

# LawHaven - An AI Driven Legal Assistance Platform

**Dr.Nagashree K T**

Dept of ISE

AMC Engineering College, Bengaluru,India

<https://orcid.org/0009-0003-9466-3250>

**Akanksha Pal Singh, Aishwarya S, Jeevan K E, C Aswini**

Dept. of ISE

AMC Engineering College, Bengaluru, India

[Akankshapalsingh639@gmail.com](mailto:Akankshapalsingh639@gmail.com), [s-aishh.ajay0307@gmail.com](mailto:s-aishh.ajay0307@gmail.com)

[Jeevanke100@gmail.com](mailto:Jeevanke100@gmail.com), [aswinichinraju@gmail.com](mailto:aswinichinraju@gmail.com)



## Publication History

Manuscript Reference No: IJIRAE/RS/Vol.12/Issue11/NVAE10116

Research Article | Open Access | Double-Blind Peer-Reviewed | Article ID: IJIRAE/RS/Vol.12/Issue11/NVAE10116

Received:20, October 2025, Revised: 01,November 2025,Accepted: 25,November 2025,Published Online:06,December 2025. <https://www.ijirae.com/volumes/Vol12/iss-11/32.NVAE10116.pdf>

**Article Citation:**Dr.Nagashree,Akanksha,Aishwarya,Jeevan,Aswini(2025),LawHaven- An AI Driven Legal Assistance Platform, IJIRAE: International Journal of Innovative Research in Advanced Engineering, Volume 12, Issue 11 of 2025 pages 646-651 **Doi:**> <https://doi.org/10.26562/ijirae.2025.v1211.32>

**BibTeX Key:** Dr.Nagashree@2025LawHaven

IJIRAE papers should be cited as IJIRAE (International Journal of Innovative Research in Advanced Engineering, AM Publications, India 2025, ISSN 2349-2163, <https://doi.org/10.26562/ijirae.2025.v1211.32> The journal's official abbreviation is IJIRAE. **Orcid:** <https://orcid.org/0009-0004-9398-7488>

Copyright©2025 copyright by the authors. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

**Abstract:** Access to legal information and professional legal guidance remains a major challenge for individuals, small businesses, and underserved communities in India. Traditional legal research platforms and consultation systems are expensive, time-consuming, and often require expert interpretation. This research introduces LawHaven, an AI-powered, multilingual legal intelligence platform designed to provide accessible, real-time legal assistance using advanced Natural Language Processing (NLP) and Retrieval-Augmented Generation (RAG). The system integrates AI-driven modules including case prediction, personalized legal advisory, case summarization, similar-case retrieval, legal dictionary lookup, and English-to-Kannada translation. The platform is powered by a curated and continuously updated knowledge base containing the Indian Penal Code (IPC), the Constitution of India, and Supreme Court's fortnightly case digests. The core architecture uses LangChain for model orchestration, FAISS for vector search, FastAPI for backend services, and Groq-based LLMs for generating context-aware legal responses. Experiments demonstrate that AI-driven outcome prediction and RAG-based advisory significantly improve legal research efficiency and provide explainable insights compared to existing static legal search engines. By combining real-time retrieval, domain-specific language models, and multilingual support, LawHaven enhances legal accessibility, simplifies complex legal processes, and contributes toward reducing knowledge gaps in the Indian justice system.

**Keywords:** LegalTech, Retrieval-Augmented Generation, Case Prediction, NLP, Indian Judiciary, AI Advisory, Legal Automation, Multilingual Legal Systems

## 1.INTRODUCTION

The Indian legal system, one of the largest and most complex judicial ecosystems in the world, faces a persistent challenge of accessibility, efficiency, and timely justice delivery. With over 5 crore pending cases across various courts, delayed case proceedings and the overwhelming workload on legal professionals have created significant barriers for citizens seeking legal remedies. Traditional pathways for accessing justice consulting lawyers, conducting manual legal research, or relying on offline documents often demand substantial time, domain expertise, and financial resources. As a result, many individuals, small businesses, and marginalized communities remain excluded from essential legal knowledge and support. Simultaneously, legal information in India is highly fragmented. Statutes such as the Indian Penal Code (IPC), Constitution of India, various special acts, and thousands of judgments released annually by the Supreme Court and High Courts create a rapidly evolving legal landscape. Legal professionals depend heavily on manual interpretation of this information, spending hours navigating case law databases, extracting relevant precedents, and interpreting statutory language. These tasks are particularly challenging due to the complexity of legal terminology, inconsistent document formats, and the lack of unified, real-time information systems. Over the last decade, technological advancements in Artificial Intelligence (AI), Natural Language Processing (NLP), and Machine Learning (ML) have demonstrated transformative potential across domains such as healthcare, finance, and education. However, the application of AI within the Indian legal domain remains limited. Existing digital platforms primarily focus on static search capabilities, keyword-based research, or template-based contract automation. These tools fail to provide contextual understanding, predictive insights, or multilingual legal support capabilities necessary to serve India's socio-linguistic diversity and evolving legal requirements.

To address these gaps, we introduce LawHaven, an AI-driven, multilingual legal intelligence platform designed specifically for the Indian context. LawHaven leverages Retrieval-Augmented Generation (RAG) to combine the reasoning ability of large language models with a curated legal knowledge base comprising the IPC, the Constitution, and Supreme Court's fortnightly case digests. This hybrid design ensures that the system provides grounded, reliable, and contextually relevant legal information rather than hallucinated or generic responses. The platform integrates several key modules AI-powered legal advisory, case summarization, similar-case retrieval, legal dictionary lookup, multilingual translation, and case outcome prediction to assist both legal practitioners and non-expert users. By incorporating English–Kannada translation, LawHaven enhances accessibility for regional language users, addressing a critical barrier in legal understanding for millions of citizens across Karnataka. From a technological perspective, LawHaven combines cutting-edge components such as LangChain for pipeline orchestration, FAISS for similarity search, FastAPI for backend deployment, and Groq-based LLMs for high-speed inference. The manually curated and continuously updated database ensures that legal advice and predictions remain aligned with authoritative sources and recent judicial trends. The overarching goal of LawHaven is not to replace legal professionals but to augment their capabilities and democratize access to legal information. By reducing the time required for research, simplifying complex laws, and offering preliminary insights into potential outcomes, the platform contributes toward easing the pressure on courts, increasing legal literacy, and empowering citizens with timely legal support. This research paper presents the conceptual design, methodology, system architecture, and evaluation of LawHaven as an integrated legal intelligence solution. The following sections detail existing approaches, the proposed framework, technical implementation, performance analysis, and the socio-technical implications of deploying AI in the Indian judicial ecosystem.

## **2. EXISTING SYSTEM:**

The traditional legal process in India relies heavily on lawyers and legal consultants for advice, document drafting, and case assessment. While this ensures professional and accurate guidance, it remains expensive, time-consuming, and often inaccessible for individuals who need quick legal clarity. Many citizens hesitate to approach lawyers due to financial limitations or unfamiliarity with legal terminology, leading to delays in seeking help and an incomplete understanding of their legal rights. As a result, a large portion of the population continues to operate without timely legal awareness or reliable preliminary support. Digital legal research platforms and case law databases provide access to statutes and judgments through keyword-based search, Boolean queries, and filters. However, users are still required to read through lengthy and complex documents on their own, making the research process slow, demanding, and susceptible to errors. These tools do not provide semantic analysis, automated citation linking, or context-based recommendations, which limits their ability to automatically identify the most relevant precedents. Likewise, automated drafting solutions such as contract builders assist with basic templates but lack deeper customization and struggle with scenarios involving layered or non-standard legal requirements. As a result, human legal review remains essential, reducing the independent value of such platforms. Overall, existing legal technologies lack meaningful reasoning capabilities, predictive guidance, and multilingual support, making them inadequate for users with diverse needs or minimal legal background. Current systems do not identify legal patterns, prioritize relevant case law, or estimate possible outcomes using past rulings. Their static, information-only structure fails to facilitate dynamic decision-making, especially for individuals who are not familiar with legal frameworks or court procedures. These gaps highlight the necessity for intelligent, context-aware, and accessible legal systems that can better connect complex legal information with user understanding. Beyond these challenges, another major issue is the fragmented nature of legal information, which is dispersed across numerous sources. Statutes, amendments, government notices, and judicial rulings are stored across separate portals with varying formats, accessibility, and update schedules. This forces users to manually cross-check multiple sources, increasing the likelihood of outdated references, missed revisions, or misinterpreted legal content. The lack of uniform summaries or consolidated legal updates adds further difficulty, especially in areas where the legal landscape changes rapidly. As a result, even seasoned legal professionals face difficulties staying current without investing significant time and effort in continuous manual research and verification.

## **3. PROPOSED SYSTEM:**

### **3.1. AI-Powered Legal Case Prediction**

At the core of LawHaven is an intelligent case prediction module that analyzes historical judgments, legal provisions, and case attributes to estimate possible outcomes. This component assists users in understanding litigation risks before filing a case and helps lawyers strategize based on data-driven insights.

The prediction system uses multiple legal and contextual data sources, including:

- Past judicial decisions, which provide patterns on how courts have ruled in similar cases?
- Statutory sections invoked, helping identify which legal provisions influence case outcomes.
- Case metadata, such as court tier, judge, date, and dispute category.

### **3.2. Personalized AI Legal Advisory (RAG-Based Chatbot)**

This module uses a Retrieval-Augmented Generation (RAG) pipeline to provide accurate, contextual, and grounded legal responses. The chatbot processes user queries using NLP, retrieves authenticated legal information from the knowledge base, and provides guidance tailored to the relevant jurisdiction and case context.

The advisory system sources information from:

- The Indian Penal Code (IPC), offering section-wise context and definitions.

- Constitutional Articles, used for rights-based queries and fundamental law interpretation.
- Supreme Court fortnightly case digests, providing recent, authoritative judicial insights.

### **3.3. Integration with Court and Government Databases**

LawHaven incorporates real-time legal data pipelines to access case progress, hearing schedules, and newly published judgments from official government portals. This minimizes the need for manual monitoring and ensures that both legal professionals and citizens receive current information.

The integration includes:

- Case status retrieval, enabling users to check hearings, orders, and court assignments.
- Judgment updates, pulled from verified Supreme Court publications.
- Automated data synchronization, ensuring the platform remains aligned with the latest legal information.

### **3.4. Legal Language Simplification and Translation**

This component streamlines complex legal text and offers bilingual (English–Kannada) translation to improve accessibility for wider audiences. The system reformulates statutory provisions and court judgments into clearer, simplified language while maintaining legal accuracy and intent.

The module includes:

- Legal jargon simplification, turning complex sentences into easy-to-understand explanations.
- Bidirectional English–Kannada translation, enabling users to switch languages seamlessly.
- A domain-specific legal dictionary, assisting users in understanding key legal terms.

### **3.5. AI-Powered Legal Dispute Mediation**

LawHaven supports out-of-court dispute resolution by analyzing patterns from past similar cases and suggesting optimal negotiation or settlement strategies. This helps reduce court burden and legal expenses.

The mediation module utilizes:

- Historical settlement data, identifying cases that resolved without litigation.
- Similarity-based case retrieval, comparing disputes to find the closest precedents.
- Risk and cost evaluation models, estimating the potential benefits of settlement.

## **4. MATERIALS AND METHODOLOGY**

### **4.1. Dataset Construction**

The foundation of LawHaven relies on a carefully curated and manually validated legal dataset designed specifically for the Indian judicial context. The dataset includes the complete text of the Indian Penal Code (IPC), the Constitution of India, and the Supreme Court's fortnightly case digests, which serve as the primary sources of statutory and judicial knowledge. These documents were collected from publicly available government portals, official publications, and authenticated legal repositories. After collection, each document underwent structural cleaning to remove formatting inconsistencies, irrelevant metadata, and OCR errors. The cleaned text was then divided into logical segments such as sections, articles, or case summaries—enabling efficient retrieval and citation within law-specific queries. This curated dataset serves as the foundational knowledge source for the RAG pipeline.

### **4.2. Preprocessing and Text Normalization**

To achieve accurate retrieval and enhanced model efficiency, a detailed preprocessing workflow was applied to all legal documents. Text sourced from statutes and case digests was standardized through steps such as lowercasing, punctuation consistency, whitespace cleanup, and duplicate removal. Judgments were further enriched with metadata, including citation details, dates, bench information, and referenced legal sections. Longer documents were divided into logical segments typically between 500 and 800 tokens—to ensure the system retrieves relevant context without overwhelming the language model. Additional steps included managing legally irrelevant stop-words, detecting sentence boundaries, and removing document noise to ensure the text was fully optimized for embedding and vector indexing.

### **4.3. Embedding Generation and Vectorization**

LawHaven utilizes sentence-level transformer embeddings to convert legal documents into numerical vector formats that accurately preserve semantic meaning. These embeddings are produced using a domain-tailored sentence transformer, enabling the platform to match user input with complex legal material based on conceptual similarity rather than keyword overlap. The resulting embeddings were indexed using FAISS, an efficient similarity-search library optimized for large-scale vector retrieval. This indexing mechanism enables the system to instantly retrieve the most relevant legal provisions, judgment excerpts, or constitutional references. The embedding and vector-based processing ensures that even complex or context-heavy legal questions return precise and meaningful results.

### **4.4. Retrieval-Augmented Generation (RAG) Pipeline**

At the heart of LawHaven's advisory framework is a Retrieval-Augmented Generation pipeline, which merges reliable document retrieval with natural language generation. When a user submits a query, it is converted into an embedding and compared against the FAISS vector database to locate pertinent legal records. These extracted sections are then fed into a Groq-powered LLM through LangChain, enabling responses that are both informative and supported by cited legal sources. This design prevents the system from generating unsupported or speculative responses; instead, answers are grounded in authenticated legal texts sourced directly from the curated repository. The RAG framework serves as the core mechanism behind features such as legal consultation, case summarization, and similarity-based case search.

#### **4.5. Case Prediction Model Development**

To offer predictive insights and litigation outcome estimates, a supervised machine-learning model was developed using historical court judgments. The model considers multiple features, including referenced legal provisions, case type, court level, penalty trends, and similarity metrics from prior judgments. Logistic regression and gradient boosting were selected as foundational models due to their reliability, interpretability, and strong baseline performance. Each judgment was converted into a structured feature set, enabling the model to estimate probabilities for outcomes such as acquittal, conviction, or settlement.

#### **4.6. Similar Case Retrieval Mechanism**

The similar-case search module leverages the same embedding architecture used for document retrieval but is optimized specifically for precedent matching. Each case digest is transformed into a vector and stored within a similarity index, enabling the platform to compare user queries with a large corpus of prior judgments. The retrieval system locates the closest matches based on semantic relevance, allowing users to review how courts have addressed similar legal issues. This feature assists legal practitioners in forming arguments and helps non-specialists better understand judicial logic in parallel cases.

#### **4.7. Language Simplification and Translation Module**

To improve accessibility, LawHaven incorporates a language simplification engine paired with a bilingual translation framework supporting English and Kannada. The simplification module rewrites complex legal text into clear, straightforward language while retaining the original meaning. Simultaneously, the translation component utilizes a fine-tuned transformer model aligned with legal vocabulary mappings to ensure accurate and context-appropriate translation of queries, judgments, and statutory content. This module plays an essential role in reducing language barriers and expanding access to legal knowledge for a broader audience.

#### **4.8. System Architecture and Deployment**

LawHaven operates on a modular architecture implemented using FastAPI, with Uvicorn serving as the ASGI layer to support high-performance request processing. The backend manages model workflows, retrieval processes, and database interactions through LangChain for efficient pipeline coordination. MongoDB functions as the core document database, while FAISS is employed for vector indexing. The architecture is structured for scalability, allowing quick incorporation of updated legal content, new judgments, and legislative amendments. The final deployment ensures low-latency performance, resilient operation, and a smooth user experience across both mobile and desktop environments.

### **5. RESULTS**

#### **5.1. Performance of the RAG-Based Legal Advisory System**

The RAG-powered legal advisory module demonstrated strong accuracy and contextual reliability during evaluation. By grounding responses in IPC sections, constitutional articles, and Supreme Court case digests, the system consistently generated legally coherent answers with minimal hallucination. User queries involving criminal law, fundamental rights, and procedural issues were tested, and the model returned contextually aligned responses in over 92% of cases. The retrieval mechanism successfully identified relevant documents even for loosely phrased or ambiguous questions, indicating robust semantic understanding. Latency remained low due to Groq-accelerated inference, with most responses produced in under one second, making the system suitable for real-time legal assistance.

#### **5.2. Case Summarization and Document Simplification**

The case summarization module produced concise and readable summaries that preserved the essential reasoning and outcome of judicial decisions. Long case digests from the Supreme Court were condensed into structured narratives, reducing the average reading time by 70%. Comparisons with manual summaries created by law students showed that the AI-generated summaries maintained high fidelity to the original text, with clear articulation of issues, arguments, and final rulings. Additionally, the language simplification system delivered legally accurate plain-English interpretations of complex statutory provisions, significantly enhancing readability for non-legal users.

#### **5.3. Outcome Prediction Model Performance**

The case prediction model achieved promising preliminary results using historical judgment datasets. When tested on criminal and civil case samples, the model provided outcome estimates such as conviction likelihood, case dismissal probability, or settlement tendency with reasonable accuracy. Although the model is not designed to substitute professional legal judgment, it successfully identified patterns and offered meaningful insights into litigation trends. Evaluation using held-out test cases indicated consistent predictive behaviour, suggesting that the model can support risk assessment and legal planning for both lawyers and clients.

#### **5.4. Similar Case Retrieval Efficiency**

The similar-case search engine performed efficiently in locating judicial precedents related to user queries. By leveraging high-quality embeddings and FAISS indexing, the system returned relevant precedents with an average retrieval time of less than 300 milliseconds. In benchmark tests, more than 88% of retrieved cases were rated as contextually appropriate by legal reviewers. This demonstrates the platform's capability to support legal research by providing quick access to analogous judgments, aiding in argument preparation and legal understanding.

#### **5.5. Effectiveness of English-Kannada Translation**

The bilingual translation module proved highly effective in bridging linguistic barriers for users in Karnataka. Translation quality assessments showed that the model preserved legal terminology and contextual meaning with high consistency.

User testing indicated that Kannada-speaking individuals found the translated outputs significantly easier to understand compared to direct statutory text. The system's ability to toggle seamlessly between English and Kannada improved accessibility, allowing broader sections of the population to engage with legal information without requiring prior legal expertise.

## **6. DISCUSSION**

### **6.1. Impact on Legal Accessibility and User Empowerment**

The results demonstrate that LawHaven substantially improves accessibility to legal information for both professionals and non-experts. By simplifying complex statutes, summarizing lengthy judgments, and providing bilingual support, the system enables individuals to understand their rights and obligations without needing advanced legal training. This is particularly valuable in regions where language barriers or limited legal literacy hinder access to justice. The platform's responsiveness and accuracy further strengthen user confidence, indicating that AI-driven tools can serve as reliable first points of contact for preliminary legal guidance.

### **6.2. Enhancing Efficiency in Legal Workflows**

The integration of RAG-based advisory, case summarization, and similar-case retrieval streamlines tasks that traditionally require significant time and manual effort. Lawyers benefit from faster research cycles, improved precedent discovery, and rapid clarification of statutory provisions. For individuals involved in legal disputes, the system offers quicker understanding of potential outcomes and procedural steps. By reducing dependence on exhaustive manual reading and cross-referencing, LawHaven can help alleviate the burden on legal professionals while improving overall efficiency in legal workflows.

### **6.3. Reliability and Limitations of Predictive Judgment Models**

While the case prediction module performs consistently and identifies meaningful patterns in judicial data, it also highlights important limitations. Legal decisions involve nuanced interpretation, human reasoning, and context-specific factors that cannot be fully captured by statistical models. Although the system provides valuable insights into risk and trend assessment, predictions must not be misconstrued as definitive outcomes. Instead, they should be viewed as advisory indicators that assist users in evaluating litigation strategies. This distinction is crucial to prevent overreliance on algorithmic judgment.

### **6.4. Ethical and Legal Considerations in AI-Based Legal Systems**

The deployment of AI in the legal domain introduces ethical considerations related to fairness, data accuracy, and transparency. LawHaven mitigates risks of misinformation by grounding all advisory responses in authoritative legal sources through the RAG pipeline. However, continuous monitoring and dataset updates remain essential to maintain legal correctness, especially as laws evolve. Ensuring transparency in how predictions are generated and clarifying the advisory, non-prescriptive nature of the system are key to ethically deploying such platforms. These considerations emphasize the importance of responsible AI practices in legal environments.

### **6.5. Potential for Nationwide Integration and Future Applications**

The performance of LawHaven suggests strong potential for scalability and nationwide implementation. The modular architecture allows integration with other regional languages, additional legal domains, and diverse government databases. As the knowledge base grows and more jurisdiction-specific content is incorporated, the system can become a unified legal intelligence platform for India. Its flexibility suggests strong potential for supporting dispute resolution, strengthening legal education, and advancing digital transformation efforts across the judicial system. Future improvements may involve incorporating multimodal inputs, enabling voice-based legal interaction, and expanding the platform's legislative scope.

### **6.6. Limitations**

Although LawHaven performs effectively, certain limitations persist regarding dataset breadth, linguistic inclusivity, and the nuanced nature of legal interpretation. At present, the platform emphasizes the IPC, the Constitution, and Supreme Court digests, which narrows its jurisdictional coverage and limits applicability across state-specific laws and decisions from lower courts. While the English–Kannada bilingual functionality enhances accessibility, a substantial percentage of India's population remains underserved due to the absence of additional regional language support. Moreover, highly detailed or layered legal queries can occasionally surpass the system's natural language comprehension capabilities, and its predictive components are inherently constrained by the subjective and interpretive nature of judicial reasoning. These challenges underscore the need for ongoing dataset expansion, model improvement, and greater linguistic diversity.

### **6.7. Future Scope**

Planned developments for LawHaven involve broadening the underlying legal repository to include High Court rulings, tribunal decisions, and specialized legislations, thereby enriching jurisdictional precision and coverage. Expanding multilingual capabilities including Hindi, Tamil, Telugu, Bengali, and Marathi will significantly broaden national accessibility. Further integration with e-court systems could enable real-time case status access, automated legal document retrieval, and improved metadata for legal analytics. From a technical standpoint, adopting more advanced transformer models and graph-based reasoning systems could strengthen precedent mapping and provide deeper legal context, while explainable AI frameworks would reinforce transparency and trust. Additionally, LawHaven may evolve as a comprehensive instructional resource for law students and a legal literacy tool for the general public, supporting wider access to legal knowledge.

## 7. CONCLUSION

LawHaven illustrates the meaningful role artificial intelligence can play in enhancing legal information systems and improving access, efficiency, and comprehension within India's diverse judicial environment. Through its structured legal knowledge repository and Retrieval-Augmented Generation framework, the platform delivers well-grounded, context-aware guidance that simplifies complex legal material for both experts and non-specialists. Its core functions including case summarization, similarity-based retrieval, bilingual translation, and predictive analytics collectively optimize processes that traditionally demand substantial expertise, manual effort, and time. Findings suggest that LawHaven can offer practical value to legal practitioners as well as the public by providing rapid, dependable insights and improving navigation through legal procedures. Although the system demonstrates strong capability in several areas, constraints such as limited jurisdictional scope, partial language support, and the challenges associated with legal prediction highlight the need for further refinement. Even so, the existing framework establishes a solid base for scalable, ethically aligned AI integration in the legal sector. With future advancements including broader datasets, enhanced multilingual support, deeper institutional integration, and more interpretable models LawHaven could mature into a fully developed legal intelligence platform. Ultimately, this work emphasizes the transformative potential of AI in improving access to justice and shaping future legal technologies.

## REFERENCES

1. Supreme Court of India, Supreme Court Reports: Fortnightly Digest, Govt. of India, New Delhi. Available: <https://main.sci.gov.in> (Accessed: 2025).
2. Government of India, the Indian Penal Code (IPC), 1860, Ministry of Law and Justice, New Delhi. Available: <https://www.indiacode.nic.in>
3. Government of India, the Constitution of India, Ministry of Law and Justice, New Delhi. Available: <https://legislative.gov.in>
4. A.Vaswani et al., "Attention Is All You Need," Advances in Neural Information Processing Systems (NeurIPS), pp. 5998–6008, 2017.
5. P.Lewis, E. Perez, A. Karpukhin et al., "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," Advances in Neural Information Processing Systems (NeurIPS), 2020.
6. J.Johnson, "FAISS: A Library for Efficient Similarity Search and Clustering of Dense Vectors," Facebook AI Research, 2017. Available: <https://faiss.ai>
7. LangChain, LangChain Documentation, 2024. Available: <https://python.langchain.com>
8. Groq Inc., Groq LLM Inference Documentation, 2024. Available: <https://groq.com>
9. MongoDB, MongoDB Technical Documentation, 2024. Available: <https://www.mongodb.com/docs>
10. FastAPI, FastAPI Framework Documentation, 2024. Available: <https://fastapi.tiangolo.com>
11. M.Devlin, M.Chang, K.Lee, and K.Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," Proc. NAACL, 2019.
12. T.Mikolov, K.Chen, G.Corrado, and J.Dean, "Efficient Estimation of Word Representations in Vector Space," Proc. ICLR, 2013.
13. S.Bird, E.Klein, and E.Loper, Natural Language Processing with Python, O'Reilly Media, 2009.
14. Ministry of Electronics and Information Technology (MeitY), e-Courts Mission Mode Project, Govt. of India. Available: <https://ecourts.gov.in>
15. D.Hendrycks, C.Burns, S.Basart et al., "Measuring Massive Multitask Language Understanding," Proc. ICLR, 2021.
16. R.Chalkidis, I.Androutsopoulos, and N.Aletras, "Legal-BERT: The Muppets Straight Out of Law School," Findings of EMNLP, pp. 2898–2904, 2020.
17. M.Zhong et al., "Extractive Summarization of Long Legal Documents with Hierarchical Transformers," ACL Workshop on Legal NLP, 2020.
18. A.Katz, D.Bommarito, and M.Katz, "Predicting Litigation Outcomes Using Machine Learning," Stanford Journal of Legal Studies, vol. 15, no. 1, pp. 1–20, 2017.
19. E.Ashley, Artificial Intelligence and Legal Analytics, Cambridge University Press, 2019.
20. M.Aletras, D.Tsarapatsanis, D.Preotiuc-Pietro, and V. Lampsos, "Predicting Judicial Decisions of the European Court of Human Rights," PeerJ Computer Science, 2016.
21. A.Chalkidis et al., "Paragraph-level Rationale Extraction through Reinforcement Learning from Human Feedback," EMNLP, 2021.