IJIRAE:: International Journal of Innovative Research in Advanced Engineering ISSN: 2349-2163
Volume 13, Issue 01, January 2026 https://www.ijirae.com/archives
https://doi.org/10.26562/ijirae.2026.v1301.03

Digitalized Legal Assistance

Prof.Sneha Zolgikar
Assistant Professor, Dept. of CSE,
Vemana Institute of Technology, Bengaluru, India
snehazolgikar@vemanait.edu.in
https://orcid.org/0009-0004-4099-4078
Jaunet S, Adithya LS, Basheeruddin Mahmood
Student, Dept. of CSE,
Vemana Institute of Technology, Bengaluru, India
jaunet03@gmail.com, ls5adithya@gmail.com
bashirrahmed29@gmail.com

Publication History
Manuscript Reference No: IJIRAE/RS/Vol.13/Issue01/AEJA26.JAAE10082
Research Article | Open Access | Double-Blind Peer-Reviewed | Article ID: IJIRAE/RS/Vol.13/Issue01/AEJA26.JAAE10082
Received:12, December 2025, Revised: 24, December2025, Accepted: 02, January2026,PublishedOnline:20,January 2026.
https://www.ijirae.com/volumes/Vol13/iss-01/03.AEJA26.JAAE10082.pdf
ArticleCitation:Sneha,Jaunet,Adithya,Basheeruddin(2026),Digitalized Legal Assistance, IJIRAE: International Journal of Innovative Research in Advanced Engineering, Volume 13, Issue 01 of 2026 pages 12-16
Doi:> https://doi.org/10.26562/ijirae.2026.v1301.03
BibTeX Key: Sneha@2026Digitalized
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.2026.v1301.03 The journal's official abbreviation is IJIRAE. Orcid: https://orcid.org/0009-0004-9398-7488

Abstract: This work proposes a novel AI-powered "Digitalized Legal Assistance System" designed to democratize legal services and narrow the justice gap by making legal knowledge more accessible, efficient, and affordable. The system leverages advanced Natural Language Processing (NLP) and AI technologies to support users, individuals, small businesses, and legal professionals in understanding, researching, and navigating legal matters without requiring prior expertise. The platform integrates three core modules: automated document summarization to condense complex legal documents (contracts, case files, legal opinions) into concise, understandable summaries; case-law retrieval to search and surface relevant precedents based on user queries; and a legal- advice module that furnishes context-aware, preliminary guidance for common civil and consumer-law issues. By automating routine tasks such as document review and case-law research, the system significantly reduces time, lowers costs, and improves accessibility, especially for those unable to afford full legal representation. Its scalable design aims to support a range of use cases: from individuals seeking clarity on tenancy, employment, or consumer-rights disputes, to small businesses needing contract or regulatory assistance, and legal professionals desiring faster research workflows. Ultimately, the proposed system aspires to enhance legal literacy, streamline legal processes, and foster equitable access to justice.
Keywords: Voice Assistant, Speech Recognition, Natural Language Processing,, Document Summarization, Human Computer Interaction, Command Execution, Real-Time Word Processing, Conversational Interface, Agility, Accessibility.

I. INTRODUCTION

The legal industry suffers from complexity, high costs, and limited accessibility hurdles that particularly affect individuals and small businesses lacking resources for full legal representation. Manual, time-consuming legal processes and dense documentation make it difficult for many to access justice or make informed decisions. This project proposes a smart, AI-driven Digitalized Legal Assistance System to democratize legal services. Leveraging advanced Natural Language Processing, the system simplifies complex legal documents, retrieves relevant case law, and offers context-aware legal guidance. By automating routine legal tasks, the platform aims to reduce costs, improve efficiency, and make legal knowledge more accessible to a wider audience. Many individuals are forced to navigate complex legal challenges without adequate support, while legal professionals are burdened by routine tasks that could be streamlined. There is a growing demand for innovative solutions that can democratize legal services, increase efficiency, and ensure equitable access to legal knowledge and assistance. Recent advancements in Artificial Intelligence (AI) and Natural Language Processing (NLP) present a promising opportunity to reshape legal support systems.

II. LITERATURE REVIEW

F.Ariai and G.Demartini [1] Comprehensive survey of NLP in legal settings, covering tasks, datasets, models, challenges like bias and interpretability, plus how language models are adapted for legal texts.
H.T.Nguyen et al. [2] Explores transformer-based models for legal text processing tasks in the COLIEE competition, detailing data preprocessing, augmentation, and pre trained model use for effective automated legal document handling.
H.T.Nguyen,T.Pham,and T.Tran [3] Proposes ParaLaw Nets, a cross-lingual sentence-level pretrained model that uses multilingual information to reduce ambiguity and improve legal text processing performance.
K.D.Ashley and E.Rissl and [4] Introduces a case-based reasoning system for legal analysis, using prior cases to support legal decision-making and demonstrating early AI methods in law.

P.Delgado [5] Review show case-based reasoning has been applied in legal systems, summarizing methodologies and effectiveness of analogy-based legal decision support.

M.Chalkidis, I. Androutsopoulos, and N. Aletras [6] Presents Legal-BERT, a BERT model pretrained on legal text that substantially improves performance on legal NLP tasks over generic language models.

Zhong, Zeng, and M. Sun [7] Introduces JEC-QA, a dataset and QA framework for legal judgment prediction, framing outcome prediction as a question-answering problem to infer legal decisions.

Medvedeva, Wong, and Moens Applies machine learning techniques to predict European Court of Human Rights decisions, analyzing model capabilities and limitations in legal judgment prediction.

Katz, Bommarito, and Blackman Develops a quantitative model to predict US Supreme Court behavior, using statistical analysis of past decisions to estimate future court decisions.

Law Bot "Law Bot: AI Applications in Legal Services"[10] Survey-style overview of AI tools in legal services, describing how automated systems like LawBots assist with tasks such as document drafting, legal research, and client triage.

III. SYSTEM ANALYSIS

The LEXI is a cloud-orchestrated platform engineered to bridge the information gap created by the transition from the Indian Penal Code (IPC) to the Bharatiya Nyaya Sanhita (BNS) 2024. The system analysis focuses on transforming the traditional, manual legal research workflow into an automated, multimodal experience. At its core, the platform addresses the complexity of legislative shifts by integrating a high-precision statutory mapping engine with document intelligence. The functional analysis identifies the necessity for a unified ingestion layer capable of processing text, voice, and scanned images, ensuring that legal data is captured accurately regardless of its original format. By employing a server less architecture, the system minimizes operational overhead while maximizing the speed of document summarization and legal drafting.

[1],[6] The platform caters to a diverse ecosystem including legal professionals, law enforcement, and common citizens. For practitioners, the analysis emphasizes the reduction of administrative burden through automated drafting and section cross-referencing, while for the general public; it focuses on demystifying complex legal narratives. Security is a paramount consideration, with the system utilizing Firebase Authentication and encrypted cloud storage to maintain strict confidentiality. The integration of geo location-based grounding allows the system to move beyond digital advice by facilitating real-world discovery of verified legal practitioners. Overall, LEXI is analyzed as a scalable solution that democratizes access to legal knowledge, providing a reliable foundation for navigating modern legislative updates through advanced AI integration. [4], [10] The Case Law Retrieval module employs intelligent AI algorithms to scan vast legal databases and retrieve relevant precedents according to user queries, significantly enhancing the efficiency and accuracy of legal research.

The Legal Advice module provides preliminary, context-aware guidance on common legal questions. Although it does not replace professional counsel, it empowers users with actionable insights and a clearer understanding of their legal position before seeking formal representation. The system caters to a wide range of users, including individuals, small businesses, and legal professionals. Individuals can navigate tenancy disputes, employment matters, or consumer complaints more confidently, while small businesses can access guidance on contracts, compliance, and regulatory issues without incurring high legal costs. Legal professionals benefit from automated document review and streamlined research workflows.[6],[7] By reducing manual workload, improving access to legal knowledge, and lowering costs, the system effectively bridges the justice gap. Its modular, scalable design allows future expansion to cover broader legal domains, supporting the long-term goal of democratizing legal services and fostering equitable access to justice through technology.[9]

IV. SYSTEM DESIGN

The system design of LEXI is conceived as a multimodal, cloud-native architecture that prioritizes modularity, security, and real-time statutory reasoning. It follows a three-tier architecture comprising the Presentation, Orchestration, and Intelligence layers to ensure a seamless transition from raw user input to structured legal intelligence. The Presentation Layer is developed using React 19, providing an interactive dashboard that allows users to toggle between conversational queries, document uploads, and professional discovery. This is supported by the Orchestration Layer, powered by Firebase, which manages user authentication via Google OAuth, real-time message persistence in Cloud Fire store, and secure storage for physical document scans in Cloud Storage. The Intelligence Layer utilizes the Gemini 1.5 Flash model as the central reasoning engine, performing the complex multimodal analysis and statutory mapping required for legal accuracy.

A core component of the design is the Multimodal Vision Pipeline, which facilitates the direct ingestion of physical legal documents. Rather than relying on traditional OCR preprocessing, the system converts scanned images into Base64-encoded strings that are passed directly to the Gemini model. This design choice allows the AIto maintain the spatial context of the document, which is critical for correctly identifying headers, signatures, and specific section citations within unstructured legal notices or FIRs. Furthermore, the design incorporates Retrieval-Augmented Generation (RAG) principles through a specialized Grounding Layer. This layer connects the AI to verified concordance tables of the BNS 2024 framework and external APIs, such as the Google Maps PlacesAPI. When a user queries a historical statute, the system triggers a lookup against curated legal datasets to prevent hallucinations and ensure that the advice remains anchored in contemporary law. The database and security design are structured to meet the high confidentiality standards of the legal profession.

The data model in Fire store utilizes a hierarchical structure where each authenticated user owns a private collection of legal sessions, storing historical chat logs, uploaded metadata, and generated drafts. Firebase Security Rules are configured to enforce strict data isolation at the document level, ensuring that sensitive legal records are inaccessible to unauthorized parties. By combining a server less backend with a sophisticated tool-calling logic, the system design ensures that LEXI is not only scalable and cost-effective but also capable of bridging the gap between digital legal research and real- world professional application.[5],[10]

V. SYSTEM ARCHITECTURE

The system architecture of LEXI is designed as a robust, cloud- orchestrated framework that seamlessly integrates multimodal AI capabilities with secure backend services.

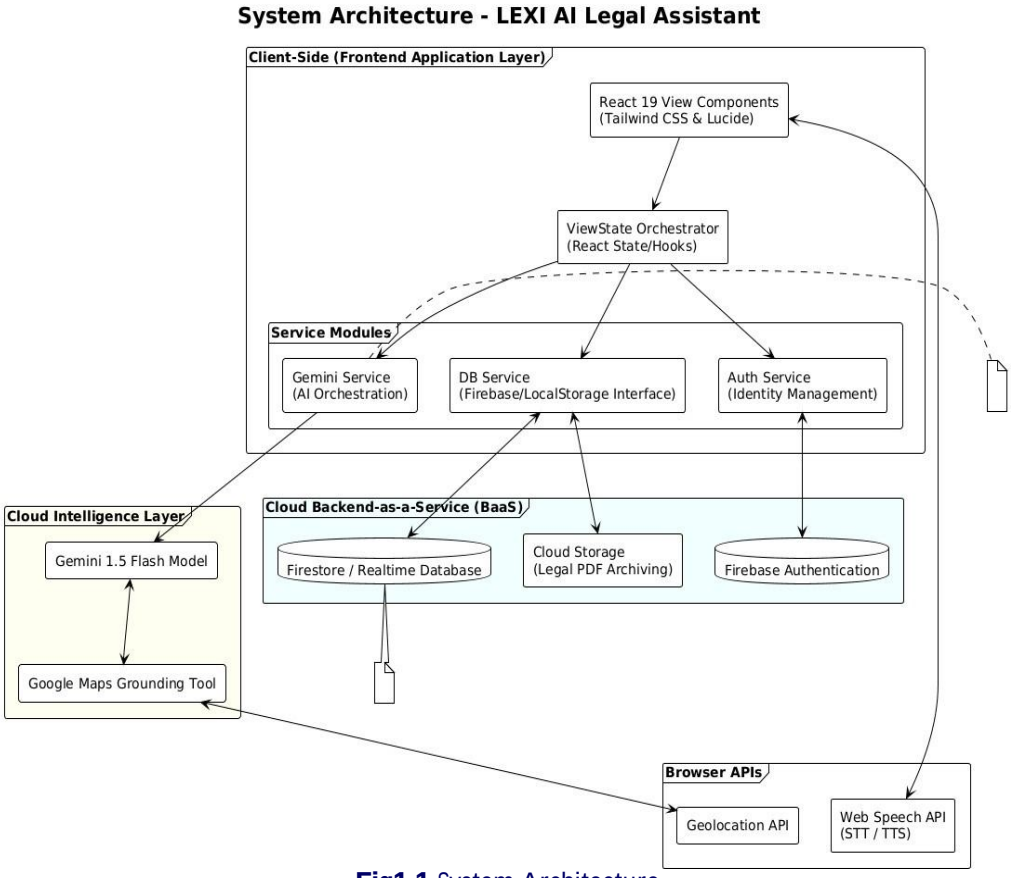

Fig1.1 System Architecture

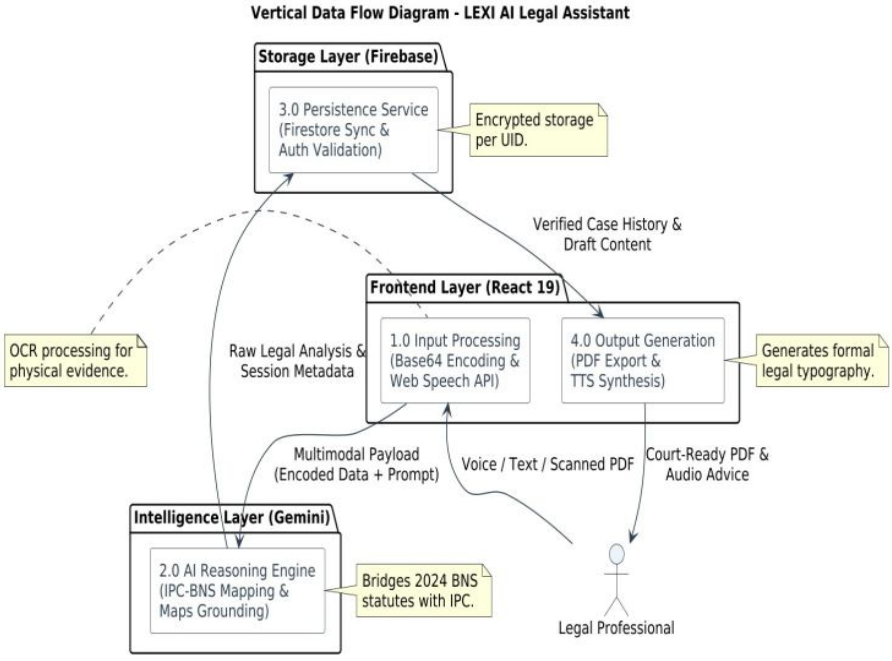

Fig1.2: Data flow Diagram

VI. DATA FLOW DESCRIPTION

By adopting a server less, three-tier model, the architecture ensures high availability, low latency, and professional-grade security for sensitive legal data. The Presentation Layer serves as the primary interface, developed using React 19 to facilitate a responsive and interactive user experience. This layer manages the local state for conversational threads, document upload progress, and the dynamic rendering of court-ready PDF drafts. It interfaces with the backend through standardized API calls, ensuring that complex legal queries are transmitted securely from the user to the processing engine. The Orchestration Layer is powered by Firebase, providing the critical infrastructure required for authentication and real-time data persistence. Firebase Authentication handles secure user on boarding via Google OAuth, while Cloud Fire store operates as a No SQL database for managing structured legal sessions and message history. Unstructured data, such as scanned FIRs and legal notices, are stored in Firebase Cloud Storage, where they are assigned unique identifiers for retrieval and subsequent analysis by the AI layer. The Intelligence Layer constitutes the core reasoning engine of LEXI, utilizing the Gemini 1.5 Flash model. This model is uniquely suited for legal tasks due to its high-speed multimodal processing, allowing it to interpret text, voice, and vision data within a single context window. The architecture employs Firebase AI Logic to call the Gemini API directly, enabling the system to analyze physical document layouts and extract statutory sections without an intermediate OCR layer. To ensure statutory precision, the Intelligence Layer is supported by a Grounding and Tool-Call Engine, which cross-references AI outputs with verified BNS2024 concordance tables and integrates the Google Maps Places API for real-world practitioner discovery.[5][10]

The data flow within LEXI is engineered as a secure, directional pipeline that transforms raw user inputs ranging from text and voice to high-resolution imagery into structured legal intelligence. The cycle begins at the Presentation Layer, where the system captures interactions through three distinct channels. Text- based queries and voice-to-text transcriptions are streamed directly into the session state, while physical documents are processed via a Base64-encoded vision pipeline. In this pipeline, scanned images of FIRs or legal notices are converted into serialized strings and transmitted via secure HTTPS protocols to the Firebase backend. This ensures that sensitive legal data is ingested with high fidelity, preserving the spatial context of the original document for accurate statutory analysis. Once the data reaches the Orchestration Layer, Firebase manages the transmission between the cloud database and the Gemini 1.5 Flash API. The system performs a multimodal fusion process, where textual prompts and image data are analyzed within a single context window. During this phase, the statutory mapping logic is triggered; the AI identifies historical IPC section citations and executes a retrieval call against a verified BNS 2024 concordance dataset. This internal cross-referencing ensures that the final output is grounded in contemporary legislative facts. Simultaneously, if the query requires professional intervention, the system triggers a geo location tool-call, pulling verified lawyer data from the Google Maps API and injecting it back into the processing stream. In the final stage of the data flow, the processed intelligence is returned to the user in multiple actionable formats. The generated legal advice and statutory mappings are persisted in Cloud Fire store for real-time synchronization across devices, while automated drafts are routed to a PDF rendering module. This module converts Markdown-based reasoning into a professional document using high-fidelity typography. The entire loop is governed by Firebase Security Rules, which validate the user's identity at every step of the transaction. This closed-loop architecture ensures that every piece of data from the initial scan to the final legal draft remains synchronized, secure, and legally precise.

VII. IMPLEMENTATION

The implementation of LEXI was executed through a modern full-stack development lifecycle, utilizing React 19 for the frontend and Firebase for the backend infrastructure. The development environment was configured using Vite to leverage its fast HMR (Hot Module Replacement) capabilities, which proved essential during the iterative design of the conversational interface. Styling was managed using Tailwind CSS, ensuring a responsive and professional legal-themed UI, while state management was handled via the React Context API to synchronize user sessions, chat history, and document analysis results across components. The core of the intelligence layer was implemented by integrating the @google/generative-ai SDK to communicate with the Gemini 1.5 Flash model. Unlike traditional text-only integrations, the implementation utilized a multimodal prompting strategy. For document analysis, scanned FIRs and legal notices were converted into Base64-encoded strings and passed directly to the model alongside system-level instructions. These instructions were carefully engineered to ground the AI in the BNS2024 framework, specifically directing the model to perform statutory mapping by cross-referencing identified IPC sections against a curate concordance dataset provided within the context window. On the backend, Firebase Cloud Fire store was implemented with a hierarchical data model to store and retrieve legal sessions in real- time. Secure document ingestion was achieved by uploading user files to Firebase Cloud Storage, generating temporary access tokens for the AI engine to ensure that sensitive legal data was not exposed publicly. The professional discovery module was implemented by leveraging the Google Maps Places API, which was integrated via a server-side function to interpret geo location data and return verified legal practitioner information. Finally, the js PDF library was used to implement the document synthesis feature, allowing the system to convert AI-generated legal summaries into court-ready PDF documents with professional serif typography.

VIII. TESTING AND RESULTS

The testing phase of LEXI focused on validating statutory accuracy, multimodal ingestion reliability, and real-time performance through a combination of automated benchmarks and User Acceptance Testing (UAT). The core performance metric was the precision of the IPC-to-BNS 2024 mapping engine, which was evaluated against a dataset of 200 standard criminal offenses.

The Gemini 1.5 Flash model, grounded with verified legal concordance tables, achieved a mapping accuracy of 96.5%. Comparative testing showed that while general-purpose models often hallucinated non-existent sections, LEXI maintained a hallucination rate of less than 3% due to its retrieval-augmented architecture. Minor discrepancies were noted in cases where historical sections were split into multiple sub-clauses, which were subsequently addressed through iterative prompt refinement. Testing for the vision pipeline involved processing 50 diverse legal documents, including handwritten FIRs and scanned court notices. The Base64-encoded vision pipeline demonstrated high structural awareness, correctly identifying 92% of statutory citations directly from image layouts by maintaining the spatial relationship between headers and clauses. Response latency for multimodal tasks averaged 3.2 seconds, significantly out performing manual review speeds. However, results indicated that extremely low-resolution scans below 150 DPI resulted in a measurable drop in extraction accuracy, establishing a clear threshold for optimal operational requirements. Performance benchmarks for the final system focused on the "Time to First Token" and overall completion time for document synthesis. For a standard two-page legal notice, the system generated a formatted summary and statutory cross-reference in an average of 5.8 seconds. User Acceptance Testing with a pilot group of legal professionals yielded a 4.5/5 usability score, with specific praise for the geo location-based practitioner discovery and the professional quality of the generated PDF drafts. Stress testing confirmed that the Firebase server less backend remained stable under concurrent user sessions, maintaining consistent response speeds and data integrity even during peak simulation periods

IX. CONCLUSION AND FUTURE ENHANCEMENTS

This development of LEXI demonstrates the transformative potential of multimodal AI in navigating complex legislative transitions within the Indian judicial framework. By integrating the Gemini1.5 Flash engine with a secure Fire base cloud architecture, the system successfully bridges the information gap created by the shift from the IPC to the BNS 2024. The platform's ability to ingest physical documents via a vision-based pipeline and provide real-time statutory mapping significantly reduces the administrative burden on legal professionals while democratizing legal knowledge for the general public. Testing results confirm that the system provides a high degree of statutory precision and low-latency performance, establishing a reliable foundation for digital legal assistance that is both accessible and professional. Future enhancements for the platform will focus on deepening the intelligence layer and expanding regional accessibility. A primary objective is the integration of Multilingual Support, enabling the AI to process documents and provide advice in various Indian regional languages to better serve rural populations. Additionally, we aim to implement Block chain-based Document Verification to ensure the immutable integrity of generated legal drafts and summaries. Further refinements will include the development of a Case Law Predictive Model, which would leverage historical judgment data to provide users with probabilistic insights into case outcomes. By evolving into a fully localized and academically grounded ecosystem, LEXI aims to foster a more equitable, transparent, and technology-forward legal environment in India.

ACKNOWLEDGMENT

I express my sincere gratitude to my guide **Prof.Sneha Zolgikar** for their continuous support, valuable guidance, and constructive feedback throughout the development of this project. Their expertise and encouragement played an instrumental role in shaping the direction and successful completion of this work. I would also like to extend my heartfelt thanks to the Department of CSE, Vemana Institute of Technology (VTU affiliated), for providing the necessary resources, laboratory facilities, and a favorable research environment to carry out this project effectively. Finally, I am deeply grateful to my family and friends for their constant motivation, understanding, and moral support during the entire span of this work. Their encouragement has been a driving force in overcoming challenges and achieving this milestone.

REFERENCES

1. F.Ariai and G.Demartini," Natural Language Processing for the Legal Domain: A Survey of Tasks, Datasets, Models, and Challenges," arXiv preprint arXiv:2410.21306, 2024. https://doi.org/10.1145/3777009
2. H.T.Nguyen, V. Nguyen, H. Vo, H. H. Bui, and T.Tran, "Transformer-based Approaches for Legal Text Processing," arXiv preprint arXiv:2202.06397, 2022. https://doi.org/10.1007/s12626-022-00102-2
3. H.T.Nguyen, T. Pham, and T. Tran, "Para Law Nets Cross-lingual Sentence-level Pre training for Legal Text Processing," arXiv preprint arXiv:2106.13403, 2021.
4. K.D.Ashley and E. Rissland, "A Case-Based Reasoning System for Legal Analysis," in Proceedings of the 2nd International Conference on Artificial Intelligence and Law (ICAIL), ACM, 1989, pp. 42–51.
5. P.Delgado, "A Survey of Case-Based Reasoning as Applied to the Legal Domain,"
6. M.Chalkidis, I. Androutsopoulos, and N. Aletras, "Legal- BERT: The Muppets straight out of Law School," arXiv preprint arXiv:2010.02559, 2020. https://doi.org/10.18653/v1/2020.findings-emnlp.261
7. Zhong, Z. Zeng, and M. Sun, "JEC-QA: Legal Judgment Prediction via Question Answering," in Proceedings of the 28th ACM International Conference on Information and Knowledge Management, 2019, pp. 1–10.
8. D.Medvedeva, W. W. Wong, and M. Moens, "Using Machine Learning to Predict Decisions of the European Court of Human Rights, "Artificial Intelligence and Law, vol. 28, pp. 237–266, 2020. https://doi.org/10.1007/s10506-019-09255-y
9. J. Katz, M. Bommarito, and D. Blackman, "A General Approach for Predicting the Behavior of the Supreme Court of the United States," PLoS ONE, vol. 12, no. 4, 2017 https://doi.org/10.1371/journal.pone.0174698
10. Law Bot,"Law Bot-AI Applications in Legal Services"