AI Based Legal Decision Support System By Using Faiss for Semantic Retrieval and Explainability

[1] Dr M Padma Priya, [2] Akil Pranav T, [3]Akshaya V, [4] Aslin R, [5] Ayush Raj S

[1]Assistant Professor, [2] Student, [3] Student, [4] Student, [5] Student,
[1] Department of Computer Science and Engineering,
[1] Bharath Institute of Science and Technology, Chennai, India
[1]padmapriya.cse@bharathuniv.ac.in, [2]akilpranav444@gmail.com, [3]akshaya492004@gmail.com,
[4] aslinr1234@gmail.com, [5] ayush7.ar11@gmail.com

***Abstract*— Legal documents can be really tough to figure out. People usually do not know where to find explanations. This project is about creating a Personal Legal AI Assistant that gives people answers about the law in a way that is easy to understand. The system uses natural language processing and a special search tool called FAISS from Facebook to look for ideas in a collection of legal information. This helps the system find the parts of the law, examples from court cases and simple explanations that people can understand when they ask a question about the law. The interface is easy to use so people who are not lawyers can get answers to their questions about the law without needing a professional to interpret the law for them. The Personal Legal AI Assistant is what this project is about. This approach is a way to make legal information more accessible and help people understand their rights and obligations better. The Personal Legal AI Assistant is about making the law easier to understand for people. The Personal Legal AI Assistant can really make a difference in peoples lives by providing legal information, about the law.**

***Index Terms*— FAISS, Vectorized legal knowledge base,semantic-similarity searches,LLMs,Transformer- Based Embeddings.**

I. INTRODUCTION

Today's legal practice involves managing growing volumes of digital documents, navigating complex regulations, and meeting increasing demands for timely and clear guidance. Generative AI has the potential to simplify routine tasks such as reviewing documents, summarizing content, assembling contracts, and preparing presentations which helps regular people and professionals who require some aid or assistance with their legal work. However, relying on general-purpose language models for critical legal contexts can be risky due to errors, outdated information, and lack of alignment with current laws.

To help with such problems, we have developed our project the Legal AI Project, a software platform built to provide fast, accurate, and context-aware legal information in a reliable and user-friendly way. The platform combines natural language processing with FAISS-based semantic search applied to a vectorized knowledge base, including various regulations, case law, and sample contracts. User queries are transformed into transformer-based embeddings to retrieve the most relevant clauses and references, which are then explained in plain language. All outputs are explicitly linked to source documents to ensure verifiability, and the system emphasizes transparency in reasoning wherever possible.

II. LITERATURE SURVEY

Wherever Times is specified, Times Roman or Times New Roman may be used. If neither is available on your word processor, please use the font closest in appearance to Times. Avoid using bit-mapped fonts. True Type 1 or Open Type fonts are required. Please embed all fonts, in particular symbol fonts, as well, for math, etc.

We are using Artificial Intelligence for transforming legal work enabling faster research,smarter document handling, and more efficient drafting and frameworking of legal documents for our users.This paper introduces Legal AI project, to be a software platform designed to deliver fast,accurate, correct and context aware legal information in an easy to understand and useful manner.

The platform is built to support common people, lawyers, researchers, and students with reliable legal knowledge for free of cost. The core idea of this project is to combine high quality NLP[natural language processing] with FAISS-based User questions are transformed into transformer-Based embeddings to enable precise retrieval of relevant references and clauses followed by a simple explanation understandable to normal non-expert users.The outputs are designed to be checked and verifiable, with transparency about all information.Privacy is maintained at all stages.

Design goals and scope

Grounded accuracy: ensure outputs can be traced to and verified against primary legal sources.
Contextual relevance: retrieve and present information that matches jurisdiction, topic, and user intent.
Privacy and security: maximize on-device processing and minimize reliance on external data sources.
Usability: deliver clear explanations in plain language suitable for non-experts while remaining rigorous for professionals.
Extensibility: support easy addition of new jurisdictions, domains, and document types.

Technical approach (high level)

Knowledge base: curate a structured collection of statutes, regulations, case law, and contract templates to serve as authoritative grounding material.

Embedding and retrieval: convert queries and document fragments into transformer-based embeddings; index with FAISS for fast, scalable semantic search.

Grounding and explanation: retrieve relevant passages and generate concise explanations that refer back to the source material.

Privacy-preserving processing: enable on-device or restricted-access processing pathways to reduce data sharing with external services.

Evaluation framework: measure grounding quality, retrieval accuracy, and user comprehension through both quantitative metrics and user studies.

Contributions

A grounding framework enabled legal tasks using retrieval-augmented generation over a legal knowledge base.

A privacy-conscious, on-device processing pipeline that minimizes external data use and exposure.

A modular, extensible architecture leveraging transformer embeddings and FAISS for precise retrieval and clear, source-grounded explanations.

An evaluation plan with representative benchmarks and potential user studies to assess accuracy, grounding, and user understanding.

Evaluation and validation (brief overview)

Quantitative metrics: retrieval precision/recall, grounding accuracy, and standard NLP metrics (e.g., BLEU/ROUGE for explanations where appropriate).

Grounding quality: alignment between generated explanations and cited sources.

User studies: assess comprehensibility, trust, and perceived usefulness among target users (students, practitioners, and educators).

Cross-jurisdictional evaluation: test performance across multiple legal domains to identify gaps and guide extension.

III. METHODOLOGY

A. Existing System

The digital legal documents are growing fast and this has created a big need for smart systems that can look at and find legal information quickly. Usually lawyers and researchers have to search through a lot of statutes, case law and regulatory documents. This takes a lot of time. You need to know a lot about legal words and what they mean. To make things easier researchers have tried to add intelligence to legal information systems.

One of the ways that people tried to use technology in law was with Legal Information Retrieval systems. These systems are made to find the important legal documents when you ask a question. They use techniques to look at legal texts and rank documents based on how relevant they are to a legal problem. Researchers have put these systems into groups like keyword-based systems, ontology-based systems and systems that use machine learning. New systems are using learning to look at unstructured legal texts and get useful information from them.

As artificial intelligence gets better, Natural Language Processing is becoming a part of looking at legal documents. This technology lets computers understand and analyze language so legal systems can summarize case documents, find important statutes and categorize legal opinions on their own. By using machine learning these tools can look at a lot of text and find patterns that would be hard to find by hand. They can do things like summarize cases, classify text and analyze precedents.

Another important development in technology is the use of semantic search systems. These systems are different from search engines because they look at the meaning and context of what you are searching for. They turn documents and searches into numbers, which lets them compare how similar things are based on what they mean not the words used. This makes legal research more accurate by finding documents that're related to what you are searching for even if they do not use the same words.

People have also been looking at using language models and special frameworks in legal AI systems. In these systems a part of the system finds legal documents first and then a language model uses those documents to generate answers. This helps make sure the answers are based on legal documents. Some systems use databases and models to make it easier to find legal documents and case precedents.

In addition to general language processing there are also models made just for legal language. Legal documents can be hard to understand because they have sentences, complex words and references to other statutes and decisions. To deal with this researchers have made models that can handle documents and do things like predict judgments, answer legal questions and find similar cases. For example some models have been really good at understanding and finding documents.

Recently people have also been studying ways to make AI systems that can handle really big datasets. These systems use machine learning and retrieval algorithms to analyze legal documents and case law. Tests have shown that these models can make classification and document retrieval more accurate, which could help legal professionals make better decisions and do research.

However there are still some problems. Many systems rely much on general language models that can make mistakes or say things that are not supported by law. Also legal reasoning can be complex. Requires looking at many documents, which can be hard for systems to do. Some researchers think that combining retrieval systems with knowledge could make legal AI systems more reliable and trustworthy.

Overall the work that has been done shows that artificial intelligence could change the way we do research and analyze documents.. There are still challenges, like making sure systems are reliable, transparent and protect data privacy, which means we need to keep working on making legal AI platforms better.

B. Disadvantages

Despite the advances in AI-driven legal research existing systems still have many challenges that limit how well they work.

First, many systems rely heavily on general language models, which are not specifically trained on texts.

This can result in made-up information, outdated data or responses that are not based on laws, which can mislead users in high-stakes legal situations.Semantic understanding is often limited.

Keyword-based retrieval methods cannot fully capture the meaning and nuance of legal queries, especially those involving multi-clause contracts, cross-references in statutes or judicial reasoning.

Source verification and transparency are often lacking.

Many AI systems generate answers without linking them to the legal documents making it hard for users to trace the reasoning or verify the accuracy of the output.

Data privacy and ethical concerns arise because most commercial AI platforms process legal data in cloud environments.This can conflict with confidentiality requirements and data protection regulations.Scalability remains an issue.

As legal texts grow in size and complexity existing systems often struggle with retrieval from large datasets leading to

slower response times and reduced usability.Adaptability to jurisdictions and languages is limited.Many systems are designed for a legal system and cannot easily accommodate variations in statutes or regulations across regions.

Finally, explainability is minimal which reduces user trust for non-experts who need clear guidance for decision-making.They need to understand AI-driven research to make informed decisions.

The systems should provide reliable information to build trust.The AI-driven legal research systems need to improve in these areas to be more effective.

C. Proposed System

The Legal AI system is designed to help people get the information they need in a simple way. It uses natural language processing and semantic document retrieval to make this happen. The main goal of the Legal AI system is to help users understand legal ideas by finding the right legal documents and making them easy to grasp.

The Legal AI system starts by creating a collection of documents like laws, regulations and court cases. These documents are cleaned up. Turned into numerical codes using special models. This helps the Legal AI system understand how different legal concepts are related and makes it easier to compare user questions to the documents.

The documents are then stored in a database that lets the Legal AI system search for similar documents quickly. The FAISS indexing framework is used to store and find these documents efficiently. This means the Legal AI system can look through thousands of documents without taking too long.

When someone asks the Legal AI system a question it processes the question using natural language processing. The question is cleaned up. Broken down into smaller parts before being turned into a numerical code. This code is then compared to the codes of the stored documents to find the relevant ones.

After finding the documents the Legal AI system creates a simple explanation of the legal information. It does not just give the user the legal text but instead summarizes it and makes it easier to understand. This helps people who do not have legal training.

The Legal AI system is designed to give answers that are based on legal sources. Each answer includes references to the documents used to find the answer. This helps users trust the information and verify its accuracy. By linking answers to legal sources the Legal AI system reduces the risk of giving wrong information.

The Legal AI system also cares about user privacy. Can do some processing on the users device. Many other AI systems rely on cloud services, which can raise concerns about data security. The Legal AI system is designed to work on the users device, which reduces the need for external services and meets ethical standards in the legal field.

The Legal AI system is also designed to be flexible and allow for new features to be added in the future. Some possible additions include predicting case outcomes, supporting languages and integrating with government legal databases. These additions could make the Legal AI system even more useful in world legal situations.

In summary the Legal AI system uses search, special codes and explanations based on real sources to provide a trustworthy platform, for legal information. By addressing the limitations of AI-based legal systems the Legal AI system aims to make legal information more accessible, transparent and reliable. The Legal AI system is designed to help people understand the law and make it easier for them to get the information they need. The Legal AI system is a tool that can help people navigate the complex world of law.

D. Advantages of this System

The proposed Legal AI system addresses limitations found in existing legal AI solutions.It uses a curated knowledge base and transformer-based embeddings to ensure retrieved information is accurate, relevant and based on verified laws, case law and regulations.This design reduces errors or fake outputs common in AI systems.

Unlike keyword-based search the system uses semantic search to understand complex legal language, including contracts with multiple clauses cross-references and nuanced queries.As a result users get precise and meaningful results improving the relevance of retrieved legal information.

Another key advantage is that all outputs are linked to legal documents ensuring transparency and verifiability.Users can trace the AIs reasoning. Validate information independently enhancing trust and reliability.

The system also generates plain-language summaries of retrieved content making it accessible to non-experts while maintaining accuracy for practitioners.

Data privacy and compliance are central to the design.The system supports on-device processing for sensitive queries ensuring compliance with privacy regulations and legal ethics.It uses FAISS-based vector search to handle volumes

of documents efficiently ensuring fast retrieval.

The architecture is modular and adaptable allowing integration of texts from multiple jurisdictions and potential multilingual data.

It supports extensions like case outcome prediction, automated contract analysis or legal trend detection.The platform can be integrated into existing workflows and knowledge management systems.Even though AI has made legal research faster and easier, current systems still have problems.Many rely on AI models not trained on legal texts leading to mistakes, outdated information or answers not based on real laws.

Traditional searches struggle to understand questions, contracts or references between statutes.

Most systems don't show where information comes from making it hard to check answers.Privacy is a concern with cloud servers, which can be risky for information.

Our system tries to solve these problems.It uses a prepared database of legal documents and transformer-based AI models to ensure accurate relevant answers based on real laws.

Semantic search helps understand questions, contracts or cross-references.Every answer is linked back to legal documents for verification.

The system explains things in language making it easier for non-experts to understand.

Privacy is protected with processing on the users device.The system quickly searches through collections of legal texts and is flexible, supporting updates with laws from different regions or languages.

Overall the Legal AI system is safer, easier to use and more reliable, than current solutions.

IV. System Architecture

Figure 1 System Architecture

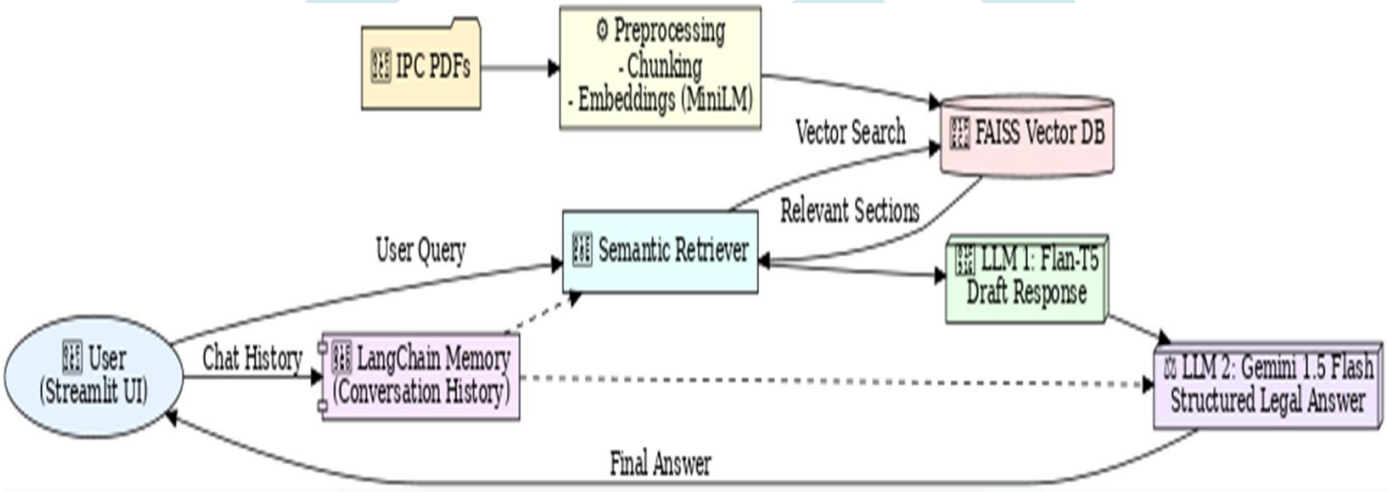

The process starts when the user sends a question through the chat interface. The question can be about an issue, a contract or a general question about laws or procedures. At this point the system just gets the text from the user. Does not do anything with it yet. This text is the beginning of the

process. The system then sends the question to the step where it is prepared for the system to understand and give a response.

Once the question gets to the step the system does a few things to it. It cleans up the text to remove any characters or symbols that could get in the way. This makes it easier for the system to understand. The system also breaks up legal documents into smaller parts. This is because legal documents can be very long and complicated so breaking them up makes it easier to find the information. Each part is then turned into a kind of code that the system can understand. This code captures the meaning of the text, not the words. All of these codes are stored in a database that allows the system to search for similar things very quickly.

After the system has prepared all the information it tries to find the parts of the legal documents that match the users question. When the user sends a question the system turns the question into the kind of code as the legal documents. This way the system can compare the question to the documents and find the parts that are most relevant. The system is not just looking for keywords. For the meaning and context of the question. This means it can find parts of the legal documents even if they do not use the exact same words as the question.

The system also keeps track of the conversation between the user and the system. This means it remembers what the user has asked before and can use that information to understand follow-up questions. For example if a user asks about a contract and then asks another question related to it the system can use the context of the question to understand the second one. This helps the system give accurate and relevant responses.

Once the system has found the parts of the legal documents it generates a response. A large language model reads the parts and turns them into a clear and concise explanation that the user can understand. The model tries to simplify the language while still keeping the original meaning. The goal is to give an answer that's both accurate and easy to understand.

After the system has generated a response it refines it to make sure it is clear and easy to read. The system also makes sure that the response is based on the documents it found earlier. This way the user can trust that the information is accurate. Comes from real legal sources.

Finally the system sends the response back to the user through the chat interface. The user gets a structured explanation of the legal information they asked about. The system combines its ability to understand the meaning of the question and remember the conversation. Summarize the

legal language to give a response that is both accurate and easy to understand. This way the user gets the guidance they need quickly and can also see the original legal sources if they want to.

V. Conclusion and Future Scope

Legal AI is a tool that helps people in India get access to knowledge. It uses something called Retrieval-Augmented Generation to understand what people are asking and to give them the answers.

This can be really helpful for people who need to find information. Sometimes it is hard to find the information and sometimes people get the wrong information. Legal AI can help with this. It can also help with the language that lawyers use. So Legal AI is not just for lawyers it can also be used by students and regular people who need help with things.

When we tested Legal AI it did a job. It was able to find the answers to people's questions more than 90% of the time. It was also fast and able to give answers. Sometimes people tried to trick it with information but it was still able to give good answers. We compared Legal AI to tools like it and it did better.

Legal AI is not perfect. Sometimes it has trouble with questions that involve countries or really long and complicated legal arguments. It also has trouble with specialized areas of law like international law. We know that we need to make it better so we are going to add training data and make it smarter.

We have plans to make Legal AI even better. We want to make it so that people can use it in languages, not just English. This will make it more helpful to people over India. We also want to make it so that it can give people information that's specific to where they live.. We want to make it better at understanding long and complicated legal arguments.

. We will make sure that peoples private information is protected and that we are clear about where our information comes from. We will make sure that we are fair to everyone. We will listen to what lawyers, teachers and regular people have to say so that we can make Legal AI the best it can be.

Legal AI is a step towards making legal information more accessible to everyone. We hope that it will help people in India get the help they need. We also hope that other people will work with us to make Legal AI even better and to use it in a way that's responsible and fair.

REFERENCES

List and number all bibliographical references in 10-point Times, single-spaced, at the end of your paper. When referenced in the text, enclose the citation number in square brackets, for example: [1]. Where appropriate, include the name(s) of editors of referenced books. The template will number citations consecutively within brackets [1]. The sentence punctuation follows the bracket [2]. Refer simply to the reference number, as in “[3]”—do not use “Ref. [3]” or “reference [3]”. Do not use reference citations as nouns of a sentence (e.g., not: “as the writer explains in [1]”).

Unless there are six authors or more give all authors’ names and do not use “et al.”. Papers that have not been published, even if they have been submitted for publication, should be cited as “unpublished” [4]. Papers that have been accepted for publication should be cited as “in press” [5]. Capitalize only the first word in a paper title, except for proper nouns and element symbols.

For papers published in translation journals, please give the English citation first, followed by the original foreign-language citation [6].

[1] Legal Information Retrieval systems: State-of-the-art and open issues, lCarlo Sansone,Giancarlo Sperlí- May 2022, 101967

[2] Natural Language Processing for Legal Document Analysis: Automating Judicial Insights Dr. Usman Khalid-Vol. 4 No. 2 (2024): Multidisciplinary Research in Computing Information Systems / Articles

[3] Revolutionizing Legal Research: How AI Saves Time and Enhances Efficiency,By Hanan Tazeem-2023

[4] AI Powered Legal Querying System using NLP Authors: Pushpa R N, Sanjana G Walke, Sharadhi D, Sharvari P K, Shreya C S, -December 2024

[5] Lawformer: A Pre-trained Language Model for Chinese Legal Long Documents, Chaojun Xiao, Xueyu Hu, Zhiyuan Liu, Cunchao Tu, Maosong Sun-May 2021

[6] Scaling Legal AI: Benchmarking Mamba and Transformers for Statutory Classification and Case Law Retrieval Anuraj Maurya-Aug 2025

[7] Domain-Partitioned Hybrid RAG for Legal Reasoning: Toward Modular and Explainable Legal AI for India Rakshita Goel, S Pranav Kumar, Anmol Agrawal, Divyan Poddar, Pratik Narang, Dhruv Kumar-,Dec 2025

[8] M. Bendre and V. Kulkarni, “AI-Driven Legal Assistance for Improved Constitutional Awareness,” in 2022 IEEE Int. Conf. on Smart Technologies, pp. 899–905, 2022.

[9] N. Srinivasan et al., “Enhancing Legal Literacy Through Conversational AI Systems,” IEEE Access, vol. 11, pp. 125460–125472, 2023. Government of India, Ministry of Law and Justice, The Constitution of India, Official Publication, 2023.