

A Smart Legal Assistant for Indian Laws

Saran Raj. S¹, Dr. C. Meenakshi²

¹MCA Student Department of Computer Application-PG VISTAS

²Assosiate Professor Department of Computer Application-PG VISTAS

DOI : <https://doi.org/10.51583/IJLTEMAS.2025.140400064>

Abstract: This project proposes a chatbot system that integrates Retrieval-Augmented Generation (RAG) and a Knowledge Graph (KG) to address queries on Indian law. The system dynamically adapts to the complexity of user queries to optimize performance and accuracy. For straightforward questions, the RAG module retrieves relevant legal documents and generates concise answers. For complex, multi-faceted queries, the system employs a multihop reasoning approach using the knowledge graph to derive accurate and context-aware responses. To enhance accessibility and inclusivity, the chatbot supports multiple regional languages, enabling users from diverse linguistic backgrounds to interact with the system in their preferred language. This hybrid design ensures efficient computation by selectively engaging the KG only when necessary, reducing resource usage while maintaining high accuracy. The chatbot provides users with seamless access to legal insights, offering quick answers for simple queries and in-depth reasoning for complex cases. By combining the strengths of RAG and KG, this system aims to revolutionize legal assistance by improving accessibility, precision, and usability for both legal professionals and the general public, empowering users to navigate Indian law with confidence and efficiency.

Keywords - The Smart Legal Advisor integrates Artificial Intelligence (AI), Retrieval-Augmented Generation (RAG), and Knowledge Graphs (KG) to provide efficient legal assistance. By leveraging natural language processing (NLP) and multi-hop reasoning, the system delivers context-aware responses to legal queries. It supports multiple regional languages, enhancing accessibility for diverse users. The system ensures accurate legal insights by dynamically adapting to query complexity, retrieving relevant legal documents, and offering structured interpretations of Indian law. Designed for law professionals and the general public, this AI-powered legal assistant improves legal research, promotes legal awareness, and facilitates real-time guidance in a resource-efficient manner.

I. Introduction

The rapid growth of legal data and the complexity of legal systems have created challenges in accessing precise and relevant legal information. To address these challenges, this project proposes a chatbot system that integrates Retrieval-Augmented Generation (RAG) and a Knowledge Graph (KG) to efficiently answer legal queries on Indian law. The system is designed to dynamically adapt to user query complexity, ensuring efficient and accurate responses. Simple queries are processed using the RAG module, which retrieves relevant

legal documents and generates concise answers. For complex, multi-faceted queries, the system utilizes a multihop reasoning approach powered by the KG, capturing intricate relationships within the legal framework. To enhance accessibility, the chatbot supports multiple regional languages, enabling users from diverse linguistic backgrounds to interact with the system in their preferred language. By combining the strengths of RAG and KG, the system offers an innovative solution that improves the accessibility, usability, and accuracy of legal assistance in India. This project aims to empower legal professionals, students, and the general public by providing reliable and efficient access to Indian legal knowledge, while optimizing computational resources and ensuring scalability.

II. Literature Survey

TITLE	AUTHOR, YEAR OF PUBLISH	LIMITATION
Retrieval-Augmented Generation for Large Language Models: A Survey	Yunfan Gao, Yun Xiong, Xinyu Gao 2024	Hallucination Problem – Even with retrieval, the model can generate incorrect or misleading responses.
Graph Neural Retrieval for Large Language Model Reasoning	Costas Mavromatis, George Karypis 2024	Many real-world knowledge graphs are missing entities and relationships, leading to sparsity issues.
Knowledge Graph Completion: A Review	Xin Zhao, Zongtao Duan 2023	Many KGC models struggle with large-scale knowledge graphs due to high computational costs

III. Methodology

The methodology for developing the Legal Query System using Retrieval-Augmented Generation (RAG) and Knowledge Graph (KG) involves multiple stages to ensure efficiency, accuracy, and scalability. The system follows a structured approach for data collection, processing, and response generation.

1. Data Collection and Preprocessing

- **Legal Document Acquisition:** Collect legal texts, case laws, statutes, and judgments from government databases, court websites, and legal repositories.
- **Data Cleaning:** Remove redundant, outdated, or irrelevant information to ensure high-quality inputs.
- **Multilingual Data Processing:** Translate and preprocess legal texts to support multiple regional languages using NLP techniques.

2. Knowledge Representation

- **Knowledge Graph Construction:**
 - Identify legal entities (e.g., laws, cases, judges, courts).
 - Establish relationships between entities for multi-hop reasoning.
 - Store structured data in a graph database (e.g., Neo4j).
- **Text Embedding for RAG:**
 - Convert legal documents into vector embeddings using Transformer-based models (e.g., BERT, LLaMA).
 - Store embeddings in a vector database (e.g., FAISS, Pinecone) for efficient retrieval.

3. Query Processing and Classification

- **Intent Recognition:** Classify user queries as simple (fact-based) or complex (multi-faceted).
- **Language Detection:** Identify the regional language of the query and translate if necessary.
- **Query Routing:**
 - **Simple Queries** → Processed using RAG, retrieving relevant legal documents and generating a concise response.
 - **Complex Queries** → Handled by Knowledge Graph, leveraging multi-hop reasoning for deeper insights.

4. Response Generation

- **RAG-based Answering:**
 - Retrieve relevant legal documents from the vector database.
 - Generate context-aware responses using LLMs (e.g., GPT, LLaMA).
- **Knowledge Graph Reasoning:**
 - Perform logical reasoning over legal entities.
 - Trace relationships and dependencies between laws and case precedents.
- **Hybrid Approach:** If a query requires both methods, the system integrates results dynamically.

5. User Interaction & Feedback Mechanism

- **Multilingual Interface:** Present answers in the user's preferred language.
- **Interactive Refinement:** Allow users to refine their queries for better responses.
- **Feedback Loop:** Improve system accuracy by integrating user feedback and retraining models.

6. System Deployment & Optimization

- **Web-Based and API Integration:** Deploy the system as a web application and provide an API for legal platforms.
- **Performance Optimization:** Monitor response time, accuracy, and resource utilization to improve efficiency.
- **Scalability and Updates:** Continuously update legal databases and refine KG/RAG models for evolving legal requirements.

Structure of the suggested work

The given architecture diagram represents a Retrieval-Augmented Generation (RAG) pipeline designed for answering legal queries using Large Language Models (LLMs) and a Vector Database (Vector DB). The process begins when a user submits a legal question. To provide an informed response, the system first processes legal documents, such as case laws and statutes, by splitting them into chunks for efficient retrieval. These chunks are then converted into vector embeddings using an embedding model and stored in a Vector Database (Vector DB) for fast and relevant retrieval.

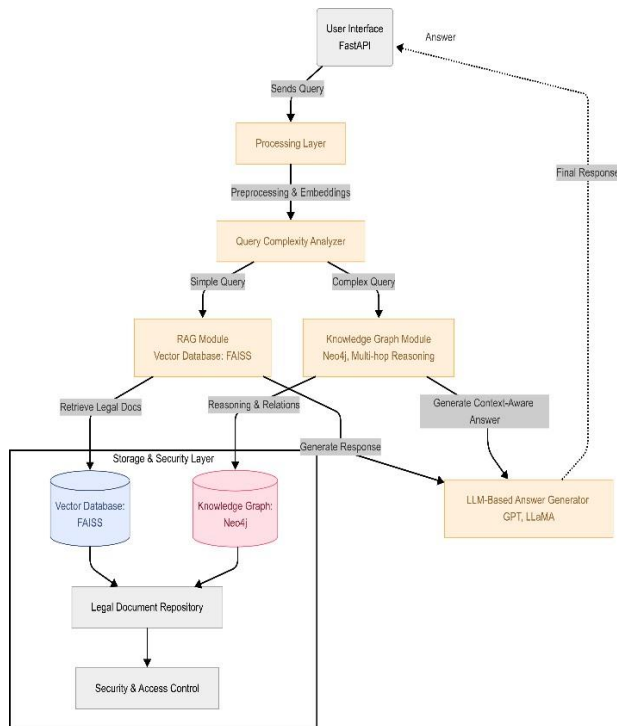


Fig.1 Architecture Diagram

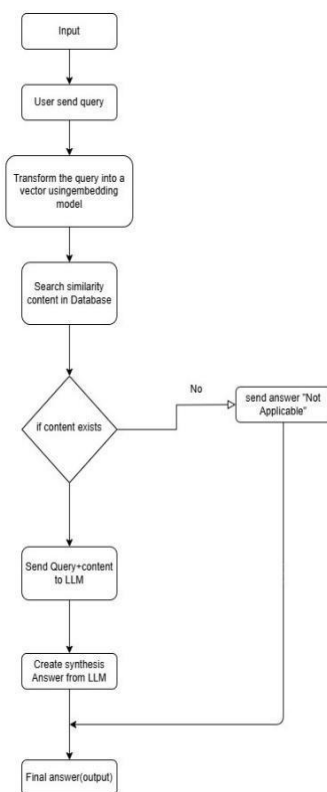
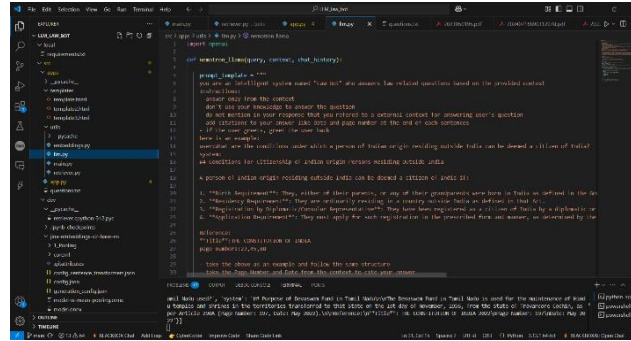
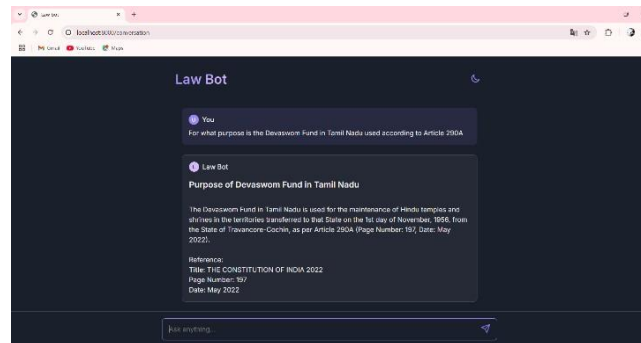


Fig.2 Data Flow Diagram

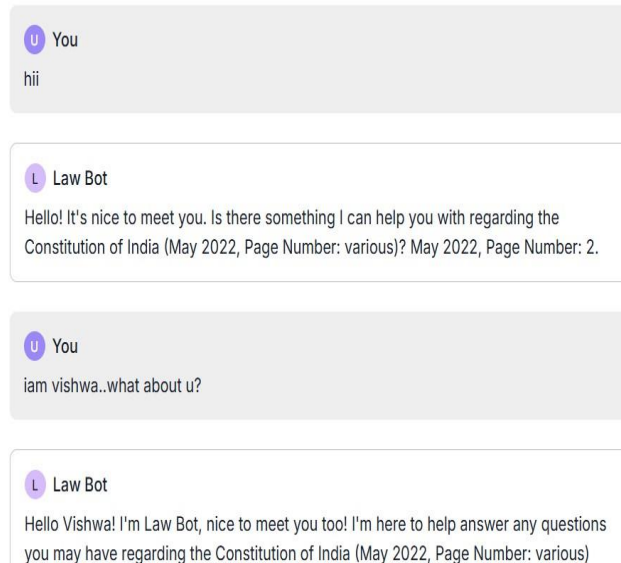
IV. Implementation



This diagram explains a modules and code implementations.



Law Bot



It represents a User Interface in our system to send a prompt and receive response quickly.

Existing And Proposed Technique

Existing Techniques

Current waste segregation methods include manual sorting, which is inefficient and hazardous. Smart dustbins use ultrasonic sensors and servo motors to automate lid opening, while magnetic separation extracts metallic waste. Infrared (IR) sensors and near-infrared spectroscopy (NIR) are used in industrial recycling to differentiate plastics and organic waste. IoT-based systems with GSM and Wi-Fi modules enable real-time waste level monitoring, and advanced models use AI and machine learning for waste classification. However, existing solutions are often costly, infrastructure-heavy, and lack hybrid power support. The

proposed system enhances efficiency with low-cost sensors, automated sorting, GSM-based monitoring, and a solar-battery hybrid power supply, making it a scalable and sustainable waste management solution.

Proposed Technique

Current systems designed to address legal queries predominantly use either Retrieval-Augmented Generation (RAG) or Knowledge Graph (KG) approaches. While RAG systems excel at retrieving relevant documents and generating answers for straightforward queries, they often struggle with complex queries requiring logical inference and multihop reasoning. Conversely, KG-based systems provide structured reasoning capabilities but can be computationally intensive and limited in scalability when handling vast amounts of unstructured data. This project aims to combine the strengths of both approaches, introducing a novel hybrid system that dynamically selects the optimal method based on query complexity, thus overcoming the limitations of existing systems.

V. Conclusion

The proposed legal query system, integrating Retrieval-Augmented Generation (RAG) and Knowledge Graph (KG), offers an innovative approach to providing accurate and context-aware legal insights. By leveraging RAG for simple queries and KG for complex, multi-hop reasoning, the system ensures efficiency, accuracy, and scalability. Additionally, its support for multiple regional languages enhances accessibility, making legal knowledge more inclusive for users across India. This dynamic framework not only optimizes resource usage but also empowers individuals—both legal professionals and the general public—to navigate Indian law with confidence and ease.

Future Scope

The proposed system combines Retrieval-Augmented Generation (RAG) and Knowledge Graph (KG) to efficiently answer legal queries on Indian law. Simple queries are handled by the RAG module, which retrieves relevant legal documents and generates concise responses. For complex, multi-faceted questions, the system utilizes a multihop reasoning approach powered by the KG, capturing intricate relationships within legal data. To enhance accessibility and inclusivity, the system supports queries in multiple regional languages, ensuring that users across India can interact with the platform in their preferred language. By dynamically switching between RAG and KG based on query complexity, the system ensures efficiency, accuracy, and scalability, offering users fast and reliable access to Indian legal knowledge while optimizing resource usage.

References

1. (2024) Yunfan Gao., Retrieval-Augmented Generation for Large Language Models.
2. (2024) Costas Mavromatis., Graph Neural Retrieval for Large Language Model Reasoning.
3. (2020) Zongtao Duan., Knowledge Graph Completion: A Review.
4. (2024) Abu-Salih, Bilal et al. A systematic literature review of knowledge graph construction and application in education.
5. (2024) Allan Almeida., LEGAL AI: A Tool to help Lawyers optimize their Performance.
6. (2024) Rakunatha Khrishanth; A Study on Judicial System and Implementation of an AI- Driven Judicial Process in India.
7. (2021) Youssra Amazou, Faouzi Tayalati. Accurate AI Assistance in Law Using RAG to Advance Legal Technology.
8. (2020) Joshua D; Osofsky, Leigh. Automated Legal Guidance.