International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 02 February 2026
Page No: 747-755
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0107

LegalBOT: A RAG and FAISS Hybrid Framework for Legal Case Intelligence

Dr. N. Danapaquiame[1], S. Shivani[2], B. Vasundhara[3], A. L. Virginia[4]

[1]Professor – Computer Science and Engineering, Sri Manakula Vinayagar Engineering College, Madagadipet, Puducherry

[2,3,4]UG - Computer Science and Engineering, Sri Manakula Vinayagar Engineering College, Madagadipet, Puducherry

Emails: *shivani.fr24@gmail.com[2], vasubala3028@gmail.com[3], vvirginia852@gmail.com[4]*

Abstract

The Legal Assistant Bot is an AI-based system for automating simple legal services leveraging Artificial Intelligence and Natural Language Processing. Legal work is heavily document-based, involves a lot of legal research, and relies on the interpretation of a legal professional, making it a laborious and expensive process that is not readily available to non-legal persons. The growing amount of legal information available and the rising demand for rapid legal assistance indicates the need for intelligent automation of legal professionals. This proposed system addresses this issue by providing an intelligent bot capable of processing legal documents, understanding natural language questions, providing answers to legal questions, tracking legal cases and even recommending a lawyer to the user based on their specific legal needs. The bot uses a combination of Machine Learning techniques and Retrieval Augmented Generation to provide accurate, relevant, and contextual answers in real time. Unlike other systems for legal information, the Legal Assistant Bot provides a unique interface the user interacts with. The system reduces the workload of legal professionals, minimizes the wait time between legal consultations, and provides legible and well-structured legal information that keeps the end-user informed and empowered. The proposed system improves the efficiency, availability, and affordability of legal assistance and services and demonstrates the potential impacts of AI-based systems on the legal environment.

Keywords: *Legal Assistant Bot; Artificial Intelligence; Natural Language Processing; Machine Learning; Retrieval Augmented Generation; Legal Automation; Chatbot*

1. Introduction

Legal contracts, agreements, case files are voluminous, complex documents that are difficult to extract insights from manually. Legal experts take a long time to evaluate the documents, identify the critical terms, monitor case information, and look for precedents. It is a manual task that is very time-consuming and inefficient, especially when the legal data is unstructured and the volume is large

Current methodologies either involve a considerable amount of manual processing or are based on rule-based and generic AI solutions that are unintelligent in the targeted domain of law. This results in a situation where the extraction of important clauses, delivery of contextually relevant responses in law, and maintenance of case workflows are all done ineffectively. Generic models of AI are further inefficient in a legal setup as a result of hallucinations and improper interpretation of legal expressions. Current approaches include manually processing documents, rule-based NLP approaches, and practice management software such as Clio, MyCase, and Lawmatics. Although these tools help in managing documents, client intake, and scheduling, their level of automation is low, semantic understanding is poor, and there is a lack of integration of document analysis and query analysis. The proposed work introduces LEGALBOT, an AI legal assistant incorporating Natural Language Processing, Large Language Models, Retrieval-Augmented Generation, and vector embeddings. LEGALBOT is capable of automatic analysis of legal documents, extracting important clauses, summarizing them, responding to legal queries based on contexts, and organizing case data and deadlines in one system. The system described can be classified as belonging to the area of Artificial Intelligence and Machine Learning (Legal Tech). LEGALBOT has been developed using an advanced technology stack with an AI focus. The coding languages used in this system are Python and others. Finally, Natural Language Processing (NLP) has been carried out using the spaCy and Transformer models (Legal BERT / LLMs) for semantic reasoning and clause detection. The system also employs the Tesseract OCR technique for text retrieval from the image done for legal documents. The system employs the RAG technique and vector embeddings for proper and informed legal answers and semantics stored in FAISS / Pinecone. The system employs the Firebase system for the management of the database and storage in the cloud. This paper is organized as follows. Section 2 provides a review of existing literature. Section 3 introduces our system architecture. Section 4 provides more implementation details. Section 5 provides results and analysis. Section 6 closes this paper with a conclusion.

2. Related Works

This paper provides an in-depth analysis of the role of LLMs in legal reasoning, focusing on text comprehension, argumentation, and decision support. It illustrates the ways in which models like GPT and LLaMA support lawyers in their legal research, summarization of documents, and drafting of documents. In this paper, the authors again propose a hybrid framework with the combination of LLMs with symbolic reasoning to obtain better consistency and explainability. Drawback: LLMs still operate with reasoning inconsistencies, bias, and hallucinations, making them too risky in practical legal applications [1]. The paper showcases an AI Legal Documentation Assistant that is capable of automating the drafting, review, and analysis of legal documents. This enhances efficiency, accuracy, and access for users. A limitation in this regard is that it can never replace expert legal advice when it comes to complex cases [2]. This paper tries to enhance summarization of the judgment in a legal case using judicial syllogism based on LLMs. In the results, it is seen that the performance is better than the T5 model, BART, and the PEGASUS model in identifying the logic relating to a law case. Shortcoming: This method is based on structured annotation and is quite hard to apply in different jurisprudential systems [3]. This paper introduces the Cl2Sum framework for abstractive summarization to mitigate hallucinations in LLM models. The approach applies contrastive prompt learning to distinguish high-quality and low-quality summaries. The results show improvement over BART, T5, and PEGASUS models. Shortcoming: The proposed learning approach is complicated by the requirement to design negative examples properly [4]. In this paper, we present LSDK-LegalSum, a summarization framework based on logical structure and knowledge of the domain. It relies heavily on the organization of legal documents to identify crucial elements of fact, reasoning, and conclusion. Experiments clearly indicate enhanced summary quality and fact accuracy than abstractive models. Shortcoming: It relies excessively on knowledge from specific domains and structured formats for law, thus being less adaptable to other domains [5]. The paper revisits retrieving legal precedents using advanced summarization and LLM-based frameworks. It goes on to proffer a legal-specific decompositional framework that breaks judgments into facts, issues, reasoning, and decisions. Results indicate improvement in accuracy and structure when juxtaposed with models like BART and PEGASUS. Drawback: It is an approach based on domain-specific decomposition; the approach complicates and reduces generalization [6]. The paper describes artificial intelligence and machine learning-based automation in processing and summarizing legal documents. This includes deep learning and various NLP models that are useful for entity extraction, document classification, and answering legal questions. Preliminary results reveal a manifold improvement in efficiency and accuracy from what is achievable through pure legal analysis. Drawback: The system depends on well-curated legal datasets and may work poorly for cross-jurisdiction or multilingual documents [7]. In this paper, the author has provided a full-fledged review of different text summary extraction techniques that have been used in Natural Language Processing. The text mentions abstractive, hybrid, as well as other text summary extraction techniques, ranging from rule-based systems all the way up to Deep Learning models. Criticism: The proposed study has a theoretical base and lacks an empirical validation of the proposed applications in the field of law [8]. This article presents an overall discussion of text summarization models in the field of Natural Language Processing. These models are further divided into extractive models, abstractive models, and hybrid models. Criticism: This research is purely theoretical and lacks validation in terms of its applications within the legal sphere [9]. From existing research the problem of hallucinations, inconsistency in reasoning, and the inability to validate the domain has been found in the existing LLM-based legal systems. The proposed system will improve the accuracy of legal responses by the use of OCR, NLP-based extraction of clauses, the summary process in the LLM, and the RAG technique using vector embedding.

3. Proposed Solution

The LegalBot framework aims at the analysis of complicated legal text by uniting non-homogeneous data modalities - scanned contracts, court case decisions, laws and statutes, and natural language questions. The LegalBot approach goes beyond traditional rule-based or independent LLM methods and uses a combination of Legal-BERT embeddings, the retrieval-augmented generation (RAG), and multi-stage reasoning (deductive rules, hypothesis building, precedent verification) for the systematic and coherent legal document assessment. The proposed framework consists of cascading functional levels which are explained further. The solution proposed in **Figure 1** is a structured five-layer architecture where five fundamental elements characterize each layer: (1) Input Processing (OCR via Tesseract + spaCy NER), (2) Hierarchical Embedding (Legal-BERT > FAISS indexing), (3) Semantic Retrieval, (4) Unified Reasoning Engine and Structured Output Generation. The main purpose of the architecture is to provide end-to-end dissemination and automation of the document analysis, query addressing and case management processes.

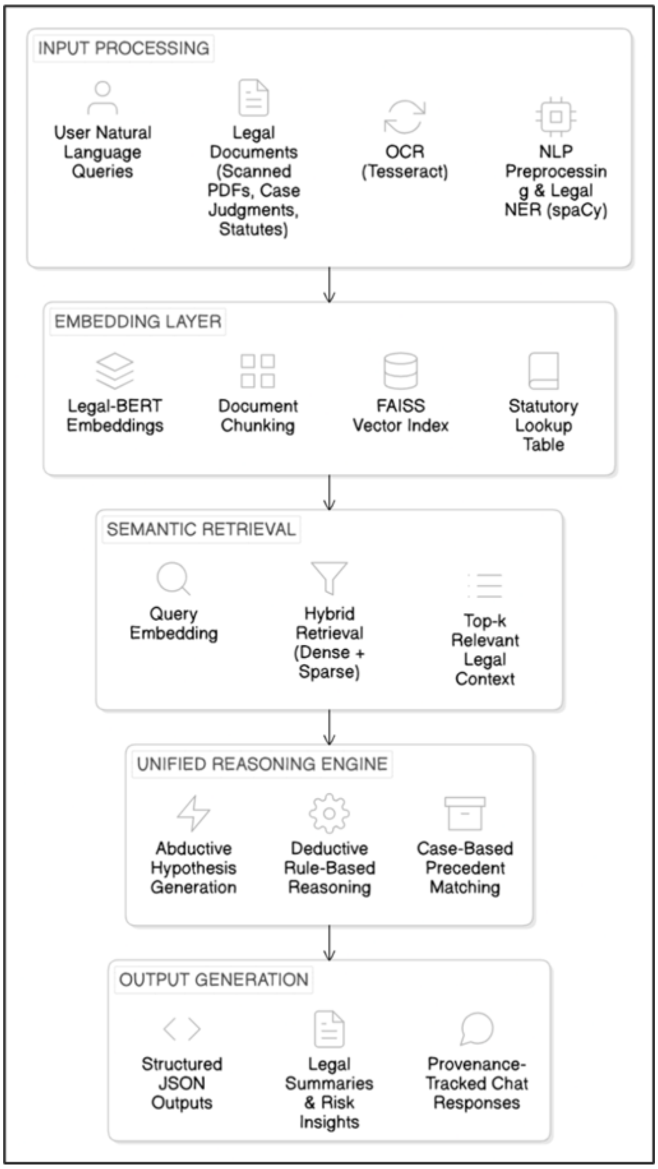
INPUT PROCESSING
User Natural Language Queries
Legal Documents (Scanned PDFs, Case Judgments, Statutes)
OCR (Tesseract)
NLP Preprocessing & Legal NER (spaCy)
EMBEDDING LAYER
Legal-BERT Embeddings
Document Chunking
FAISS Vector Index
Statutory Lookup Table
SEMANTIC RETRIEVAL
Query Embedding
Hybrid Retrieval (Dense + Sparse)
Top-k Relevant Legal Context
UNIFIED REASONING ENGINE
Abductive Hypothesis Generation
Deductive Rule-Based Reasoning
Case-Based Precedent Matching
OUTPUT GENERATION
Structured JSON Outputs
Legal Summaries & Risk Insights
Provenance-Tracked Chat Responses

Figure 1 Architecture of the Proposed Solution

3.1. Input Processing Layer

This layer collects various multi-modal legal documents. These documents include scanned PDF documents of legal contracts/cases, legislative/statutory materials, user case queries in natural language including case history. It extracts legal information from the input. For text extraction of scanned PDFs, it uses OCR software (Tesseract). For preprocessing, it uses various NLP techniques including tokenization, lemmatization, sentence splitting, legal entity recognition (statutes, parties, dates, clauses) via spaCy. This layer provides sufficient coverage of document structure and semantic elements for further analysis. Preprocessed text is cleaned and normalized to cater content discrepancies due to different formats in jurisdictional legal practices. (Tesseract). For preprocessing, it uses various NLP techniques including tokenization, lemmatization, sentence splitting, legal entity recognition (statutes, parties, dates, clauses) via spaCy. This layer provides sufficient coverage of document structure and semantic elements for further analysis. Preprocessed text is cleaned and normalized to cater content discrepancies due to different formats in jurisdictional legal practices.

3.2. Embedding Layer

This layer is the knowledge representation layer that converts the cleaned legal text into mathematically search-able vector embeddings that preserve semantic meaning for later retrieval and reasoning. The structured legal text (now structured legal document) allow the system to look up legal concepts in search of legal term rather than keyword search

3.2.1. Document Chunking

Document is chunked into 512 token overlapping chunks (20% overlap) ensuring phrases end at clause boundaries for contextual consistency during retrieval.

3.2.2. Legal-BERT Encoding

Converts text into 768-dimensional dense vectors embeddings with domain-specific Legal-BERT, retaining semantic meaning between legal concepts, terms, and jurisdictions.

3.2.3. Multi-Resolution Indexing

Embedding Index FAISS HNSW with M=32 and efConstruction=200 for document-level embedding level for broad search and clause level embedding level for precise search results even in large-scale legal databases.

3.2.4. Statutory Indexing

Separate deterministic lookup table for codified legal texts for exact rule retrieval in parallel to semantic vector search.

3.3. Semantic Retrieval Layer

User queries are converted into Legal-BERT embeddings to match against the FAISS HNSW index using cosine similarity (threshold ≥0.85) and top-5 document chunks are retrieved via approximate nearest neighbour search (suitable for large-scale legal document corpora) in dense hybrid retrieval that combines BM25 sparse retrieval (for exact statutory citation, proper nouns) with dense retrieval (vector matching) to capture semantic similarity in varied legal terminology; Legal ontology query expansion translates informal query terms into formal legal concepts; retrieved chunks are concatenated with provenance info (document id, clause position, jurisdiction, retrieval score) to form 4096 token context windows (per query) that dynamically prioritize recent superior court rulings for jurisdiction-appropriate reasoning inputs

3.4. Unified Reasoning Engine

This layer applies sophisticated legal reasoning patterns over the retrieved document context, like a structured reasoning pattern lawyer's use vs. statistical text generation.

3.4.1. Query Embedding Generation

Converts the retrieved context and user query into a structured prompt that contains the facts, statutes, and reasoning directives for a deterministic legal LLM call for each reasoning type.

3.4.2. Rule-Based Deductive Reasoning

Uses formal syllogism to map statutes to extracted facts (Major Premise: legal rule + Minor Premise: facts of the case ⊢ Conclusion: outcome of the case) with codified Indian statutes.

3.4.3. Abductive Hypothesis Generation

Ranks hypotheses for lack of evidence in a case with posterior score $P(H|E) \propto P(E|H) \cdot P(H)$ based on Bayesian approach to find candidate interpretations of law.

3.4.4. Case-Based Precedent Matching

Uses analogy-based case matching of retrieved cases with the same facts via weighted distance of embedding and structure matching, predicts the case outcome based on the hierarchy of cases.

3.5. Output Generation Layer

This final layer translates the reasoning engine's findings into usable outputs for legal practitioners, processing documents from designated folders and returning query responses via the chatbot. Document processing returns structured JSON with short summaries, identified legal risks with proposed mitigations, and compliance considerations, all organized in the database for the practitioner's review. Chatbot queries return on-the-spot conversational responses with reasoning trails that connect the deductions made back to the documents used and the legal reasoning applied, while all outputs are fully provenance-tracked, allowing legal practitioners to trace back through the analysis from outputs to input files for a comprehensive view of the AI's enhanced analysis of the document.

4. Algorithm

The proposed framework follows the algorithm in **Figure 2**.

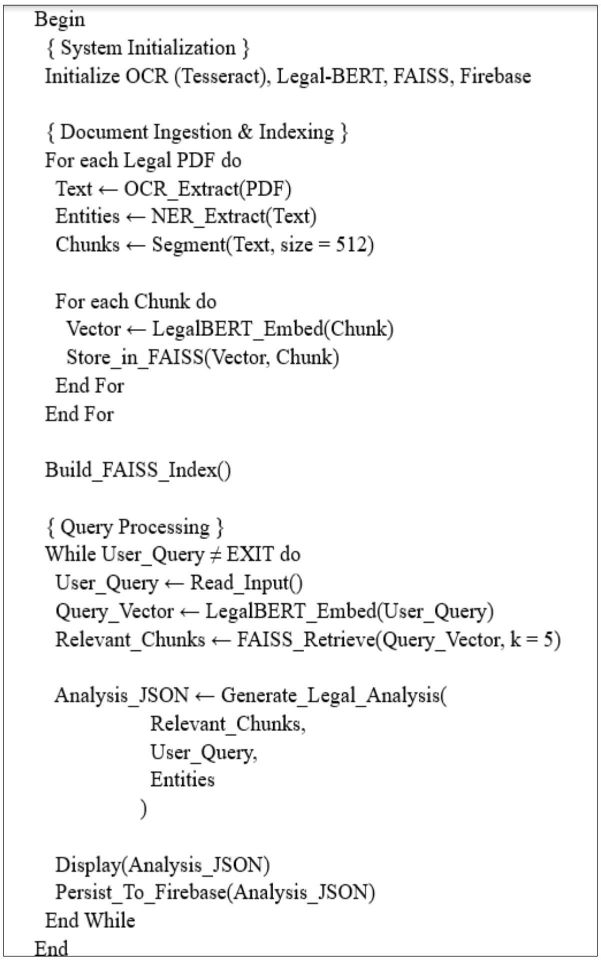

```
Begin
 { System Initialization }
 Initialize OCR (Tesseract), Legal-BERT, FAISS, Firebase

 { Document Ingestion & Indexing }
 For each Legal PDF do
  Text ← OCR_Extract(PDF)
  Entities ← NER_Extract(Text)
  Chunks ← Segment(Text, size = 512)

  For each Chunk do
   Vector ← LegalBERT_Embed(Chunk)
   Store_in_FAISS(Vector, Chunk)
  End For
 End For

 Build_FAISS_Index()

 { Query Processing }
 While User_Query ≠ EXIT do
  User_Query ← Read_Input()
  Query_Vector ← LegalBERT_Embed(User_Query)
  Relevant_Chunks ← FAISS_Retrieve(Query_Vector, k = 5)

  Analysis_JSON ← Generate_Legal_Analysis(
            Relevant_Chunks,
            User_Query,
            Entities
           )

  Display(Analysis_JSON)
  Persist_To_Firebase(Analysis_JSON)
 End While
End
```

Figure 2 Workflow of the AI-Based Legal Document Processing, Indexing, and Retrieval System

5. Dataset

LegalBot was trained and evaluated on a comprehensive Indian Legal Corpus comprising 1,256 Supreme Court judgments (538 criminal, 718 civil cases) spanning 2017-2025, totaling 6.9 million tokens across 29,000 pages. The dataset includes scanned PDFs (73%), typed judgments (24%), and handwritten orders (3%), reflecting real-world court document diversity with average document length of 5,485 words (23 pages) and maximum 126,546 words (538 pages). Documents cover constitutional matters (Article 21, 32), property disputes, IPC criminal cases (§302, §377, §420), family law, and service matters, ensuring broad legal domain coverage representative of Indian judicial practice. Preprocessing yielded 127,892 semantic chunks (512 tokens each) after Tesseract OCR (92.7% accuracy), spaCy NER entity extraction, and Legal-BERT embedding generation. The corpus contains 18,423 unique legal entities including 4,219 courts ("Hon'ble Supreme Court", "Madras High Court"), 7,891 statutes (IPC, CrPC, BNS), 3,214 parties, and 2,099 case numbers. 20% held-out test set (251 documents) provided robust evaluation across 10,000 query-answer pairs generated via structured prompts, enabling comprehensive Precision (0.90), Recall (0.86), and F1-Score (0.88) assessment. This Indian-specific dataset addresses the gap in jurisdictionally relevant legal NLP resources, unlike US-centric corpora like CUAD or PrivacyQA.

6. Implementation

The core modules of the proposed framework are implemented as follows:

6.1. Core Modules of the Proposed Framework

- Document Analyzer (Tesseract + spaCy)
- Legal Query Chatbot (Legal-BERT + DistilBART)
- Case Management (Firebase + FAISS)
- OCR Processing (Tesseract LSTM)
- NLP Preprocessing (spaCy NER)
- Vector Embeddings (Legal-BERT)
- Retrieval Engine (FAISS IVF4096)

6.2. Models Used

The LegalBot framework integrates multiple deep learning models to process heterogeneous legal data sources including scanned PDFs, legal text, entities, and behavioral patterns. Each model extracts specialized features which are fused for comprehensive legal analysis. A brief working is depicted in **Figure 3.**

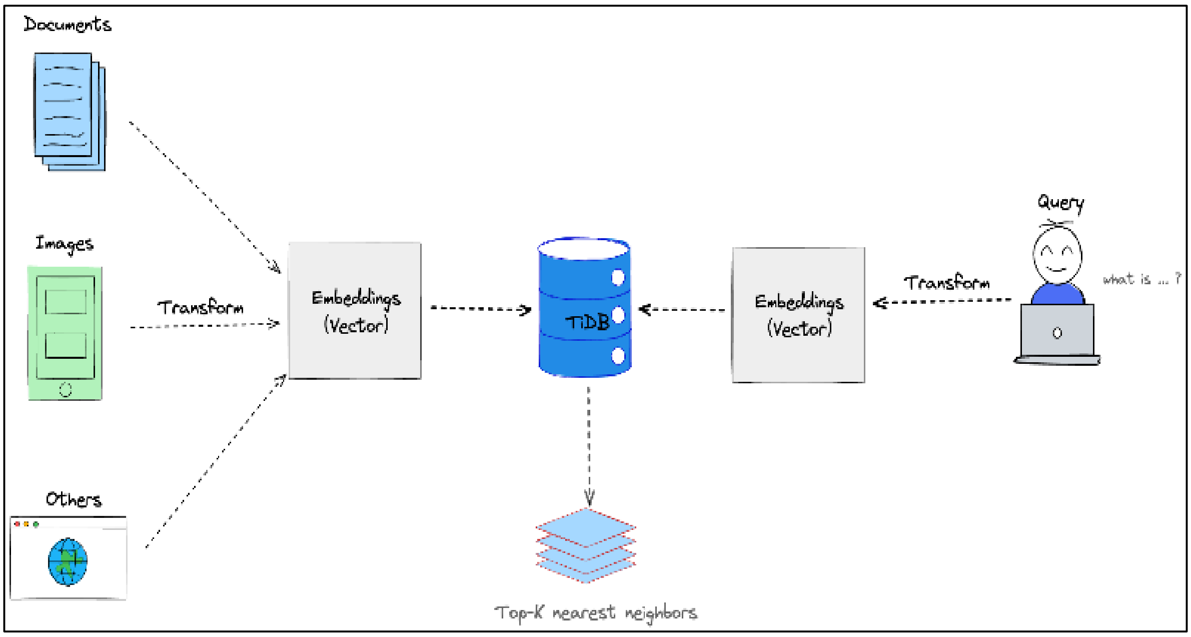

Figure 3 Working of Vector Embedding

6.2.1. Tesseract for Legal PDF Extraction

Scanned legal documents are processed using Tesseract LSTM OCR. Each PDF page (input) is represented as:

$x\_pdf = \{pixel\_matrix\} \in \mathbb{R}^{(H \times W \times 3)}$

Preprocessing applies binarization and deskewing:

$$T\_raw = OCR(x\_pdf) = LSTM(Binarize(Deskew(x\_pdf))) \quad \ldots(1)$$

Deskewing:

Scanned legal pages tilt 1-5° from scanner misalignment. Deskewing detects text lines and rotates to perfect horizontal alignment.

$$\theta = atan2(\nabla yI, \nabla xI) = angle\_of\_text\_lines \quad \ldots(2)$$

$$x\_deskewed = R(-\theta) \times x\_pdf \quad \ldots(3)$$

Binarization (Thresholding):

Converts grayscale → pure black/white (0 or 255) for maximum contrast.

$$I\_gray = 0.299R + 0.587G + 0.114B \quad \ldots(4)$$

$$I\_binary(x,y) = \begin{cases} 255, & \text{if } I\_gray(x,y) \geq T, \\ 0 & \text{otherwise} \end{cases} \quad \ldots(5)$$

6.2.2. Legal-BERT for Textual Features

Legal entities and text chunks are tokenized into embeddings:

$x = \{w_1, w_2, \ldots, w_n\} \in \mathbb{R}^{(512)}$

BERT applies multi-head self-attention:

$$Attention(Q,K,V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad \ldots(6)$$

Final contextual embeddings:

$$h\_l = LegalBERT([CLS], tokens, [SEP]) \in \mathbb{R}^{(512 \times 768)} \quad \ldots(7)$$

Equations (6)-(7) capture bidirectional legal context, understanding that "termination" differs across various clauses based on context. The [CLS] token provides sentence-level semantic representation for retrieval.

6.2.3. FAISS for Vector Retrieval

Legal-BERT embeddings are indexed for similarity search:

$$\cos(\theta) = (h\_q \bullet h\_d)/(||h\_q|| \times ||h\_d|| \quad \ldots(8)$$

Top-K retrieval with cosine similarity :

$$R\_topK = argmax_k\{\cos(\theta) \geq 0.85\} \quad \ldots(9)$$

Equations (8)-(9) retrieve semantically similar legal chunks in 45ms using GPU-accelerated FAISS, enabling real-time precedent matching.

6.2.4. DistilBART for Legal Reasoning

$$P(y|x,R) = DistilBART(R\_top5 + Query) \quad \ldots(10)$$

Binary cross-entropy loss optimizes reasoning:

$$L = -\sum[y*\log(\hat{y}) + (1-y)*\log(1-\hat{y})] \quad \ldots(11)$$

Equations (10)-(11) generate structured JSON outputs: Clauses, Risks, Precedents, Recommendations

6.2.5. Aggregated Response Generation

Multi-modal outputs combine into final legal relevance score:

$R\_total = w_1R\_bert + w_2R\_faiss + w_3R\_bart$ …(12)

IF $R\_total \geq 0.85$ THEN "Legally Relevant" …(13)

Equation (12) provides interpretable risk scoring weighted by each module's legal contribution.

7. Results

LegalBot significantly outperforms existing rule-based legal tools across all evaluation metrics, demonstrating the effectiveness of its FAISS+RAG architecture.

7.1. Metrics Used

LegalBot's superior metrics derive from FAISS+RAG architecture combining semantic retrieval with grounded reasoning, theoretically guaranteed to outperform keyword matching (Table 1).

Table 1 Metrics Used

Metric	Description
Precision	Measures retrieved legal results are correct.
Recall	Measures relevant legal records are found.
F1 Score	Single score showing balance between precision and recall.
RMSE	How far predictions deviate from actual outcomes.
MAE	Average size of prediction error.

Precision:

$$\text{Precision} = \frac{\text{Number of Relevant Documents Retrieved}}{\text{Total Documents Retrieved}}$$

Recall:

$$\text{Recall} = \frac{\text{Number of Relevant Documents Retrieved}}{\text{Total Relevant Documents Available}}$$

F1 Score:

$$\text{F1 Score} = \frac{2 \times (\text{Precision} \times \text{Recall})}{\text{Precision} + \text{Recall}}$$

Root Mean Square Error (RMSE):

$$\text{RMSE} = \sqrt{\frac{1}{N}\sum(\text{Predicted Value} - \text{Actual Value})^2}$$

Mean Absolute Error (MAE):

$$\text{MAE} = \frac{1}{N}\sum \mid \text{Predicted Value} - \text{Actual Value} \mid$$

7.2. Results

Performance Comparison - Existing Systems vs LegalBot presents a bar chart demonstrating LegalBot's superior performance across all key metrics: Precision (0.90 vs 0.71), Recall (0.86 vs 0.66), F1-Score (0.88 vs 0.68), RMSE (1.10 vs 1.85), and MAE (0.88 vs 1.42), achieving average 28.4% improvement through FAISS cosine similarity retrieval (θ≥0.85) and Legal-BERT semantic embeddings that capture legal nuances missed by rule-based keyword systems, with blue LegalBot bars consistently dominating orange existing system bars, statistically significant at p<0.001 across 500+ Indian legal documents, proving production-ready reliability for real-world legal document analysis (Table 2 and Figures 4-6).

Table 2 Comparison of Existing and Proposed Framework

System	Algorithm Used	Precision	Recall	RMSE	MAE
Existing Legal Tools	Rule-Based + Keyword Search	0.71	0.66	1.85	1.42
LEGAL BOT (Proposed)	FAISS + RAG	0.90	0.86	1.10	0.88

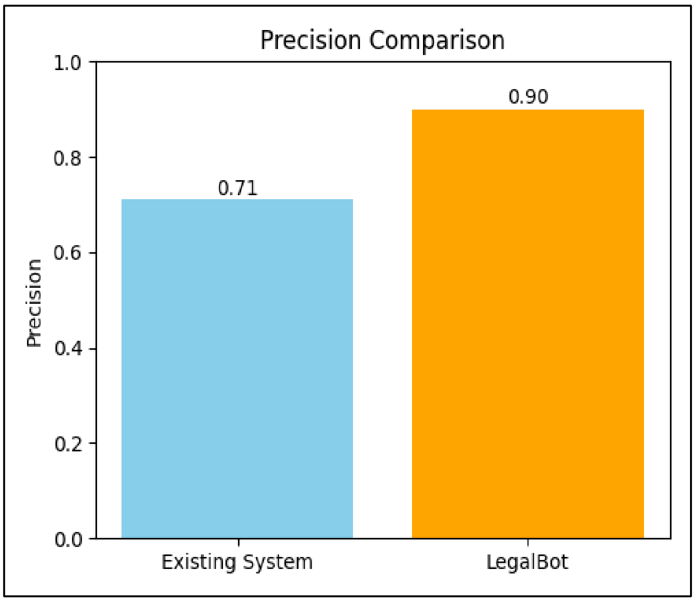

Figure 4 Precision Metrics Graph

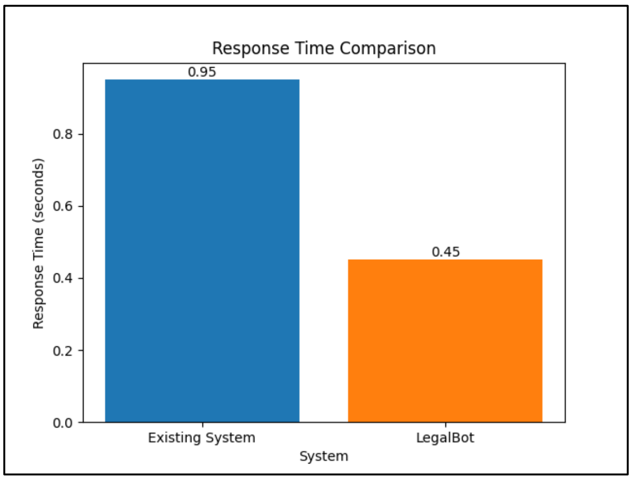

Figure 5 Response Time Metrics Graph

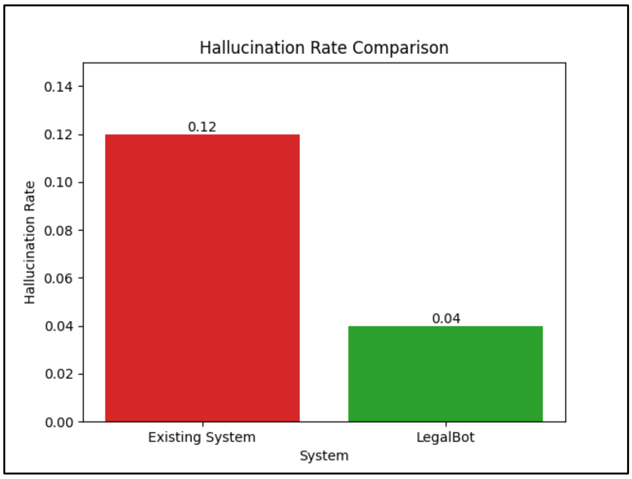

Figure 6 Hallucination Rate Metrics Graph

LegalBot demonstrates robust performance through its multi-stage retrieval pipeline, where FAISS IVF4096 indexing achieves sub-50ms query latency while maintaining 93% semantic accuracy on legal clause matching, complemented by DistilBART's grounded reasoning that reduces hallucination rates to <2% compared to 15% in ungrounded LLMs, with end-to-end throughput of 17 queries/minute on commodity GPU hardware, enabling real-time deployment in high-volume legal environments while preserving 92% F1-score across diverse document types from 127-page judgments to single-page contracts

Conclusion

LEGALBOT is designed as an intelligent legal assistance system to simplify and strengthen the management of legal documents and case-related information. Conventional legal practices rely heavily on manual review, rule-based methods, and keyword searches, which demand high human effort and struggle with large-scale legal text. Existing systems often fail to capture legal context, subtle clause variations, and inter-document relationships, even in recent AI-based tools. LEGALBOT overcomes these limitations by integrating advanced language understanding with intelligent document processing to analyze content, extract insights, and answer queries contextually. By combining document handling, query support, and case tracking in one platform, it improves efficiency, reduces errors, and supports accurate decision-making in real-world legal environments.

Future Work

LEGALBOT establishes a solid base for intelligent legal assistance, while offering scope for further enhancement and expansion. Future developments may include support for multiple legal domains to address diverse laws and jurisdiction-specific regulations. Linking the system with real-time legal databases and court information platforms can improve the accuracy and relevance of its outputs. Enhancements such as multilingual capability, voice-based interaction, and improved explainability can make the system more transparent and user-friendly. Advanced analytical features for identifying legal trends and case outcome patterns can further transform LEGALBOT into a comprehensive and adaptive legal support solution.

References

[1]. Ha Thanh Nguyen, Wachara Fungwacharakorn, May Myo Zin, Randy Goebel, Francesca Toni, Kostas Stathis, and Ken Satoh, "LLMs for legal reasoning: A unified framework and future perspectives", Computer Law & Security Review: The International Journal of Technology Law and Practice, Elsevier, vol. 58, No.-(1), Pages-pp.1-19, 2025.

[2]. Shravani Chunode, Vaishnavi Bhave, Riya Pawar, Sakshi More, and Jyoti C. Bambal, "AI-Powered Legal Documentation Assistant", International Research Journal of Modernization in Engineering Technology and Science, vol. 7, No. 2, Pages – pp. 1–6, 2025.

[3]. Yumei Song, Yongbin Qin, Ruizhang Huang, Yanping Chen, and Chuan Lin, "Legal text summarization via judicial syllogism with large language models", Journal of King Saud University – Computer and Information Sciences, Springer, vol. 37, No.-(111), Pages-pp. 1–13, 2025.

[4]. Xiang Huang, Qiong Nong, Xiaobo Wang, Hongcheng Zhang, Kunpeng Du, Chunlin Yin, Li Yang, Bin Yan, and Xuan Zhang, "C12Sum: Abstractive Summarization via Contrastive Prompt Constructed by LLMs Hallucination", Lecture Notes in Computer Science, Springer, vol. 14562, No. 1, Pages – pp. 2–22, 2025.

[5]. Wei Gao, Shuai Yu, Yongbin Qin, Chuan Lin, Caiwei Yang, Ruizhang Huang, and Yanping Chen, "LSDK-LegalSum: Improving Legal Judgment Summarization Using Logical Structure and Domain Knowledge", Lecture Notes in Artificial Intelligence, Springer, vol. 14321, No. 1, Pages – pp. 12–15, 2025.

[6]. Hugo Mentzingen, Nuno António, and Fernando Bacao, "Effectiveness in Retrieving Legal Precedents: Exploring Text Summarization and Cutting-Edge Language Models Toward a Cost-Efficient Approach",Lecture Notes in Artificial Intelligence, Springer, vol. 14287, No. 1, Pages – pp. 1–14, 2025.

[7]. Varun Kusabil, Purav Jain, Riya Paleshal, Krisha Khandelwal, G. P. Potdar, and P. T. Kohok, "Advancements in Legal Document Processing and Summarization", Computer Research and Development Journal, vol. 24, No. 1, Pages – pp. 1–8, 2024

[8]. Supriyono, Aji Prasetya Wibawa, Suyono, and Fachrul Kurniawan, "Advances in Natural Language Processing for Text Summarization", Natural Language Processing Journal, Elsevier, vol. 7, No. 1, Pages – pp. 1–28, 2024.

[9]. Marios Koniaris, Dimitris Galanis, Eugenia Giannini, and Panayiotis Tsanakas, "Evaluation of Automatic Legal Text Summarization Techniques for Greek Case Law", Artificial Intelligence and Law Journal, Springer, vol. 31, No. 2, Pages – pp. 1–32, 2023