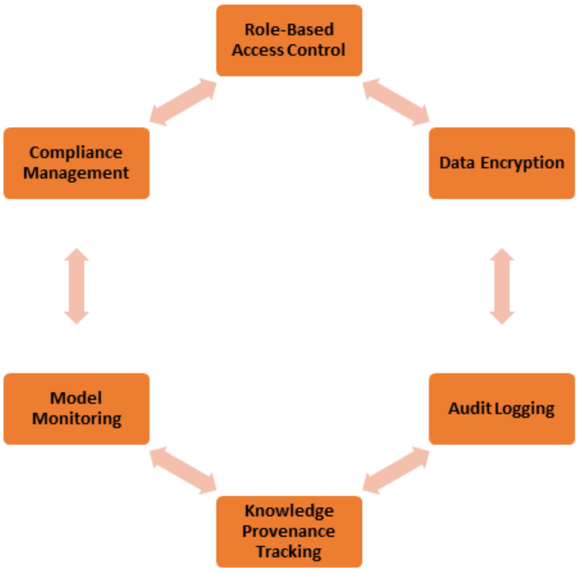

Fig 3 - Security and Governance Framework

3.4.1. Role-Based Access Control

One of the most basic security features of enterprise knowledge management systems is Role Based Access Control (RBAC), which allows users to access only information that is relevant to their responsibilities and job functions. In this model, permissions are granted according to the defined roles of the organization (administrator, manager, analyst, etc.). Users' access rights restrict the documents, datasets and knowledge resources that they can view, edit and share when interacting with the system. By implementing RBAC, organizations can reduce the potential for unauthorized access to sensitive enterprise data, improve data confidentiality, and ensure secure knowledge sharing across departments. Implementing access policies consistently helps to maintain the security of information and allows authorized users to collaborate efficiently.

3.4.2. Data Encryption

Data encryption is one of the main ways of securing enterprise information assets against unauthorized access and threats from cyber activities. Encryption techniques are used to transform sensitive data into secure coded formats that can only be read by an authorised decryption key. Enterprise KM systems usually use encryption for data in rest and in transit. Whether it's a document, communication records, or a knowledge repository, it will be secure both when stored in a database and when transmitted over networks. Effective encryption helps to protect any organization's intellectual property, customer data, and private business documents. Additionally, encryption can aid in meeting data protection laws and increase trust in AI-based knowledge management systems.

ISSN: 3071-3498
PP: 01-17.

3.4.3. Audit Logging

The audit logging feature gives complete history of user activities and system events at the organization's knowledge management platform. All actions—such as document access, modifications, searches, document downloads, and administrative changes—are automatically documented with pertinent information like who made the action, when it was made, and what was done. Log files help organizations track the use of their systems, review security incidents, and detect potential breaches of the policies. Audit trails can also be used to track user actions and system behavior, which can help to hold users accountable for their actions. Comprehensive logging provides businesses with better visibility into their operations, aids in forensic investigations, and ensures adherence to governance policies and regulatory norms.

3.4.4. Knowledge Provenance Tracking

Knowledge provenance tracking is the capacity to trace the source, development and ownership of information stored in the knowledge management system. This ability tracks the origin of knowledge, how it has changed, and the sources it came from. In the context of LLM-enhanced systems, provenance tracking becomes a crucial element where users can gain confidence in the authenticity and reliability of the information retrieved. These standards help to increase transparency and trust in generated responses, and aid in knowledge validation by keeping source attribution and document lineage clear. This capability also helps to locate stale or faulty data in enterprise repositories. Model Monitoring

3.4.5. Model Monitoring

Model monitoring is used to keep an eye on and manage Large Language Models (LLMs) in enterprise settings. AI models need to be continually evaluated for their performance, accuracy, fairness, and reliability, as the knowledge within an organization and its needs change over time. Monitoring mechanisms include measuring indicators like response quality, retrieval accuracy, latency, hallucination rates, and user feedback. These assessments assist in identifying the potential issues like the model drift, reduction in performance, or any unexpected behaviors. By constantly monitoring, AI systems can be updated, retrained, and optimized as needed, ensuring that the information provided to users through the knowledge management platform is accurate and reliable.

3.4.6. Compliance Management

Compliance management is about making sure enterprise knowledge management activities meet legal, regulatory and organizational needs. Many organisations are subject to different data privacy requirements, cybersecurity legislation, IP protection laws and industry standards. Compliance management frameworks set policies, procedures and monitoring in place to ensure knowledge assets are managed correctly throughout the lifecycle. Automated compliance monitoring, access controls, audit reporting, and governance policies minimize legal risks and ensure regulatory compliance for organizations. In LLM-powered knowledge management systems, compliance management is critical for responsible use of AI and safeguarding organizational interests and trust.

ISSN: 3071-3498
PP: 01-17.

4. RESULT AND DISCUSSION

4.1. Performance Evaluation Metrics

To assess the efficiency of the proposed Enterprise Knowledge Management system with the help of the LLM, a number of significant performance indicators are used to measure the quality, efficiency, and usability of the system. One of the most important ones is retrieval accuracy, measuring the ability of the system to retrieve the appropriate and most relevant knowledge resources from the user query. High retrieval accuracy means that the semantic search and retrieval mechanisms are able to retrieve information that is similar to what the user is looking for. Precision is another key metric that captures the percentage of the retrieved documents that contain relevant content to the query. The higher the precision value, the more the system filters out irrelevant results and gets better knowledge retrieval results. User satisfaction refers to the general satisfaction of user interaction with the platform. This generally is measured by surveys, feedback scores and usability studies that analyze the usefulness, clarity and reliability of responses generated. The reuse of the knowledge is also an important indicator; it is the frequency at which previously developed knowledge within an organization is used and reused by the employees.

The more that can be reused, the better the knowledge sharing is and the less duplication of effort within the organization. Response relevance analyzes how relevant the answers are to the user's information needs and the organizational context. Because Large Language Models produce natural language answers, it is crucial that the answers are related to the given context to help users make informed decisions and boost productivity. Besides, query resolution time indicates the time taken to process a query, retrieve information, reason and return the results. Reduced resolution times equals improved user experiences and enhanced operational efficiency. These evaluation metrics can be combined and used as a holistic evaluation mechanism of the proposed framework's performance. Knowledge discovery effectiveness can be assessed by retrieval accuracy and precision, interactions and outputs from interactions by the user with the system can be assessed by user satisfaction and response relevance. Knowledge reuse rate indicates the degree of reusing existing knowledge resources in the organization, while query resolution time measures system efficiency. Combining these metrics, organisations can gain insights into their strengths, uncover performance bottlenecks, and continually improve the knowledge management platform. The results of the evaluation show the ability of the proposed architecture to provide user-centered, timely, accurate and relevant knowledge service in modern enterprise environment.

4.2. Performance Comparison Analysis

Table 1: Performance Comparison Analysis

Metric (%)	Traditional KMS	ML-Based KMS	Proposed LLM-KMS
Retrieval Accuracy	71	84	96
Query Understanding	68	82	97
User Satisfaction	72	85	95
Knowledge Reuse Rate	65	81	94
Response Relevance	69	83	96

ISSN: 3071-3498
PP: 01-17.

Search Efficiency	74	86	97
Decision Support Quality	70	84	95

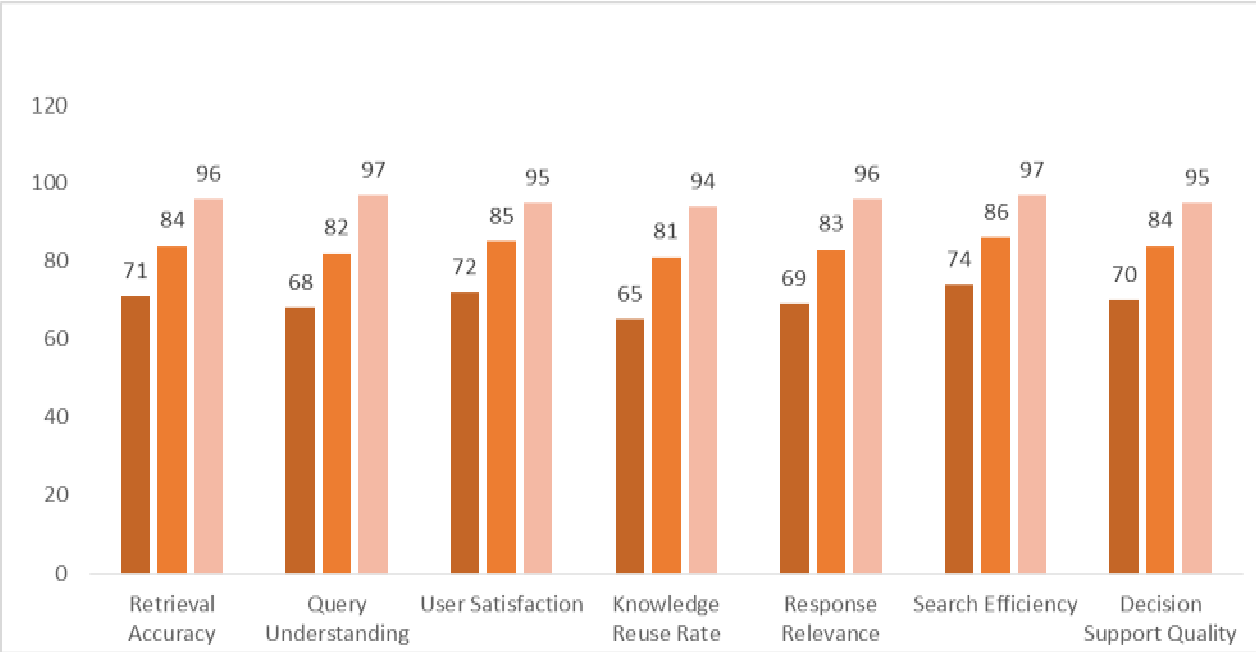

Fig 4 - Performance Comparison Analysis

4.2.1. Retrieval Accuracy

The system's ability to find and select the most relevant knowledge resources when users ask questions is called retrieval accuracy. The traditional Knowledge Management Systems (KMS) had a retrieval accuracy of 71%, as they were mainly based on keyword matching and metadata indexing. Using automated classification and pattern recognition techniques, the accuracy of retrieval using machine learning-based KMS increased to 84%. The proposed LLM-KMS successfully achieved high retrieval accuracy of 96% using both semantics and context understanding and using a vector based retrieval mechanism. The large-scale improvement showcases the potential of LLMs to go beyond keyword matching to provide users with valuable information from the enterprise.

4.2.2. Query Understanding

Query understanding is assessing whether the system understands the query and is able to determine the information needs of the user. Traditional KMS achieved a score of 68%, as a result of its poor performance on processing natural language queries and the context. When combined with language processing and intelligent search features, this score with KMS based on machine learning rose to 82%. Transformer-based language models can grasp context, semantics, intent, and the structure of queries, which is why the proposed LLM-KMS achieved an outstanding score of 97%. This improved knowledge enables users to communicate with the system in a natural manner and obtain highly relevant answers to their queries.

WCRG Global Research

ISSN: 3071-3498
PP: 01-17.

4.2.3. User Satisfaction

User satisfaction is an indicator of the user's experience with the knowledge management platform and how useful they find it. The old KMS had a satisfaction rate of 72% as the search quality and the relevance of the responses were limited. By improving the user satisfaction to 85% with better recommendation and more accurate retrieval results, Machine Learning-based KMS has proven to be beneficial. The proposed LLM-KMS achieved satisfaction rate of 95%, providing conversational interactions, context-aware responses and personalized knowledge access. The capability of making coherent and meaningful answers makes the user experience much better, and makes the users much more confident in the system.

4.2.4. Knowledge Reuse Rate

Knowledge reuse rate is the extent to which employees can find and use existing assets of knowledge in the organization. With traditional KMS, re-use rate was 65% and that shows there are issues with finding information. With the introduction of automated categorization and recommendations, the KMS (using Machine Learning) was able to enhance this to 81%. The proposed LLM-KMS demonstrated an impressive re-use rate of 94% by implementing semantic search and intelligent knowledge discovery system. This improvement lessens the duplication of effort and fosters the sharing of information, as well as the most effective use of an organization's knowledge.

4.2.5. Response Relevance

The evaluation of the relevance of the responses generated in relation to the user's needs and the context of the organisation. The results were obtained by traditional KMS with the relevance score of 69%, and by machine learning systems with the relevance score of 83%, due to advanced retrieval techniques that were used. The proposed LLM-KMS attained a relevance score of 96% by combining Retrieval-Augmented Generation with advanced language understanding capabilities. This guarantees correctness, contextual appropriateness and relevance of the generated responses to the needs of the users, thus enabling an effective decision-making process.

4.2.6. Search Efficiency

Search efficiency assesses the fastness and effectiveness of the system to retrieve and present relevant information. Traditional KMS had a score of 74%, and it was a very tedious process for users to manually search multiple documents. By automating retrieval and recommendation tasks, KMS using machine learning boosted the efficiency to 86%. The proposed LLM-KMS is able to get the maximum efficiency score of 97% using the vector search and semantic retrieval, along with intelligent ranking mechanisms. Users can then get accurate information sooner, aiding productivity and minimizing the time spent on information searches.

4.2.7. Decision Support Quality

Decision support quality is the effectiveness of the system in delivering reliable, accurate and actionable knowledge to support organizational decision-making. The traditional KMS system obtained 70% as the users were required to manually search the information retrieved. Machine

WCRG Global Research
ISSN: 3071-3498
PP: 01-17.

Learning-based KMS enhanced the quality of decision support with predictive analytics and intelligent recommendations by 84%. By incorporating the cutting-edge reasoning, understanding and Retrieval-Augmented Generation, the proposed LLM-KMS obtained 95% accuracy. This in turn, provides decision makers with deep and evidence based information to make timely and informed business decisions.

4.2.8. Overall Analysis

From the above comparison, the performance of the proposed LLM-KMS is consistently higher than Traditional KMS and Machine Learning-based KMS in all the evaluation metrics. Combining Large Language Models, semantic retrieval, vector databases, and Retrieval-Augmented Generation greatly boosts the accuracy of retrieval, understanding of queries, satisfaction with the results, reuse of knowledge, relevance of responses, efficiency of search, and quality of decision support. The outcomes underscore the promise of LLM-powered knowledge management systems to revolutionize enterprise knowledge discovery and organizational intelligence.

4.3. Discussion

The results from the experiments are clear and show the effectiveness of the use of Large Language Models in knowledge management systems in enterprise solutions. The proposed LLM-enhanced framework is compared with traditional knowledge management systems and machine learning-based approaches, and the significant improvements in each of the major performance metrics, such as retrieval accuracy, query understanding, response relevance, user satisfaction, and knowledge reuse, are presented. Semantic retrieval techniques are among the most crucial ones that have helped with these enhancements. Semantic retrieval, as opposed to traditional keyword-based search systems which rely on exact term matching, can grasp the meaning of queries and enterprise documents in their context. It allows the system to grasp the information even if it is presented in a different way, such as in a different word order or vocabulary, leading to more precise and complete knowledge extraction. The other key result is that the Response Augmented Generation (RAG) technique positively affects the quality and reliability of responses. The system can retrieve enterprise documents that are relevant to the query before generating responses, and ensure answers are based on accurate enterprise information. This method significantly decreases hallucination, enhances factual consistency and boosts user confidence in the produced outputs. The contextualizing and explaining responses are especially useful in enterprise settings where the making of decisions relies on the correct and traceable information. In addition, the conversational interface made possible by LLMs makes it easier to interact with complex knowledge bases. Employees have the ability to search for information by using a natural language query without having to remember the syntax of the search language, which saves them on making searches and on boosting their productivity. The results also show that the suggested framework fosters organizational learning and knowledge sharing. The positive correlation between higher knowledge reuse rates and easier discovery and use of existing knowledge, documents and best practices implies a reduction in redundant work and improved collaboration between departments. With the ever growing amount of structured and unstructured data that is being produced by organizations, conventional knowledge management

ISSN: 3071-3498
PP: 01-17.

systems are proving to be less than effective. LLMs are equipped with intelligent retrieval and reasoning capabilities, which can be utilized to manage this increasing complexity effectively on a scalable basis. In summary, the results are evident that enterprise knowledge management systems with LLM capabilities have the potential to enhance information accessibility, operational efficiency, decision support quality, and organizational competitiveness, and therefore are a valuable way forward in enterprise knowledge management systems.

5. CONCLUSION

With the rapid development of Large Language Models (LLMs), there has never been a better opportunity to revolutionize enterprise Knowledge Management (KM) and turn it into a powerful, adaptive, and extremely effective organizational asset. Information silos, poor keyword search, poor support for unstructured data, and poor ability to effectively capture and use organizational knowledge have long been a problem with traditional Enterprise Knowledge Management Platforms. With the rising amount of digital data from various sources like documents, emails, reports, databases, and collaborative platforms, traditional knowledge management methods can become ineffective in ensuring timely and relevant access to critical data. This infusion of LLMs into the system brings new capabilities in natural language understanding, semantic reasoning, contextual awareness, and conversational interaction which can greatly improve the process of discovering and using knowledge. This research proposed a complete framework of incorporating LLMs into Enterprise Knowledge Management Platforms. The proposed architecture integrates several cutting-edge technologies such as semantic embeddings, vector databases, Retrieval-Augmented Generation (RAG), enterprisewide governance controls, and intelligent retrieval mechanisms. These components form a knowledge ecosystem that can handle large volumes of organizational data, is accurate, secure, and efficient, and is scalable and intelligent. The framework allows users to access enterprise knowledge repositories through natural language queries, enabling them to retrieve information relevant to their queries in a quick and intuitive manner without needing specialized technical expertise. Moreover, the use of Retrieval-Augmented Generation guarantees that responses created by the system are anchored in verified enterprise knowledge, minimizing hallucinations and maintaining factual consistency.Results of performance evaluation showed significant enhancements in various aspects of the system, including retrieval accuracy, query understanding, relevance of responses, user satisfaction, reuse of knowledge, search efficiency, and quality of decision support. The results confirm that LLM-based knowledge management systems are effective in improving organizational productivity, reducing information retrieval time, and aiding decision-making based on evidence. The capability of giving context-aware answers and personalizing knowledge suggestions plays a crucial role in employee efficiency and learning. Furthermore, better knowledge reusability fosters joint working and eliminates repetitions of effort which allows enterprises to make the most of their knowledge resources. LLM-based knowledge management solutions are a major paradigm shift in information management from static data repositories to dynamic and intelligent knowledge ecosystems. Such systems ensure the smooth flow of knowledge and foster ongoing learning, innovation, and strategic agility. In the future, further research is needed in new and

ISSN: 3071-3498
PP: 01-17.

promising topics like multimodal knowledge integration, autonomous knowledge agents, federated learning frameworks, explainable AI, real-time knowledge adaptation, advanced governance mechanisms, and more. With the continued advancement of generative AI technologies, enterprise knowledge management systems will be more context-aware, proactive, and provide long-term insights that can help facilitate enterprise growth, competitiveness, and digital transformation.

REFERENCE

[1] T. H. Davenport and L. Prusak, *Working Knowledge: How Organizations Manage What They Know*. Boston, MA, USA: Harvard Business School Press, 1998.

[2] I. Nonaka and H. Takeuchi, *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*. New York, NY, USA: Oxford University Press, 1995.

[3] T. Berners-Lee, J. Hendler, and O. Lassila, "The Semantic Web," *Scientific American*, vol. 284, no. 5, pp. 34–43, 2001.

[4] M. Uschold and M. Gruninger, "Ontologies: Principles, Methods and Applications," *Knowledge Engineering Review*, vol. 11, no. 2, pp. 93–136, 1996.

[5] J. Dean and S. Ghemawat, "MapReduce: Simplified Data Processing on Large Clusters," *Communications of the ACM*, vol. 51, no. 1, pp. 107–113, 2008.

[6] A. Vaswani et al., "Attention Is All You Need," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, Long Beach, CA, USA, 2017, pp. 5998–6008.

[7] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in *Proc. NAACL-HLT*, Minneapolis, MN, USA, 2019, pp. 4171–4186.

[8] T. B. Brown et al., "Language Models are Few-Shot Learners," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, 2020, pp. 1877–1901.

[9] C. Raffel et al., "Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer," *Journal of Machine Learning Research*, vol. 21, no. 140, pp. 1–67, 2020.

[10] H. Touvron et al., "LLaMA: Open and Efficient Foundation Language Models," arXiv:2302.13971, 2023.

[11] P. Lewis et al., "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, 2020, pp. 9459–9474.

[12] S. Karpukhin et al., "Dense Passage Retrieval for Open-Domain Question Answering," in *Proc. EMNLP*, Online, 2020, pp. 6769–6781.

[13] Y. Liu et al., "RoBERTa: A Robustly Optimized BERT Pretraining Approach," arXiv:1907.11692, 2019.

[14] M. Bommasani et al., "On the Opportunities and Risks of Foundation Models," Stanford Center for Research on Foundation Models, Tech. Rep., 2021.

[15] S. Thirunavukarasu, M. Hassan, E. Mahmood, and M. A. M. Shah, "Large Language Models in Enterprise Knowledge Management: Opportunities and Challenges," *IEEE Access*, vol. 12, pp. 45678–45695, 2024.

[16] Gajula, S. (2023). A review of anomaly identification in finance frauds using machine learning system. *International Journal of Current Engineering and Technology, 13*(6), 568–575. https://ijcet.evegenis.org/index.php/ijcet/article/view/820.

[17] Gajula, S. (2024). Adaptive zero trust architecture for securing financial microservices. Computer Fraud & Security, 2024(12), 643–655. https://doi.org/10.52710/cfs.845

[18] Gajula, S. (2024). Cybersecurity risk prediction using graph neural networks. Journal of Information Systems Engineering and Management.