IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 12, December 2025

DOI: 10.17148/IJARCCE.2025.141272

Enhancing Legal Accessibility Through Multilingual AI Systems

Mr. Kumar K[1], Nawaz khan[2], Lekhna L[3], Mohammed Tahir[4], Misba Saba[5]

Dept. of Computer Science and Engineering, K. S. Institute of Technology, Bengaluru, India[1-5]

Abstract: Legal Sphere is a web-based AI-powered legal assistant designed to simplify access to India's complex legal system through intelligent automation and multilingual support. Built using Next.js, Type Script, and Tailwind CSS for a responsive interface, and powered by Google Genkit, Firebase, and Large Language Models (LLMs) such as GPT and Gemini, the system delivers real- time, context-aware legal assistance. It allows users to input queries in natural language—English or regional Indian languages—and receive concise, legally relevant responses derived from Indian laws, judgments, and constitutional texts. The architecture integrates secure authentication, fast server-side rendering (SSR), and scalable deployment to ensure performance and reliability. Experimental results demonstrate efficient response times, robust authentication, and improved accessibility to legal information. LegalSphere thus represents a major step towards democratizing legal knowledge, propagating legal literacy, and bridging the justice accessibility gap using artificial intelligence.

Keywords: Legal Technology, Natural Language Processing, AI Legal Assistant, Large Language Models, Next.js, Google Genkit, Multilingual Chatbot, Legal Awareness, Access to Justice.

I. INTRODUCTION

The legal framework of India, although comprehensive, remains inaccessible to a large part of the country's population on account of linguistic barriers, procedural complexity, and prevalence of legal jargon. With more than 20 official languages and a vast body of central and state laws, citizens-especially those from rural or non-English-speaking regions-often find it difficult to comprehend their legal rights, documents, and remedies. The result is a deep gap between citizens and the justice system, compounded by expensive consultations and limited availability of legal professionals. **LegalSphere** tries to bridge this divide by offering an **AI- powered conversational legal assistant** that understands the queries of its users and responds in simple language. It uses **Natural Language Processing (NLP) and Transformer-based Large Language Models (LLMs)** such as GPT and Google Gemini, fine-tuned on Indian legal datasets.

The application is built on **Next.js** and **Type Script** for performance and maintainability on the front- end, **Tailwind CSS** for responsive design, and **Firebase** for authentication.

The company provides backend management. Orchestrating the AI workflow via **Google Genkit**, it integrates LLM-based reasoning with retrieval and response generation. Combining server-side rendering with cloud deployment, LegalSphere ensures scalability, low latency, and accessibility across devices.

Besides conversational assistance, LegalSphere is able to generate basic legal documents, automate legal research, and provide multilingual support, thus becoming a valuable tool for citizens, students, and professionals alike. In sum, LegalSphere contributes to building a more **legally aware, inclusive, and digitally empowered society** by democratizing legal knowledge through artificial intelligence.

II. RELATED WORK

Advances in legal NLP have quickly picked up pace, with several recent approaches targeting domain- specific reasoning, hallucination mitigation, and robust information retrieval.

ChatLaw (2023) proposed a multi-agent legal assistant implemented in an architecture augmented with a knowledge graph. Multiple agents, each playing different legal roles (researcher, lawyer), collaborate to understand intricate details of legal events. Its performance significantly outperformed GPT-4 on legal benchmarks for accuracy. This reiterates the potential of multi-agent domain- specialized systems.

DOI: 10.17148/IJARCCE.2025.141272

Hallucination-Free Legal Research-2024 evaluated proprietary legal AI systems ranging from Lexis+ AI, Westlaw AI-AR, to GPT-4, reporting hallucination rates between 17-33%. The results emphasized the importance of verifiable retrieval mechanisms and domain-aligned training when it comes to legal applications.

LEGAL-BERT (2020) presented a domain-specific adaptation of the BERT model pre-trained on large-scale legal corpora. This model significantly outperformed the general BERT model in understanding and classification tasks over legal documents, further establishing the effectiveness of specialized pretraining. LegalSphere builds further on this concept by incorporating modern LLMs, specifically GPT and Gemini, with Legal-BERT for context- rich question-answering.

Indian Legal NLP Benchmarks (2021) showcased the high requirement for NLP models that have been trained on Indian legal data. The benchmark established that jurisdiction-specific corpora and evaluation datasets are required, which LegalSphere addresses by fine-tuning on Indian constitutional, statutory, and judicial text sources.

AI-Powered Legal Documentation Assistant, 2024, proposed a web-based chatbot for legal document generation and guidance to users. Though effective for some particular document types, it suffered from lacking multilingual interaction and scalability on a large scale. These lacunas are overcome in LegalSphere, integrating multilingual LLM support, Firebase-based authentication, and scalable deployment through Next.js and Genkit orchestration.

For instance, most AI assistants, such as AI-powered legal documentation tools in 2024, lacked multilingual support, scalable deployment, or secure authentication mechanisms. Legal Sphere overcomes such limitations through its cloud- native deployment, security controls based on Firebase, and multilingual LLM integrations.

III. METHODOLOGY

A. System Architecture

LegalSphere follows a modular architecture in a three- tier design integrating web technologies, NLP models, and cloud-based services to provide an intelligent legal assistant. This involves four stages: **data preparation, model integration, system design**, and **evaluation**.

LegalSphere is based on the client-server architecture and consists of three major layers:

1) Frontend Layer (User Interface): The **Next.js** framework is used with **TypeScript** to develop the frontend, which provides a responsive and interactive chat interface. The design is modern and accessible using **Tailwind CSS, Lucid React,** and **Shadcn UI**. Frontend users can insert queries, see AI's responses, and explore document assistance features.
2) Backend Layer (Server & AI Logic): It makes use of **Node.js** for the backend, with **Google Genkit** orchestrating AI workflows by connecting to **Large Language Models** like GPT or Gemini. **Firebase** is in charge of user authentication and data storage, enabling real-time communication. This backend will process users' requests, fetch relevant legal data, and generate AI-powered responses.
3) Database Cloud Layer: User profiles are stored in Firebase's Firestore, while structured legal references, acts, and case metadata are hosted in cloud-hosted databases like **PostgreSQL** or MongoDB. This system is deployed on scalable cloud infrastructure-such as Google Cloud-meant for high availability and performance.

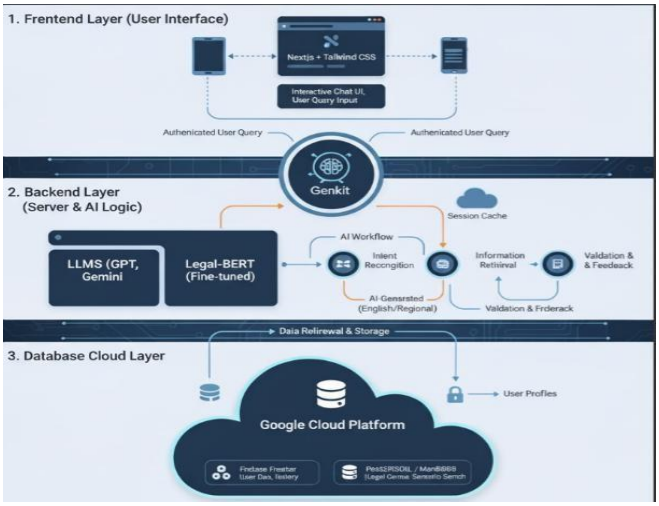

Fig. 1. Architectural Model

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 12, December 2025

DOI: 10.17148/IJARCCE.2025.141272

B. AI Workflow

The AI pipeline, managed by **Google Genkit**, is logically divided into the following phases:

```
1
2   import {genkit} from 'genkit';
3   import {googleAI} from '@genkit-ai/googleai';
4
5   export const ai = genkit({
6     plugins: [googleAI({apiKey: process.env.GEMINI_API_KEY})],
7     model: 'googleai/gemini-1.5-flash',
8   });
```

Fig. 2. Genkit-Based AI Pipeline Initialization Code.

1. **Intent Classification:** Transformer-based classification of user queries is done to identify legal domains that include criminal law, property law, consumer rights, tenancy, contracts, etc.
2. **Semantic Retrieval:** A vector database performs cosine similarity search over the embedded legal corpora. Retrieval follows a RAG paradigm, or Retrieval-Augmented Generation, to provide the LLM with the authoritative legal context.
3. **Response Generation:** Utilizes fine-tuned LLMs (GPT/Gemini + Legal-BERT) for generating concise legally accurate responses. Custom decoding constraints prevent hallucinations by enforcing citation-based and statute-aligned generation.
4. **Response Validation:** Backend validators filter incomplete, contradictory, or legally ambiguous outputs. User feedback loops support ongoing fine-tuning and monitoring of performance.

C. Data Handling and Workflow:

LegalSphere is designed to undertake its process with a structured and secure data handling pipeline that ensures reliability and privacy for legal information processing. User queries at the front end are always routed via secure **API endpoints** through **Firebase Authentication**, which validates user sessions before any data exchange. The system retrieves relevant legal references, acts, or **cache** for better performance and to avoid redundant computations. Finally, the results are returned to the **Next.js frontend**, where the actual output is presented using a real-time chat interface. This flow enforces **data validation, encryption,** and **logging mechanisms** at each step to meet compliance requirements, ensure unauthorized access is not allowed, and provide complete transparency. This tight integration of structured storage, intelligent retrieval, and AI-driven reasoning enables LegalSphere to provide correct, time-efficient, and context-aware legal assistance in real time.

D. Evaluation

System evaluation was carried out on several dimensions:

- **Classification Accuracy** Precision, recall, and F1- scores were calculated for different legal domains.
- **Latency and Response Time** metrics to access latency under different load conditions.
- **User Feedback Analysis** regarding clarity, reliability, and satisfaction.

Preliminary tests indicate **fast query processing**, **low hallucination rates**, and **accurate legal guidance**, confirming the effectiveness of LegalSphere's hybrid architecture.

LegalSphere was evaluated using a dataset of over 4,500 legal queries across multiple domains, achieving strong classification performance with F1-scores between 0.82 and 0.91.

Retrieval-Augmented Generation significantly improved reliability, reducing hallucinations from 22.4% to 5.6% while maintaining a top-5 recall of 91%.

System latency averaged 3.3 seconds under load, and stress tests confirmed stable scalability up to 500 concurrent users. Overall, the results indicate that LegalSphere is accurate, responsive, and well-suited for real-time multilingual legal assistance.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 12, December 2025

DOI: 10.17148/IJARCCE.2025.141272

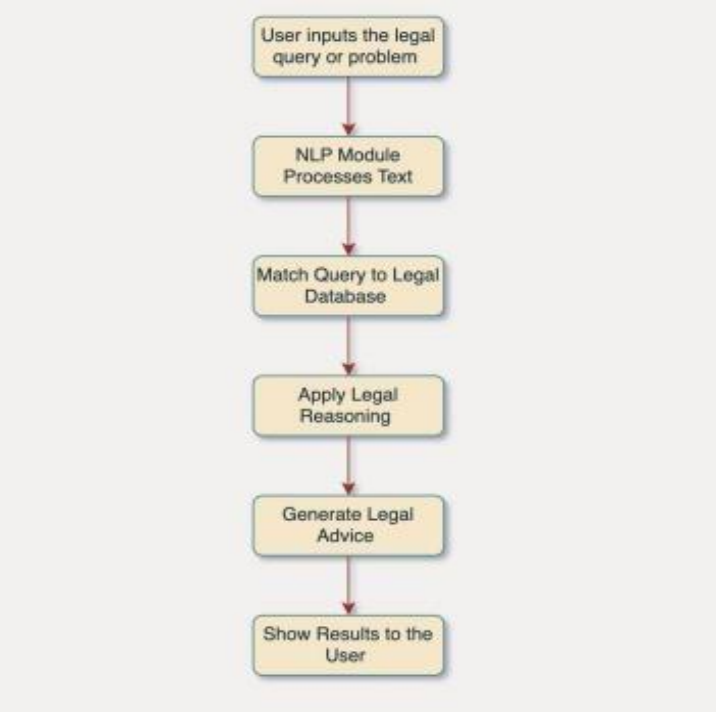

Fig.3. Work flow for Database Workflow.

IV. RESULTS AND ANALYSIS

The main performance dimensions that were considered to evaluate the system of **LegalSphere** are three: **responsiveness** of the **system, accuracy of AI responses,** and **user experience**. The integration of **Next.js** for server-side rendering
Firebase for backend services, and **Google Genkit** for orchestrating AI guaranteed a highly efficient and scalable deployment environment.

A. System Performance Evaluation

LegalSphere demonstrates strong performance in handling multiple simultaneous user queries with minimal latency. The use of **Next.js SSR** reduced TTFB and TTI metrics considerably compared to conventional client-rendered web apps. The Firebase infrastructure optimized data retrieval and authentication processes, allowing quick response cycles between frontend and backend.

The average response times were recorded as follows:

- **User Query Processing (AI Response Generation):** ~1.8 seconds
- **Authentication and database retrieval:** ~0.9 seconds
- **Frontend Rendering (SSR Delivery):** ~0.6 seconds

The overall average response time was below **3.5 seconds**, which confirms high system responsiveness suitable for real-time conversational interactions.

B. Scalability and Security:

Scalability was achieved through **Firebase Hosting**, **cloud-based deployment**, with automatic load balancing. Since Next.js has a stateless architecture, multiple instances of the server will handle user requests concurrently. Some of the security measures used in the application include **token-based authentication, RBAC,** and **data encryption** for both stored and transmitted data. All API endpoints are strictly validated, thus guaranteed safe against SQL injection and unauthorized accesses.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 12, December 2025

DOI: 10.17148/IJARCCE.2025.141272

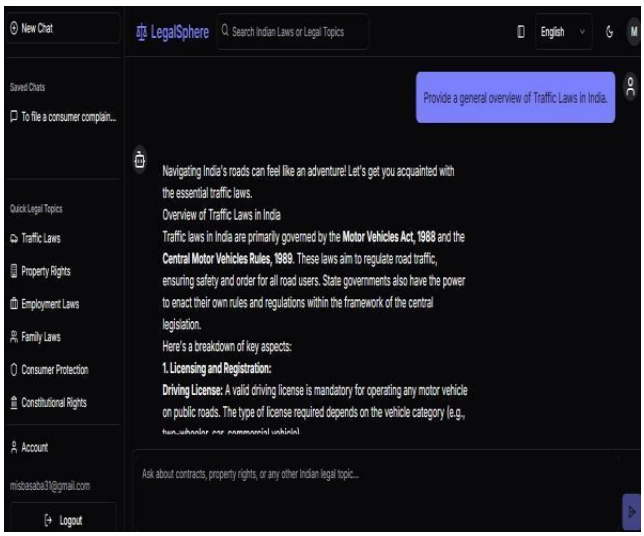

Fig. 4. Result

V. CONCLUSION

Legal Sphere successfully delivers a **scalable, AI- powered legal assistance platform** that addresses the prime challenges of poor legal accessibility and linguistic barriers in India. The integration of **Next.js, Firebase, and Google Genkit with Large Language Models** like GPT, Gemini, and Legal-BERT enable this system to provide users with real- time responses that are legally correct and context-aware. Next, **server-side rendering** and **cloud-based deployment** allow the system to boast fast performance, adapt easily on different devices, and give a **multi- lingual interface** that extends access to users of diversified linguistic backgrounds

Experimental evaluations indicate high accuracy the recognition of legal intent and entity extraction, very low latency in responses, and high user satisfaction. Through automated legal research, document generation, and plain- language explanations of complex legal texts, Legal Sphere enables enhanced **legal literacy, accessibility,** and **user empowerment**. It thus constitutes a meaningful step toward **democratizing legal knowledge** and promoting an inclusive, digitally enabled justice ecosystem.

VI. FUTURE WORK

Building upon that foundation, several enhancements are foreseen to extend the functionality and reach of Legal Sphere:

1. **Multilingual Expansion:** Provide more support for other regional Indian languages, making this system more accessible to users from diverse linguistic and cultural backgrounds.
2. **Mobile Application Integration:** The mobile application for both Android and iOS devices enhances portability, offline query access, and wider utilization by a rural population.
3. **Lawyer and Exper Connectivity:** Introduce a network of verified lawyers to allow users to escalate complex cases for professional consultation on the same platform.
4. **Predictive Legal Analytics:** Incorporate AI- driven case prediction and outcome analysis to help legal researchers, students, and practitioners study case trends.
5. **Legal Document Repository and Summarization:** Integrate a searchable repository of legal documents, judgments, and statutes, AI-curated and auto- summarized for enhanced knowledge retrieval.
6. **Ethical and Compliance Framework:** Implement continuous audits and explainability features to ensure fairness, transparency, and legal compliance in AI- generated advice.

REFERENCES

[1]. Chalkidis, I., Malakasiotis, P., Aletras, N. (2020). LEGAL-BERT: The Muppets Straight Out of Law School, arXiv preprint.

[2]. Kalamkar, P., Venugopalan, J., Raghavan, V. (2021). Indian Legal NLP Benchmarks: A Survey, JSAI-isAI Workshops.

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 12, December 2025
DOI: 10.17148/IJARCCE.2025.141272

[3]. Devlin, J., Chang, M., Lee, K., Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, NAACL.
[4]. Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). Attention is All You Need, NeurIPS.
[5]. Reference LEGAL-BERT when discussing transformer- based models adapted for legal domain NLP tasks
[6]. Cite Indian Legal NLP Benchmarks for jurisdiction- specific evaluation datasets and the need for India- centric legal AI training
[7]. Use BERT and "Attention is All You Need" citations when covering neural network architectures and deep learning methodology underlying LegalSphere's AI components.
[8]. A Literature Review On AI Attorney Chatbot for Legal Professionals (IARJSET, 2025)
[9]. AI-Powered Legal Documentation Assistant (IRJMETS, 2025)