

Legal Lens: An AI-Based Judicial Platform for India

Abhiyanshu Ruhela¹, Aman Pratap Solanki², Ayush Kaushik³, Hariom Giri⁴

Nitin Goyal⁵

^{1,2,3,4} UG Student, CSE Dept. , R.D. Engineering College Research Centre
Ghaziabad, India

⁵Assistant Professor , CSE Dept. , R.D. Engineering College Ghaziabad, India

ABSTRACT

Judicial systems continue to struggle with increasing case backlogs, inefficient access to legal records, and limited public engagement with legal resources. These challenges hinder timely justice delivery and reduce overall system transparency. This paper introduces *Legal Lens*, a unified framework aimed at improving the accessibility and usability of legal information. The proposed approach focuses on streamlining information retrieval and presenting legal content in a more user-friendly manner. The study outlines the system architecture, key functionalities, and design considerations with an emphasis on ethical aspects such as fairness and data protection. By enhancing access to legal knowledge, the framework seeks to support informed decision-making and

contribute toward a more efficient and inclusive legal ecosystem.

KEYWORDS

Intelligent Legal support Systems, Information Retrieval Supported Generation (RAG), FastAPI, FAISS, LangChain, Large Language Models, Legal Information Retrieval.

I. INTRODUCTION

The cases pending and unheard are already in abundance which ultimately delays the declaration of justice to the people of our country. The large amounts of legal data being unavailable or distant from the people results in unawareness and non-guided legal procedures.

This makes the complete process slow, inefficient, and costly.

Legal Lens, referred to as the “Justice Bridge”, acts as an intelligent platform that helps individuals understand legal issues and access relevant legal information without requiring extensive legal expertise. This particular system seeks to:

- Maximize the automation of the legal research through NLP(Natural Language Processing).
- Properly utilize the LLM(Large Language Models) fashions for understanding and detailed processing of the felony, against which the person seeks justice.
- Provide the person with the proper and all the possible solutions, data which can be initiated in a process resulting in justice.
- Provide the facilities to legal professionals and litigants in making informed decisions.
- Enhance easy and cost-effective methods to gain justice with interactive AI-driven platform.

II. LITERATURE SURVEY

[1] Over the past few years, technological advancements have increasingly influenced the legal domain, enabling automation in criminal studies for tasks like evaluating agreements and anticipating case outcomes.

[2] This section summarizes earlier work and notable contributions in the area of AI-supported legal resource systems, structured into the following parts.

- AI in legal studies
- AI enabled chatbots and assistants
- Regulations in documents handling and settlement parsing
- Predictive AI in regulation

[3] Use of NLP in legal studies have been set as an example in many legitimate ways. Various cases entities, their laws and case precedents are specifically mapped onto the graphs in order to

generate the common points and relations between them. Lack of specific domain results in worsening of prison interpretation and reasoning in AI fashions.

[4] Artificial Intelligence chatbots along with other digital assistants are significantly being used by the users.

[5] ‘Do not pay’ in 2015 by Browder: It gives suggestions depending upon the customer rights and small claiming disputes. Ross Intelligence in 2016 by Lewis: Its an AI powered assistant that gives cases regulation suggestions using the ways of IBM Watson.

[6] All the data possibly which can be collected is needed and helpful for study and best evaluation.

[7] FAISS is used for massive data retrievals. Indexing the legal embeddings: All of these vector embeddings are made from each legal file.

[8]The case regulations are thoroughly supported by the multi step reasoning in Llama3-8B-8192.An up to date dataset of the Indian Judicial context is required. Asynchronous API management results in high performance.

[9]Recent advancements in **Large Language Models (LLMs)** have further improved legal AI systems. Models such as GPT and LLaMA demonstrate strong reasoning capabilities and can generate context-aware responses when combined with legal datasets.

[10] The proposed **Legal Lens platform** builds upon these advancements by integrating FastAPI, LangChain, HuggingFace embeddings, FAISS, and the Llama3 model to provide an AI-driven legal assistance platform specifically designed for the Indian judicial ecosystem.

| Ref | Paper | Method Used | Limitation |
|-----|-------------------------------|----------------------------|-------------------------------------|
| [1] | Streamlining Criminal Justice | AI case backlog management | Focused only on case prioritization |
| [2] | LawPal | RAG + LLM | Limited dataset scope |
| [3] | Legal Word Embedding Dataset | NLP embeddings | Not focused on full legal assistant |

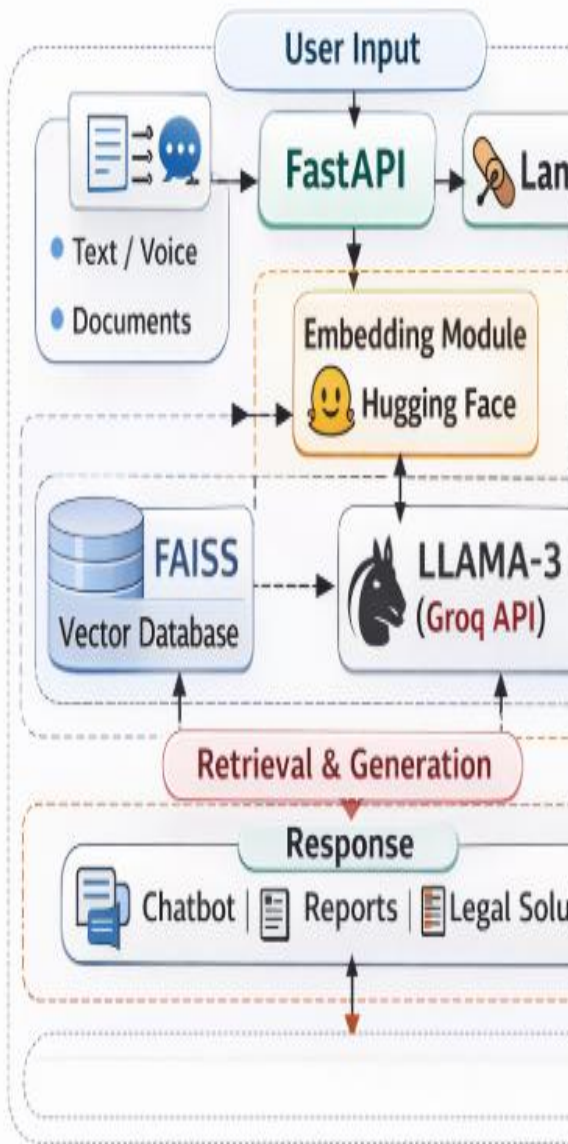
| Ref | Paper | Method Used | Limitation |
|-----|--------------------------------|--------------------------|---------------------------------|
| [4] | LLM as Tax Attorneys | LLM reasoning | Domain-specific legal tasks |
| [5] | Neural Embedding Pipeline | FAISS document retrieval | No legal reasoning |
| [6] | RAG LLM Systems | RAG architecture | General document systems |
| [7] | LawLLM | Domain specific LLM | Requires large training dataset |
| [8] | AI Legal Assistance Bangladesh | LLM legal assistant | Limited implementation scale |

Research Gap

| Problem | Existing Systems | Proposed Solution (Legal Lens) |
|-------------------------------|----------------------------|--------------------------------------|
| Legal knowledge accessibility | Limited public platforms | AI chatbot based legal assistance |
| Slow legal research | Manual case search | FAISS based semantic search |
| Lack of contextual reasoning | Traditional search systems | LLM-based legal reasoning |
| Complex legal language | Difficult for common users | AI generated simplified explanations |
| Limited Indian legal systems | Mostly foreign datasets | Focus on Indian legal context |

III. Architecture Diagram

Figure 1. System architecture of the proposed Legal Lens platform.



robust legal assistance platform when compared to other older projects.

FastAPI, LangChain, FAISS, Hugging Face Embeddings, Groq API, and Llama3-8B-8192 are the ones applied in this research. This section outlines how the system is developed, with a stepwise explanation of its individual components.

Refer to the points mention below:

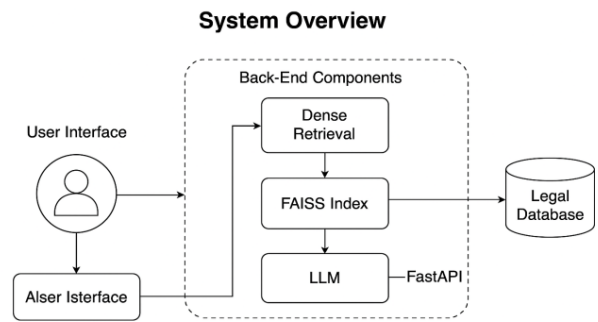


Figure 1. System Overview

IV. METHODOLOGY

By applying advanced language processing methods, the Legal Lens platform offers an efficient and scalable way to retrieve and analyze legal information. This system is designed to provide a scalable and

IV.1. System Architecture

- The five essential elements which make up the LegalLens machine architecture are as follows:
 - The User Interface: This Layer is responsible to respond to
 - consume inputs, which if specifically said, are the Questions raised by the users and then comes the answers provided by the system, known as the System's Answers.
 - API layer: This layer is responsible for fetching of and

- management of the consumer interface.
- The Embeddings are created from the criminal document in order to have a semantic and safe search. This is done by the presence and operative nature of the Hugging face +LangChain.
- Information retrieval is performed by using the FAISS Vector.
- The legal reasoning module is implemented using a large language model accessed through the Groq API, where the LLaMA3-8B-8192 model is employed to generate responses tailored to specific case inputs.

• **IV.2. Gathering and getting ready facts**

Legal databases and repositories serve as the primary sources of legal documents. Text Content cleaning: This includes removal or ignorance of the whitespaces, HTML tags and other special characters. The Lowercasing and Uppercasing of the letters is considered and normal and accepted. In this project, use of Hugging Face embeddings is mandatory for its unique and unseen efficiency and properties, to transform the processed text into numerical vectors. Also use of LangChain for embedding generation is important here.

A previously processed criminal document. The embeddings are created using the LangChain process. A numerical vector with 768 dimensions is the output.

IV.3. Use of FAISS for Indexing and looking into

A fast, scalable and efficient vector Database is created only by using the FAISS library.

- Index Production: The device creates an L2 flat index, just in order to shop the document embeddings.
- File insertion: The latest record of vector is held inside the index.
- Chronic Storage: The index is always stored on the disk in order to meet the needs in upcoming queries. The proper evaluation of the user queries, the presence of embeddings such that the most pertinent instances are found for this evaluation with the listed vectors by using the L2 distance metric.

IV.4. LLM based prison reasoning(Llama3-8B-8192)

The Groq API is further integrated with the very particular Large Language Model Llama3-8B-8192 for providing a better version of stylish prison reasoning.

- Input: FAISS associated instances and consumer queries.
- Spark off Creation: The system creates and intensifying prompt by fusion of query and legal pertinent contexts.

- Inference: Groq’s Llama3-8B-8192 is used to create reactions.

IV.5. Workflow present day the device

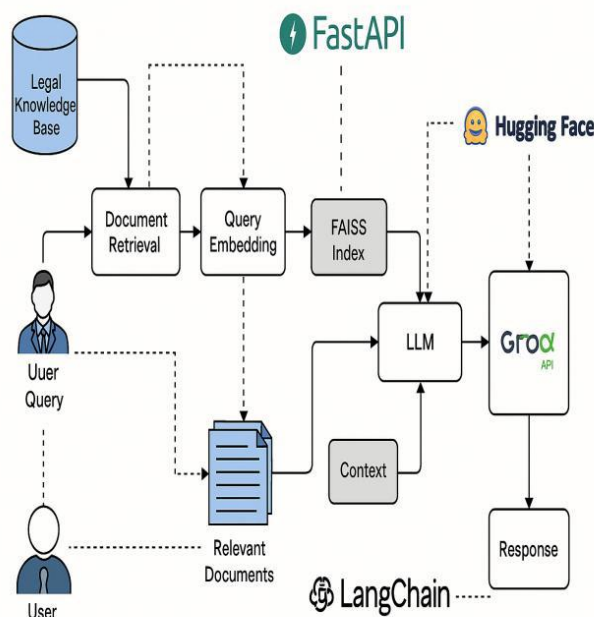
- User Input: Users publish the criminal queries but using the FastAPI interface.
- Embedding Generations: LangChain converts the question into embeddings.
- Contextual Response: The LLM is the one which receives the customer queries and instances for in depth justification.
- Output technologies: Formatted reaction is received by the user.

IV.6. Verification cutting edge the machine

- Dataset of atleast 1000 court cases have been

worked upon by the specialists in order

- to check and maintain the systems effectiveness. As per the findings show, an average 92% fee is fixed for finding a pertinent courtroom instances.
- Each and every query has a mean reaction time of 1-2 seconds.
- Effective Quality Feedbacks at the version from prison expertise end.



IV.6. Algorithm Used : -

Input: User legal query
Output: AI-generated legal guidance

Step 1: Receive user query

- Step 2: Generate embeddings using HuggingFace
- Step 3: Retrieve relevant documents using FAISS
- Step 4: Construct prompt using retrieved context
- Step 5: Send prompt to Llama3 via Groq API
- Step 6: Generate legal response
- Step 7: Return formatted answer to user

```

Step 1: User submits legal query
Step 2: Perform NLP preprocessing
- Tokenization
- Stopword removal
- Query normalization

Step 3: Generate semantic embedding
using Sentence Transformer model

Step 4: Search vector database
using FAISS similarity search

Step 5: Retrieve relevant legal documents
and statutes

Step 6: Apply Retrieval Augmented
Generation (RAG)

Step 7: Send context to Large Language Model

Step 8: LLM generates final legal response

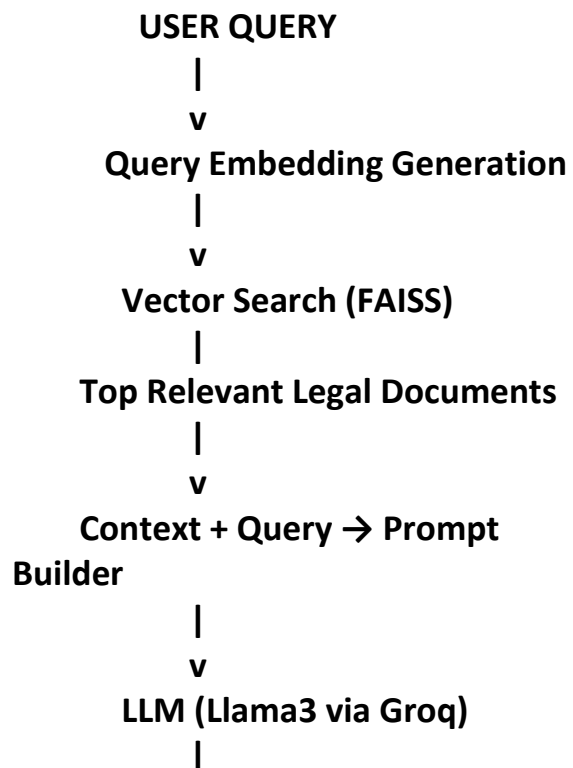
Step 9: Return answer to user interface
    
```

Dataset Description

The dataset applied in this study includes a range of Indian legal sources. Court judgments were gathered from publicly accessible repositories, including records from the High Courts and the Supreme Court of India. In addition, legal acts and statutory materials were sourced from official government portals.

A total of roughly 1450 legal documents, such as judgments, statutes, and articles, form the dataset. Alongside this, around 500 queries generated by users were included to measure the effectiveness of the Legal Lens system.

All documents were preprocessed using text cleaning techniques such as removal of HTML tags, normalization, tokenization, and stopword filtering before generating embeddings using HuggingFace sentence transformers.



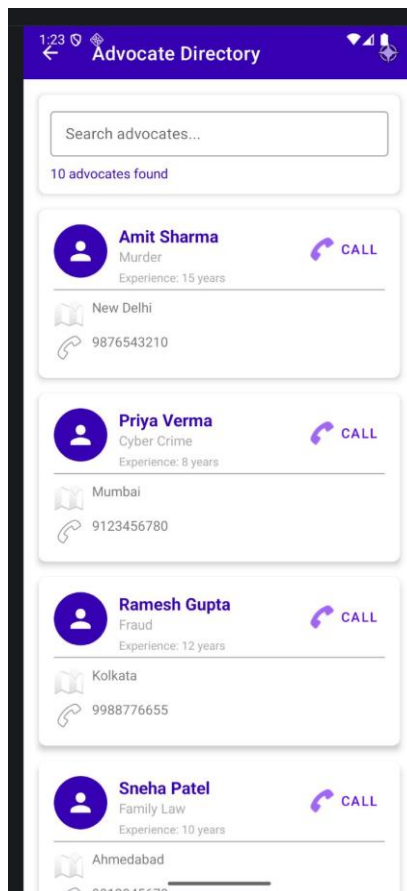
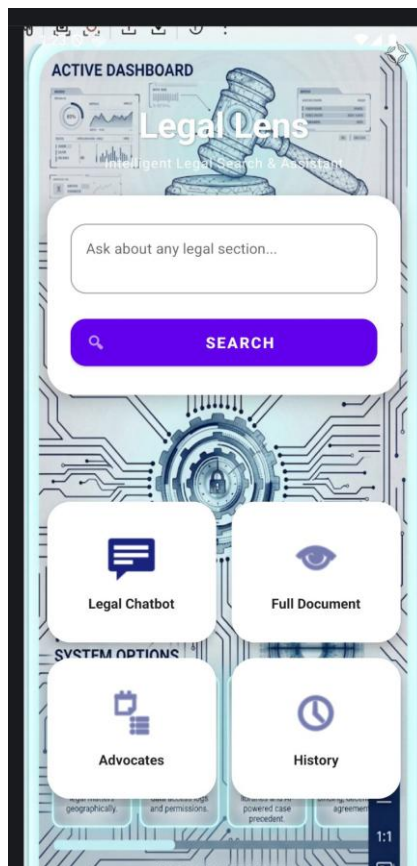
V
AI Generated Legal Answer

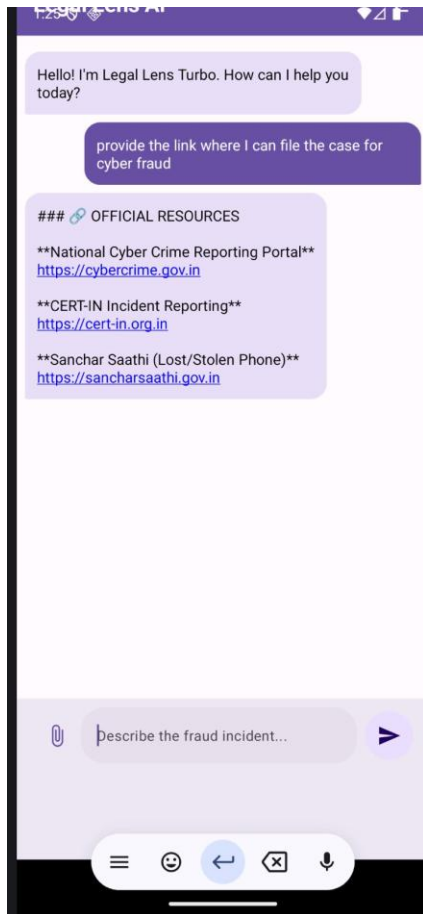
V. RESULT

A huge variety of the modern days legal documents, legal guidelines of cases and the queries generated by the users have been used to access LegalLens’s overall performance.

The consequences have shown how nicely the machine has been working in order to provide the unique legal records, processing of the questions and very reliable case retrievals. All the principal areas have been evaluated surely and accurately.

The following images show the detailed results of our project successfully replying to every prompt in detail, along with the sections that are applicable after complete analysis of the input prompt.





VI. Conclusion

This paper presented **Legal Lens**, an AI-driven legal assistance platform designed to improve accessibility to legal information within the Indian judicial ecosystem. The system integrates modern technologies including FastAPI, LangChain, FAISS vector search, HuggingFace embeddings, and the Llama3 large language model through the Groq API.

Findings from the experiments suggest that the system performs reliably in identifying relevant legal documents and generating explanations aligned with the given context.

By utilizing a Retrieval-Augmented Generation architecture, Legal Lens improves both the efficiency and reliability of legal information retrieval compared to traditional keyword-based and NLP-based systems.

The results indicate that AI-based legal assistance platforms can significantly support individuals in understanding legal processes and accessing relevant case information.

Future work will focus on expanding the dataset, incorporating multilingual legal support for Indian languages, and enhancing explainability to improve transparency and trust in AI-assisted legal systems.

VII. References

- [1] R. Marcotte, "AI-Assisted Legal Research," *Legal Information Management*, Cambridge Univ. Press, 2025.
- [2] S. Gupta, R. Pataskar, R. Dhattrak, R. Fernandes, and N. Narkar, "Smart AI Legal Advisor," *International Journal for Research in Applied*

Science and Engineering Technology (IJRASET), 2025.

[3] J. Gupta, A. Sharma, S. Singhania, and A. I. Abidi, "Legal Assist AI: Leveraging Transformer-Based Model for Effective Legal Assistance," *arXiv preprint*, arXiv:2505.22003, 2025.

[4] Z. Hou, Z. Ye, N. Zeng, and T. Hao, "Large Language Models Meet Legal Artificial Intelligence: A Survey," *arXiv preprint*, 2025.

[5] D. Panchal, A. Gole, V. Narute, and R. Joshi, "LawPal: A retrieval-based generation system for improved legal accessibility in India," *arXiv preprint*, doi:10.48550/arXiv.2502.16573, 2025.

[6] D. K. V. Dev, "Streamlining criminal justice: using intelligent methods for effective case backlog management," *International Journal of Science and Research (IJSR)*, vol. 13, no. 7, pp. 1557–1561, 2024.

[7] S. Sharma, R. Gupta, and P. Singh, "End-to-End Neural Embedding Pipeline for Large-Scale PDF Document Retrieval Using Distributed FAISS and Sentence Transformer Models," *ResearchGate*, 2024.

[8] A. Mishra, R. Patel, and S. Kumar, "Developing Retrieval Augmented Generation (RAG) Based LLM Systems Using PDF Documents as Primary Data Source," *arXiv preprint*, arXiv:2410.15944, 2024.

[9] Z. Zhong, J. Yang, and H. Chen, "LawLLM: A language-based model designed for the U.S. legal domain," *arXiv preprint*, arXiv:2407.21065, 2024.

[10] M. Rahman, T. Ahmed, and S. Islam, "Exploring Possibilities of AI-Powered Legal Assistance in Bangladesh through Large Language Modeling," *arXiv preprint*, arXiv:2410.17210, 2024.

[11] Meta AI Research Team, "The Llama 3 Herd of Models," *arXiv preprint*, arXiv:2407.21783, 2024.

[12] A. Bommarito and D. Katz, "Large Language Models as Tax Attorneys: A Case Study in Legal Capabilities Emergence," *arXiv preprint*, arXiv:2306.07075, 2023.

[13] H. Li, W. Su, C. Wang, Y. Wu, Q. Ai, and Y. Liu, "Enhancing legal case retrieval using structure-aware language models," *arXiv preprint*, arXiv:2305.06812, 2023.

[14] Y. Zhang, X. Wang, and J. Liu, "A dataset for evaluating legal word representations: a case study based on Chinese legal texts," *arXiv preprint*, arXiv:2203.15173, 2022.