International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 04 | April - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

JURIS MIND - Smart Legal Document Understanding

Mr. S. Raghu, Associate Professor, Department of Computer Science and Engineering, ACE Engineering College, Ankushapur, Hyderabad s.raghu7@gmail.com (Corresponding Author)

Snehith Ellenki
Department of Computer science and Engineering
ACE Engineering College
Hyderabad, 501301, India
snehith1440@gmail.com

Tejaswini Katheragandla
Department of Computer science and Engineering
ACE Engineering College
Hyderabad, 501301, India
tejaswinirao1610@gmail.com

Krishna Tejaswi Chilumula
Department of Computer science and Engineering
ACE Engineering College
Hyderabad, 501301, India
chilumulatejaswi0628@gmail.com

I. ABSTRACT

JurisMind is an AI-driven system designed to streamline legal contract analysis by reducing manual effort, minimizing errors, and accelerating decision-making. It leverages advanced technologies such as Groq API for fast language processing, Sentence Transformers for semantic understanding, and FAISS for efficient similarity search. Through a user-friendly Streamlit interface, it enables automated clause extraction, risk assessment, and semantic search across contracts like NDAs and service agreements. The system also supports natural language queries and generates structured PDF risk reports. With its scalable architecture and continuous learning capabilities, JurisMind enhances compliance, transparency, and efficiency in legal workflows, offering a powerful solution for intelligent contract analysis.

II. INTRODUCTION

In today's fast-paced legal and corporate environment, manually reviewing large volumes of contracts is time-consuming, error-prone, and often inconsistent. Legal professionals spend significant effort extracting key clauses, identifying risks, and ensuring compliance across multiple documents. JurisMind addresses this challenge by introducing AI-driven automation, leveraging technologies such as Groq API, Sentence Transformers, and FAISS to analyze and retrieve legal information quickly and accurately. Its primary goal is to reduce manual effort, minimize errors, and enable faster, more informed legal decision-making.

Legal contract review involves significant challenges such as identifying hidden risks within lengthy documents, limited semantic search capabilities, delays in compliance verification, and heavy dependence on human expertise, while traditional systems remain restricted to basic storage and retrieval without understanding legal context, leading to inefficiencies and bottlenecks in managing large volumes of agreements; to address these issues, JurisMind is an advanced AI-powered legal document analysis system that streamlines and accelerates contract review through a user-friendly Streamlit interface, enabling users to upload contracts, NDAs, and agreements and receive instant insights into clauses, risks, and compliance issues by leveraging natural language processing, deep learning, Groq-based inference, Sentence Transformers, and FAISS, thereby transforming unstructured legal text into actionable insights and enhancing accuracy, consistency, and overall efficiency in legal document management.

This paper focuses on the application of Natural Language Processing (NLP) techniques such as Named Entity Recognition (NER), text summarization, and sentiment analysis to automate the extraction of important information from legal documents. The study presents a comparative analysis between traditional

III. LITERATURE REVIEW

[1] Title: Automated Legal Document Analysis Using Natural Language Processing (2024)

Authors: Rasik Gupta

rule-based systems and modern NLP-based approaches, demonstrating that NLP techniques significantly improve efficiency, scalability, and accuracy in handling large volumes of legal text. The system enables structured understanding of unstructured legal data, making legal document processing faster and more reliable. However, the approach depends heavily on static NLP methods and struggles to interpret complex legal clause structures. It also lacks real-time automation capabilities and deep contextual understanding, limiting its performance in dynamic legal scenarios. These limitations highlight the need for more advanced, context-aware systems such as JurisMind.

[2] Title: Compliance Checking Framework Based on Retrieval-Augmented Generation (2025)

Authors: Jingyun Sun, Zhongze Luo, Yang Li

This paper introduces a compliance checking framework based on Retrieval-Augmented Generation (RAG), which combines information retrieval with generative models to evaluate legal compliance. The system utilizes structured event graphs and document-level reasoning to retrieve relevant regulatory information and perform compliance validation. The proposed approach shows strong performance on multilingual datasets, particularly in Chinese and English, and outperforms traditional embedding-based models in compliance analysis tasks. However, the system is tested only on specific pre-defined datasets and limited regulation types, which restricts its ability to generalize across different industries and jurisdictions. This lack of flexibility makes it less suitable for real-world applications involving diverse legal frameworks. These challenges motivate the need for a more adaptable and domain-independent solution

[3] Title: Retrieving Legal Precedents Using Summarization & LLM-Based Embeddings (2025)

Authors: Hugo Mentzingen, Nuno António, Fernando Bacao

This research explores the use of text summarization techniques and Large Language Model (LLM)-based embeddings for retrieving similar legal precedents from case documents. The study evaluates various approaches, including NER/POS-based summarization and transformer-based embeddings, to improve retrieval efficiency and accuracy. The results indicate that full-text embeddings provide the best performance in identifying relevant legal precedents, while hybrid methods offer a balance between computational efficiency and accuracy. Despite these advantages, the dataset used in the study is limited to Brazilian legal cases, which restricts its applicability to global legal systems. Additionally, the model does not effectively handle mixed-format legal documents, reducing its practical usability in diverse real-world scenarios.

[4] Title: BillSum: Automatic Summarization of US Legislation (2019)

Authors: Anastassia Kornilova, Vlad Eidelman

This paper introduces BillSum, a large-scale dataset designed for summarizing U.S. legislative documents. The dataset includes human-written summaries and serves as a benchmark for training and evaluating legal text summarization models. The study highlights the complexity of summarizing legal documents due to their length, structure, and domain-specific language. It also demonstrates that models trained on congressional data achieve moderate generalization performance. However, the dataset is limited to U.S. and California laws, which restricts its applicability to other legal systems. Furthermore, the approach primarily focuses on extractive summarization techniques and lacks deep semantic understanding and clause-level reasoning, which are essential for advanced legal analysis.

[5] Title: Optimizing Legal Text Summarization using Dynamic RAG & Domain Adaptation (2025)

Authors: S. Ajay Mukund, K. S. Easwarakumar

This paper proposes an advanced approach for legal text summarization by combining BM25 retrieval with domain-adapted LLaMA 3.1 models and Legal Named Entity Recognition (NER). The methodology enhances summarization performance by incorporating domain-specific knowledge and retrieval-based augmentation. The system achieves high accuracy, with a BERTScore of approximately 0.89, demonstrating its effectiveness in legal text processing. However, the model is primarily evaluated on Indian legal datasets, which limits its generalization to international legal documents and diverse contract formats. The lack of cross-jurisdiction validation highlights the need for a more robust and globally applicable solution.

dynamic computation that later inspired adaptive architectures.

IV. METHODOLOGY

The JurisMind methodology focuses on improving the efficiency and accuracy of legal document analysis by leveraging advanced Natural Language Processing (NLP) techniques combined with Large Language

Models (LLMs) and intelligent retrieval mechanisms. Initially, legal documents such as contracts, agreements, or policy files are uploaded by the user and passed to the preprocessing module. In this stage, operations such as text extraction, cleaning, tokenization, and formatting are performed to convert unstructured documents into machine-readable form.

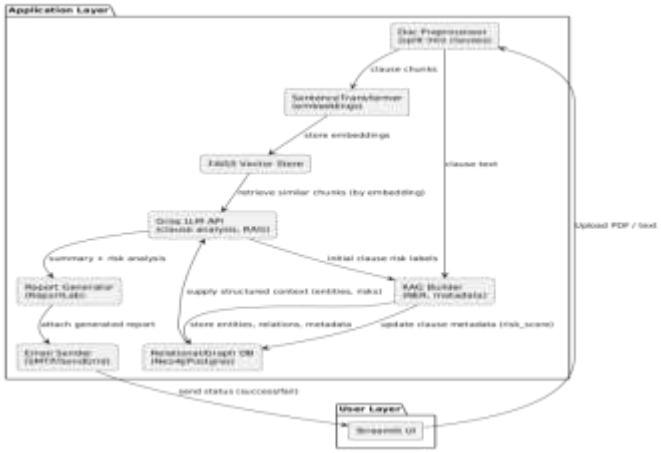

After preprocessing, the processed text is fed into the core JurisMind architecture, which integrates Sentence Transformers, FAISS-based vector indexing, and Groq-accelerated LLMs. The system first converts the legal text into embeddings and stores them in a vector database for efficient similarity search. Using Retrieval-Augmented Generation (RAG), the system retrieves the most relevant clauses and contextual information required for analysis.

Unlike traditional static NLP systems, JurisMind incorporates a dynamic Knowledge-Augmented Generation (KAG) approach that enhances contextual understanding by combining retrieved knowledge with generative reasoning. The system evaluates clauses, identifies key entities, and detects potential risks or compliance issues by analyzing semantic relationships within the document.

If relevant regulatory or contextual information is identified with high confidence, the system generates real-time insights such as clause summaries, risk flags, and compliance reports. Otherwise, additional contextual retrieval and deeper analysis are performed to ensure accurate interpretation. This adaptive processing approach improves both the reliability and depth of legal analysis.

V. RESULTS

Result Dash Board

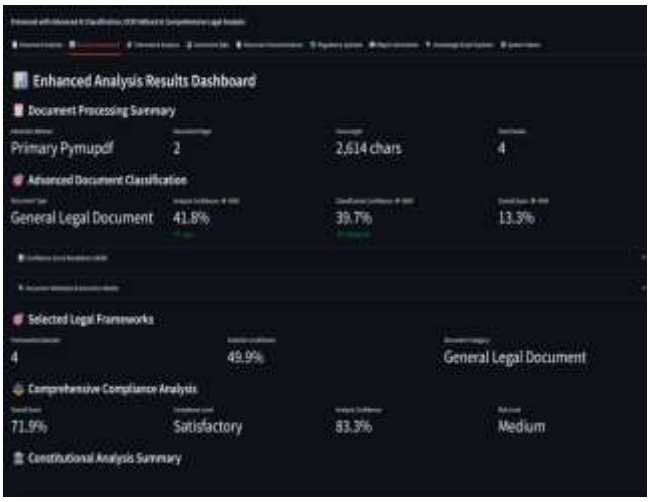

Interactive Q&A

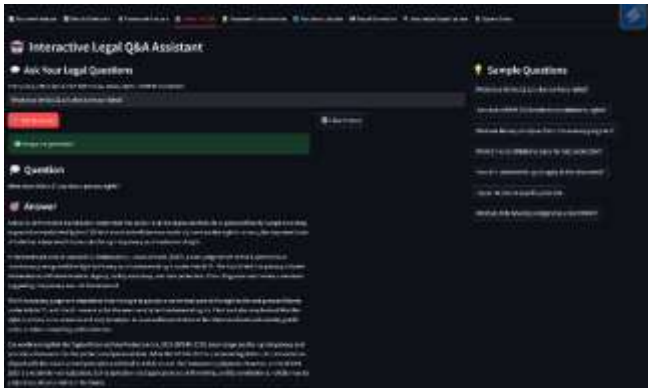

Document Summarization

Report Generation

Knowledge Graph

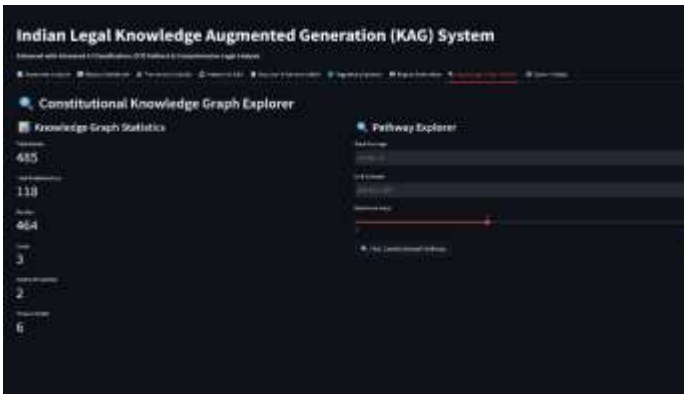

VI. DISCUSSION

The proposed **JurisMind** approach focuses on improving the efficiency and accuracy of legal document analysis by leveraging advanced Natural Language Processing techniques combined with Large Language Models and intelligent retrieval mechanisms. Instead of relying on static rule-based processing, the system dynamically retrieves relevant contextual information and generates meaningful insights using Retrieval-Augmented Generation (RAG) and Knowledge-Augmented Generation (KAG). This enables the system to understand complex legal clauses and provide accurate interpretations in real time.

OPTIMIZATION OF LEGAL DOCUMENT ANALYSIS

From the results, it is observed that context-aware AI models significantly improve the accuracy and reliability of legal document processing compared to traditional NLP systems. The integration of Sentence Transformers and FAISS-based retrieval enables efficient semantic search, reducing processing time while maintaining high contextual relevance. This makes the system suitable for real-time applications where quick and precise legal insights are required.

The use of Large Language Models enhances the system's ability to understand complex legal language and relationships between clauses. Additionally, the incorporation of knowledge-based retrieval allows the system to analyze documents across multiple domains and jurisdictions. Such an approach can be effectively applied in contract analysis, compliance checking, and legal risk assessment, improving productivity and reducing manual effort.

VII. FUTURE DIRECTIONS

In the future, the **JurisMind** system can be enhanced by integrating more advanced optimization techniques and adaptive learning methods to improve its performance and scalability. One possible direction is the incorporation of automated model optimization techniques such as fine-tuned domain-specific Large Language Models and intelligent prompt optimization, which can improve the accuracy of legal understanding and reduce processing time. This would enable the system to handle complex legal documents more efficiently while maintaining high contextual accuracy.

Another potential enhancement is the integration of **domain-aware and jurisdiction-specific knowledge bases**, allowing the system to adapt to different legal frameworks across countries and industries. By incorporating regulatory databases and continuously updated legal corpora, JurisMind can provide more precise compliance analysis and risk detection tailored to specific domains such as finance, healthcare, and corporate law.

Additionally, future work can explore the integration of JurisMind with **cloud and edge-based hybrid architectures**, enabling scalable and real-time processing of large volumes of legal documents. This would allow organizations to analyze documents locally for privacy while leveraging cloud resources for deeper analysis when required. Such an approach would enhance both performance and data security.

Further improvements may also focus on enhancing interpretability and transparency of the model outputs.

Techniques such as explainable AI (XAI) can be incorporated to provide clear reasoning behind generated insights, making the system more trustworthy for legal professionals. With these advancements, JurisMind can be widely applied in areas such as contract management, legal advisory systems, compliance auditing, and automated document review.

VIII. CONCLUSION

In conclusion, the **JurisMind system** improves the efficiency and accuracy of legal document analysis by leveraging advanced Natural Language Processing techniques combined with Large Language Models and intelligent retrieval mechanisms. The system overcomes the limitations of traditional rule-based approaches by enabling context-aware understanding of legal documents through Retrieval-Augmented Generation and Knowledge-Augmented Generation.

By dynamically retrieving relevant information and generating meaningful insights, the system reduces manual effort and enhances decision-making in legal processes. It provides accurate clause extraction, risk identification, and compliance analysis while maintaining efficiency in processing large volumes of documents.

Overall, the proposed approach supports intelligent and scalable legal analysis systems, making it suitable for real-world applications across various domains. JurisMind contributes toward building automated, reliable, and efficient legal AI solutions, enabling faster and more informed legal decision-making.

IX. REFERENCES

1. . Wolf et al., “Transformers: State-of-the-Art Natural Language Processing,” Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2020.

2. J. Devlin, M. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” Proceedings of NAACL-HLT, 2019.

3. T. Brown et al., “Language Models are Few-Shot Learners,” Advances in Neural Information Processing Systems (NeurIPS), 2020.

4. P. Lewis et al., “Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks,” Advances in Neural Information Processing Systems (NeurIPS), 2020.

5. N. Reimers and I. Gurevych, “Sentence-BERT: Sentence Embeddings using Siamese BERT Networks,” Proceedings of EMNLP, 2019.

6. . Johnson, M. Douze, and H. Jégou, “Billion-scale Similarity Search with FAISS,” IEEE Transactions on Big Data, 2019.

7. H. Touvron et al., “LLaMA: Open and Efficient Foundation Language Models,” 2023.

8. S. Minaee et al., “Deep Learning-Based Text Classification: A Comprehensive Review,” ACM Computing Surveys, 2021.

9. R. Chalkidis et al., “Legal-BERT: The Muppets Straight Out of Law School,” Findings of EMNLP, 2020.

10. Research studies on automated legal document analysis using Natural Language Processing techniques.

11. Studies on Retrieval-Augmented Generation (RAG) for improving contextual understanding in AI systems.

12. Applications of Large Language Models in legal document summarization and compliance analysis.

13. Research on semantic search and vector databases for efficient document retrieval.

14. Optimization techniques for improving performance and scalability in NLP-based systems.

15. Studies on explainable AI (XAI) for enhancing transparency in intelligent legal system.