International Journal of Computer Applications (0975 - 8887)
Volume 187 - No.23, July 2025

NyayaAI: Multi-Talented AI for Legal Aid, Smart Investing, and Scam Shielding

Swarupa Balaji
Department of AI & DS
MNM Jain Engineering College, Chennai

Saranya Maghendran
Department of AI & DS
MNM Jain Engineering College, Chennai

R. Ramya
Department of AI & DS
MNM Jain Engineering College, Chennai

ABSTRACT

The Legal and financial systems in India are often complex and inaccessible, especially for middle-class, economically weaker sections, and illiterate individuals. These groups struggle with common issues such as property fraud, contract disputes, tenancy conflicts, employment disagreements, and consumer rights violations due to limited legal awareness and the high cost of legal counsel. Financial challenges such as fraudulent investment schemes, lack of awareness about mutual funds and government schemes, fake job listings, and unverified companies further add to their vulnerability. NyayaAI addresses these gaps through an AI-powered legal and financial assistant that provides free, instant, and reliable guidance. Using advanced Natural Language Processing (NLP) and deep learning models, NyayaAI interprets user queries and helps to understand legal terms, as it undergoes personalized training in the Indian legal system. It also offers financial advisory support by suggesting suitable investment options, flags potential scams, verifying company authenticity, and identifying fake job postings.

Keywords

AI-Powered Assistant, Legal Aid, Financial Advisory, Natural Language Processing, Investment Fraud Detection, Government Schemes.

1. INTRODUCTION

1.1 Background and Motivation

India's legal system is one of the most comprehensive and intricate frameworks in the world, consisting of vast codes, complex statutes, and detailed procedural guidelines. While this complexity ensures thorough governance and regulation, it also makes navigating the legal landscape exceptionally challenging for ordinary citizens. Particularly, middle-class individuals, those from economically weaker sections, and illiterate populations often find themselves overwhelmed by legal terminologies and procedures that are difficult to comprehend. Property transactions, tenancy agreements, employment contracts, consumer rights issues, and other everyday legal matters become potential pitfalls, where misunderstandings or lack of awareness can lead to serious consequences.

The growing prevalence of fraud and exploitation further exacerbates this problem. In recent years, India has seen a significant rise in fraudulent activities involving real estate, deceptive employment schemes, financial investments, and online scams. Many individuals, unaware of the intricacies of legal documentation or lacking the means to verify authenticity, fall prey to well-crafted fraudulent contracts. These scams are often executed by fraudsters who exploit loopholes in existing laws and create documentation that appears legally sound on the surface. Victims only realize the deception after incurring significant financial or legal damage. The situation is further worsened by the prohibitive costs of legal counsel, which many from vulnerable populations cannot afford. As a result, the absence of accessible, reliable, and affordable legal assistance leaves millions defenseless against such exploitation.

Existing legal applications and technology platforms are often designed for legal professionals or prosecutors and are generally focused on global legal systems. While these platforms are undoubtedly powerful, they do not cater to the everyday legal needs of common citizens. The legal terminology used is often too technical, and the interfaces are not designed with accessibility in mind. This leaves a vast gap between legal knowledge and those who need it most. Individuals without a legal background continue to struggle with understanding their rights, interpreting legal clauses, and taking preventive steps against fraudulent schemes.

1.2 Introducing NyayaAI: AI-Powered Legal Assistant

Addressing the legal awareness gap requires a solution that is both technologically advanced and user-centric. NyayaAI is an innovative AI-powered legal assistant designed to provide free, instant, and reliable legal guidance tailored to the Indian legal system. Unlike existing tools that cater mainly to legal professionals, NyayaAI focuses on helping clients and laypersons.

It uses advanced Natural Language Processing (NLP) models trained on Indian legal documents, case histories, and procedures to deliver contextually accurate guidance. Through a user-friendly

chatbot interface, individuals can ask legal questions in natural language. NyayaAI interprets these queries, analyzes them against its legal database, and provides clear explanations, potential outcomes, and actionable advice. Key features also include legal document analysis, preliminary document drafting, and case outcome prediction using deep learning—helping users make informed decisions.

1.3 Finance Advisor and Fraud Detection Overview

To combat rising fraud in legal and financial matters, NyayaAI integrates a powerful Fraud Detector Module. It identifies scams related to investments, companies, and employment by analyzing suspicious terms, verifying entities through registries and trusted platforms, and flagging fake job ads with common scam traits. This helps users detect and avoid fraudulent schemes.

The Finance Advisor provides personalized guidance on mutual fund investments by assessing user preferences, risk levels, and market trends. It also informs users about government schemes, tax benefits, and financial rights and responsibilities. Together, these tools empower users to make safer, smarter legal and financial decisions with confidence and clarity.

Given the complexity of the Indian legal system and the rising incidence of fraud and exploitation, the demand for accessible, reliable, and context-specific legal assistance is greater than ever. NyayaAI addresses this need by offering:

—**AI-powered legal guidance** tailored specifically to the Indian legal framework, including IPC, CrPC, CPC, and other statutes.
—**Natural language processing (NLP)** to simplify complex legal terminology and provide clear, understandable advice.
—**Fraud detection capabilities** focused on investment schemes, fake companies, and fraudulent job postings.
—**Legal document analysis** and preliminary document generation to assist users with common legal paperwork.
—**Predictive analytics** to help users anticipate possible case outcomes based on deep learning models trained on Indian case histories.
—**User-friendly chatbot interface** that enables natural language queries for immediate legal assistance.

2. RELATED WORK

Recent advancements in artificial intelligence have led to the development of AI-powered tools for legal and financial assistance, focusing on enhancing accessibility, transparency, and efficiency. Various systems have employed Natural Language Processing (NLP), hybrid deep learning architectures, and decision-tree-based models to automate legal analysis, interpret legislation, and offer real-time legal support. The Analysis Table in Table 1. gives a clear overview.

2.1 Overview of AI in Legal and Financial

A prominent approach involves developing chatbots trained on comprehensive legal datasets, including statutes, case laws, and regulations. These systems use Natural Language Processing (NLP) for data pre-processing and structuring, enabling the chatbot to understand user queries and respond with relevant information. User intent recognition is central—distinguishing whether a user seeks legal advice, general knowledge, or professional referrals. These systems are refined through user feedback to ensure legal accuracy and technical robustness. LegalAsst, a full-process system, adopts a modular architecture covering case analysis, legislation analysis, judicial decision-making, and human oversight. It constructs structured case graphs from unstructured descriptions via information extraction, improving interpretability. Legislation graphs, aligned with specific case types like fraud or smuggling, guide decisions on penalties and sentencing. To address bias and ensure fairness, a human intervention module reviews outputs to maintain ethical and legal standards. Methodologically, legal AI research often employs systematic literature reviews using bibliometric techniques. Tools like Biblioshiny and VOSviewer support science mapping and visual network analysis, helping researchers identify gaps and emerging trends in legal technology.

Recently, there has been a growing interest in using graph-based approaches to uncover fraud patterns in transaction networks. Instead of treating each data point in isolation, this method models transactions as graph-where entities (like accounts or users) are represented as nodes, and interactions (such as money transfers) form the edges. To evaluate its effectiveness, this setup is compared against both classical machine learning models (such as logistic regression and random forests) and standard deep learning architectures (such as dense neural nets and autoencoders). What sets the graph-based approach apart is its ability to capture complex relational structures. Techniques like Graph Convolutional Networks (GCNs) and Graph Attention Networks (GATs) are central to this.

2.2 Explainable Methodologies for Decision Making

Recent advancements in legal NLP have emphasized the importance of domain-specific pre-training and sequential modeling for improved performance on legal tasks. InCaseLawBERT and CustomInLawBERT, which are variants adapted to the Indian legal domain of pre-trained language models (PLM), have recently been introduced. These models are fine-tuned or trained from scratch on a large Indian legal corpus and demonstrate significant improvements over the baseline in terms of perplexity and task-specific performance on legal statute identification, semantic segmentation, and court judgment prediction.

To ensure the legal advisory system operates with transparency and accountability, explainability is embedded at multiple stages of the pipeline — from document ingestion to user interaction. Given the complexity of legal language and the critical implications of legal guidance, it is imperative that the system not only delivers accurate responses but also provides interpretable reasoning behind its outputs. On the user-facing side, explainability is reinforced through structured dialogue flows and intent recognition. The chatbot does not respond with generic output; instead, it guides users through a logical sequence of questions tailored to their legal context. User intents — such as seeking advice, clarifications, or professional referrals — are classified through trained NLP models, allowing for relevant responses and transparent interaction trails. Graph-based models like Graph Convolutional Networks (GCNs) and Graph Attention Networks (GATs) are integrated for advanced scenarios such as fraud detection or legal network analysis. These models are chosen for their inherent interpretability in understanding node-level importance. GATs provide insight into which relationships or connections in a legal or financial graph were most influential in arriving at a decision. The use of attention mechanisms makes it possible to highlight key dependencies that shaped the output, contributing to trust in AI-powered legal tools.

In line with ethical standards, a disclaimer is presented to inform users that the system offers guidance based on structured legal data

Table 1. Analysis Table

Title	Description	Advantages	Tech Stack
LegalAsst: Human-centered and AI-empowered machine to enhance court productivity and legal assistance	This paper presents an AI-powered, yet human-centered system designed to enhance legal productivity by streamlining case analysis, legislation review, and judicial decision-making. It offers explainable and traceable support using structured representations and decision trees, enabling transparent and up-to-date legal judgments.	Combines modification capabilities, which help update the most up-to-date legislation and societal factors to generate more adaptable judgment outcomes.	SVD, Graph-based learning, Lawformer
Pre-trained Language Models for the Legal Domain: A Case Study on Indian Law	This study focuses on pre-training legal language models on Indian legal texts. The model is evaluated across multiple legal NLP tasks and demonstrates improved explainability through qualitative attention-based analysis.	Improved cross-domain performance (EU/UK texts), Custom vocabulary boosts relevance for local terminology	Pre-trained BERT models, KL Divergence
AI-Powered Fraud Detection in Financial Services: GNN, Compliance Challenges, and Risk Mitigation.	This study presents a Graph Neural Network (GNN) framework for fraud detection in finance. It uncovers complex fraud patterns while ensuring real-time feasibility in high-frequency financial environments.	GNNs better identify hidden, collusive, and synthetic fraud compared to conventional models. The framework supports AML/KYC compliance and is computationally efficient for real-time, high-volume financial data.	Graph neural networks, Explanable AI
Enlightening Justice: Empowering Society Through AI-driven Legal Assistance	This is an AI-powered legal aid system aimed at reducing legal ignorance and promoting justice for underserved populations. By combining NLP, it classifies legal documents and offers real-time, case-specific legal assistance with high accuracy and accessibility.	Equips people, especially the poor and illiterate—with essential legal knowledge and terminology. Outperforms traditional methods in legal document classification.	SpaCy, NLTK, Transformers
LEGALBOT - AI LAW ADVISOR CHATBOT	LEGALBOT is an AI-powered legal chatbot designed to democratize legal services by offering affordable and accurate legal information using advanced NLP.	Makes legal help affordable and widely accessible, Automates routine legal Q&A for lawyers and firms.	NLP, Large language models
Artificial intelligence in judicial adjudication: Semantic biasness classification and identification in legal judgement (SBCILJ)	This research addresses the issue of semantic bias in legal judgments, particularly within the Chinese Artificial Intelligence and Law (CAIL) dataset. The study aims to detect and classify semantic biases that may affect AI predictions in legal decisions.	Helps legal practitioners by providing AI-driven tools to identify and mitigate biases, promoting fairness and ethical AI usage	SVM, KNN, Multilayer Perceptron, Naïve-Bayes
A Chatbot Framework for the Children's Legal Centre	This chatbot framework designed to help children access legal rights information and create cases for legal advisors.	High classification accuracy (up to 98.24%), automatically generates structured case reports for advisors	Two-layer RNN, Softmax, F1 Score

and predefined scopes but does not replace certified legal advice. This transparency in system scope and limitation further supports responsible deployment and user trust.

2.3 Decision Paths and Legal Outcome Visibility

Following data acquisition and preprocessing, the text extraction and structural understanding phase plays a pivotal role in transforming complex PDF content into analyzable digital form. This process ensures that the semantic and syntactic structure of the documents is preserved, allowing downstream models to perform accurate information retrieval and embedding generation. The procedure begins with parsing the PDF files using robust Python libraries such as PyPDF2 and pdfminer. These tools facilitate granular extraction, handling the content on a per-page basis while preserving logical boundaries like paragraphs and sentences.

Special attention is given to documents with intricate formats, such as multi-column layouts, by leveraging spatial positioning data to differentiate between column breaks and line continuations. Regular expressions and pattern matching are used to delineate and classify textual elements into titles, subtitles, bullet points, and narrative content.

Legal decision-making tools must provide clarity on how specific outcomes are reached and which parts of the legal document influenced those outcomes. To address this need, recent research emphasizes model interpretability through attention-based traceability and human-model agreement evaluation. One effective methodology involves analyzing model-assigned attention scores to determine which portions of a legal document were most influential in the decision. For instance, in Court Judgment Prediction tasks, transformer-based models are fine-tuned to classify outcomes (e.g., favorable/unfavorable judgments). However, the output alone, typically a binary classification—is insufficient for real-world adoption unless the underlying rationale can be validated. To ensure traceability, legal documents are tokenized and divided into chunks, each of which receives an attention weight from the model's upper encoder layers (such as BiLSTM + Attention). These attention scores form a probability distribution over the document, effectively highlighting which sections the model focused on to arrive at its decision.

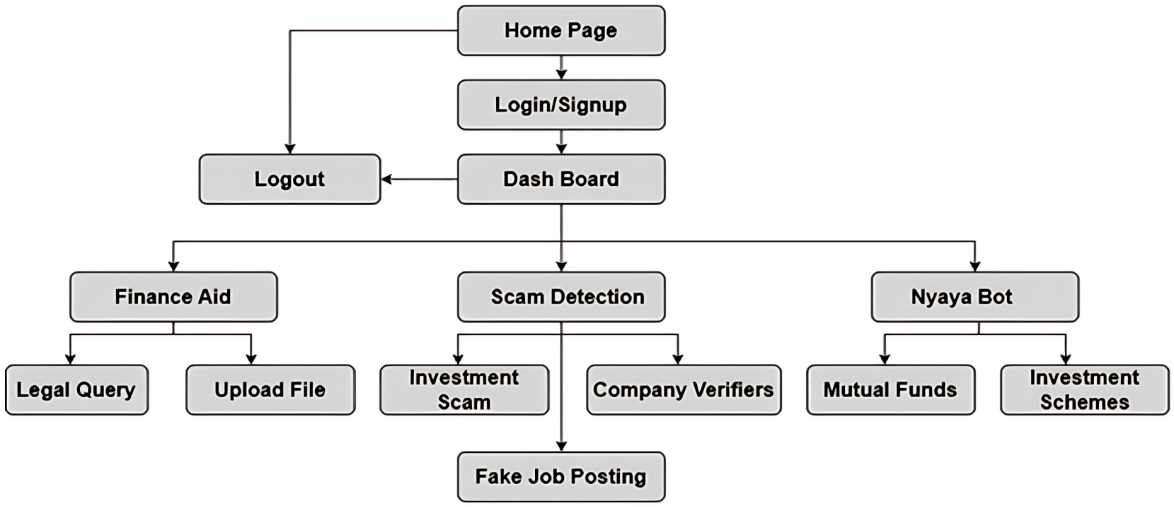

Fig. 1. Project Flowchart

2.4 Limitations, Ethical Concerns, and Future Scope

Despite notable advancements in legal and financial AI systems, several limitations and ethical considerations remain that must be addressed to ensure practical scalability, fairness, and reliability. A key technical challenge lies in the scalability of indexing and inference, especially when using high-dimensional vector search libraries like FAISS or complex models like Graph Neural Networks (GNNs). As dataset sizes grow into millions of document chunks or transactions, computational overhead and memory constraints increase substantially. Optimization techniques such as model distillation, quantization, and distributed graph processing are essential for improving efficiency in real-time applications. From an ethical standpoint, data privacy and security are paramount. Legal and financial documents often contain sensitive information, requiring strict compliance with data protection laws such as GDPR. Access control mechanisms, audit logs, and transparency about data handling practices must be built into these systems.

It is important to remember that chatbots have limitations and cannot replace the expertise of a qualified human lawyer. Additionally, the accuracy and reliability of a Legal Bot depend on the quality of its programming and the currency of its legal knowledge, which can become ineffective if not updated in a timely manner. Users should use these chatbots as a starting point for basic legal inquiries and seek professional legal counsel for complex or critical legal issues.

3. METHODOLOGY

NyayaAI is a comprehensive AI-powered legal assistance platform developed with the aim of enhancing access to legal information, promoting financial awareness, and mitigating fraud in digital interactions. At its core, the system integrates custom-trained language models on domain-specific data with intelligent pipelines that process user queries, legal documents, and financial inputs in a contextual and efficient manner.

The foundation of NyayaAI rests on a legal chatbot trained on Indian legislative texts, case law summaries, and procedural codes. The legal content is preprocessed and structured before being routed through a semantic transformation pipeline that enhances interpretability and relevance. This enables the chatbot to respond with clarity and legal precision while adapting to the user's query style. In parallel, NyayaAI includes a dedicated financial advisory interface—designed as an interactive dashboard—which provides personalized insights on mutual funds, statutory compliance, and investment patterns. It leverages a lighter language model optimized for financial datasets and is enhanced by a refinement pipeline that contextualizes and explains complex data. This results in a user-friendly experience where financial guidance is delivered in a conversational yet analytically robust manner.

A significant part of the system is its fraud detection module. This component scans a wide range of textual inputs, including job offers, corporate claims, and investment pitches—to detect inconsistencies and warning signs. It integrates classification layers and pattern recognition stages that flag potentially fraudulent elements by matching them against known scam patterns and behavioral outliers. This architecture empowers NyayaAI to assist users across legal, financial, and compliance domains with a high degree of reliability and relevance.

3.1 Data Collection

Data collection is the paramount step of the methodology, as the project calls for a variety of legal, financial, and scam-related data from multiple reliable sources. This diverse and domain-specific data ensures NyayaAI remains context-aware, accurate, and responsive across its three primary modules.

—**Legal Module:** The model is trained on several statutory texts including the IPC (Indian Penal Code), CrPC (Code of Criminal Procedure), IEA (Indian Evidence Act), and key parts of the Constitution of India. These are parsed into JSON format with fields such as Section Number, Title, Description, Exceptions, and Illustrations. The system is also enriched with definitions of legal terminologies and colloquial legal phrases to improve query understanding. Sources include government repositories, eSCR, and Indian Kanoon. Additionally, recent reforms under

Bharatiya Nyaya Sanhita (BNS) are integrated to enable pre-2023 and post-2023 case differentiation.

—**Document Processing Dataset:** Real-world legal documents such as rent agreements, offer letters, notices, and contracts are used to train the document summarization and QA modules. These are uploaded by early testers or curated from anonymized public datasets. OCR pipelines are tested using scanned PDFs of varying quality to make the system robust.

—**Finance Module:** JSON-formatted data of mutual funds, tax laws, investment types, SEBI regulations, and government schemes have been collected from MoneyControl, SEBI's portal, ClearTax, and the National Portal of India. This includes parameters like NAV, returns, fund type, risk levels, tax benefits, and redemption rules.

—**Fraud Detection Module:** This module is supported by multiple datasets:

—A Kaggle dataset for fake job postings, cleaned and enriched with features like exaggerated pay claims, work-from-home keywords, and missing company credentials.

—A custom-built investment scam dataset labeled for fraud/legitimacy, built using scraped schemes from Reddit forums, scam tracker sites, and financial blogs.

—Web-scraped data of company verification information via platforms like LinkedIn, Google Maps, MCA India, MSME directories, Zauba Corp, and domain WHOIS lookup. Parameters such as incorporation status, domain age, and business registration are used to identify red flags.

—**User Simulation Dataset (for Evaluation):** A synthetic dataset containing hundreds of sample user queries — legal, financial, and fraud-related — has been created to simulate real-world interaction. It helps test the chatbot's response relevance, latency, and accuracy.

—**Multilingual Data (Planned):** To enable future regional language support, scraped parallel corpora of legal content in Hindi, Tamil, and Telugu are being collected. This will be used for training multilingual NLP pipelines.

3.2 Data Preprocessing

For NyayaAI's preprocessing pipeline, legal documents such as IPC, CrPC, and Articles are first extracted from PDFs using parsers like PymuPDF and fitz, then cleaned by removing unnecessary whitespace, headers/footers, superscripted references, and noise using regex. The text is segmented into sections based on patterns like "Section 302:" and parsed into structured JSON containing the section number, title, description, exceptions, illustrations, and subsections. Simultaneously, financial data is scraped from trusted sources like SEBI, cleartax, and MoneyControl using BeautifulSoup. This data—covering mutual funds and potential scams—is cleaned by stripping HTML, normalizing formats, and filtering red flags like "guaranteed returns" or missing SEBI registrations. It is then structured into JSON with fields such as fund name, type, NAV, returns, risk level, strategy, and scam indicators. This combined structured preprocessing ensures NyayaAI can perform accurate legal analysis, fraud detection, and financial advising.

Using pandas and numpy, the fake job postings dataset was preprocessed to remove null values, duplicate rows, and irrelevant columns such as job ID or location metadata in the dataset. To enhance the model's understanding of fraud patterns, features such as presence of suspicious phrases ("work from home and earn $5000/week", "no experience needed"), excessive exclamation marks, and unrealistic benefits were flagged.

3.3 Model and System Design

3.3.1 Module I: NYAYA AI LEGAL AID CHATBOT

NyayaAI's legal chatbot is an advanced AI-powered system designed to provide instant, accurate, and context-aware legal assistance. As shown in Fig. 2, it uses a hybrid model architecture combining inLegalBERT, Mistral-7B, and Gemini 1.5 Flash to interpret queries, retrieve precise legal information, and summarize legal documents in real time.

—The core of NyayaAI's chatbot is a fine-tuned version of inLegalBERT, a legal language model trained extensively on JSON-structured data from the Indian Penal Code, the Criminal Procedure Code (CrPC), the Indian Evidence Act, and related statutes. This model is responsible for retrieving precise legal provisions,

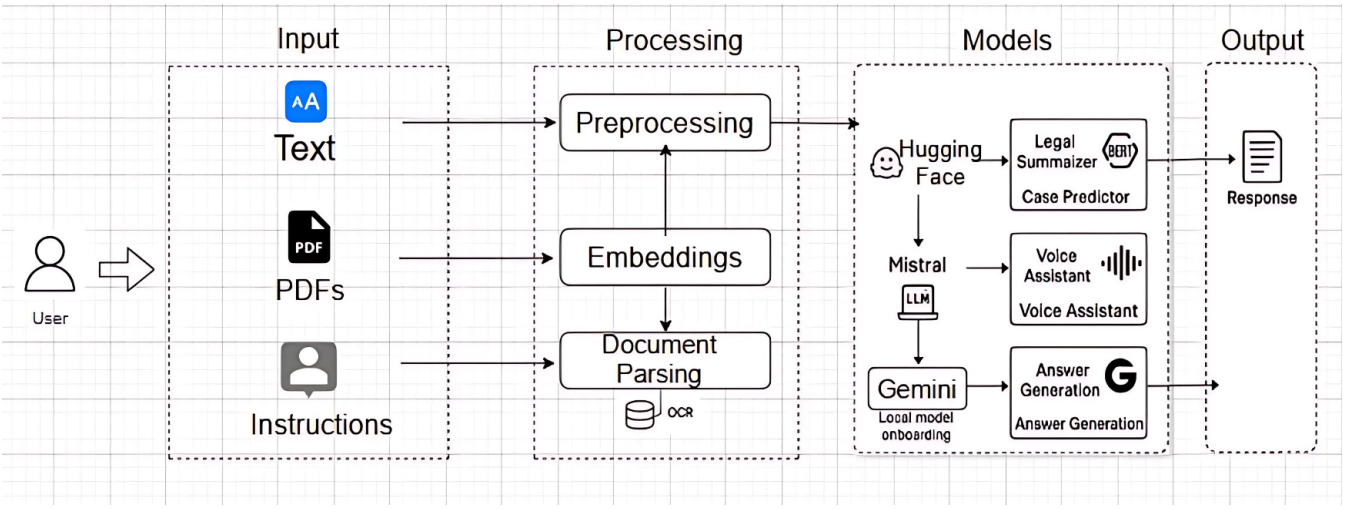

Fig. 2. Legal Aid Bot Architectural Diagram

Table 2. Data Collection

Dataset Name	Source	Description
IPC, CrPC, IEA, Legal Terminologies	Government of India, eSCR	Full articles, sections, subsections, exceptions, and illustrations structured in JSON form.
BNS (Bharatiya Nyaya Sanhita)	Govt. 2023 reforms	Used for classification of pre and post 2023 legal application.
Investment Scam Corpus for DistilBERT	Custom	Tokenized version of investments that are classified as fraud or legit based on the expected returns, description, duration etc.
Mutual Funds JSON (Concepts and Definitions) and Investment Schemes	MoneyControl, cleartax, National Portal of India	Includes types of mutual funds, fund names, FAQs etc.
Fake job posting	Kaggle	Used Machine Learning model XGBoost for fake job detection.
User Simulation Data	Custom synthetic generation	Query-response pairs to evaluate system behavior and benchmarking across modules.
Regional Language Data (Planned)	Wikipedia + Govt portals	Parallel corpora in Hindi, Tamil, Telugu for future multilingual training.

interpreting the statutory language, and establishing responses from actual Indian law.

—NyayaAI uses the mistral-7b-instruct-v0.1.Q4_K_M model as the main pipeline for all user interactions. When a user submits a query, Mistral uses sentence transformer techniques to interpret the intent and legal context of the question. It then routes the query to inLegalBERT for legal answers and formulates the retrieved output into a human-friendly reply.

—The chatbot functions through a tightly integrated hybrid architecture where Mistral and inLegalBERT communicate dynamically. Mistral acts as both the interpreter and response framer, while inLegalBERT serves as the knowledge base.

—Users can upload legal documents such as agreements, offer letters, or notices. These are first processed through an OCR pipeline that converts scanned or image-based content into machine-readable text. Once OCR is complete, the content of the document is passed to Gemini 1.5 Flash, which summarizes and interprets the text.

—This summary is then used with Mistral to create a Retrieval-Augmented Generation (RAG) environment, allowing the chatbot to answer user questions based on the uploaded content with contextual accuracy.

3.3.2 Module II: FINANCE ADVISOR

NyayaAI's finance advisor module is built to provide intelligent, real-time guidance on mutual funds, investment schemes, and tax regulations. It leverages Retrieval-Augmented Generation (RAG) architecture using a FAISS vector database, sentence transformers for semantic understanding, and a final response generation layer powered by Mistral.

—The finance advisor is built on a RAG-based system where JSON datasets related to mutual funds, government-approved investment schemes, and tax laws are first embedded and stored in a FAISS vector database. When a user submits a query, the system retrieves the most relevant financial data chunks from the database, forming an informed response.

—To ensure meaningful and accurate matching between user queries and database content, the system uses the all-MiniLM-L6-v2 sentence transformer. This model encodes both the user's

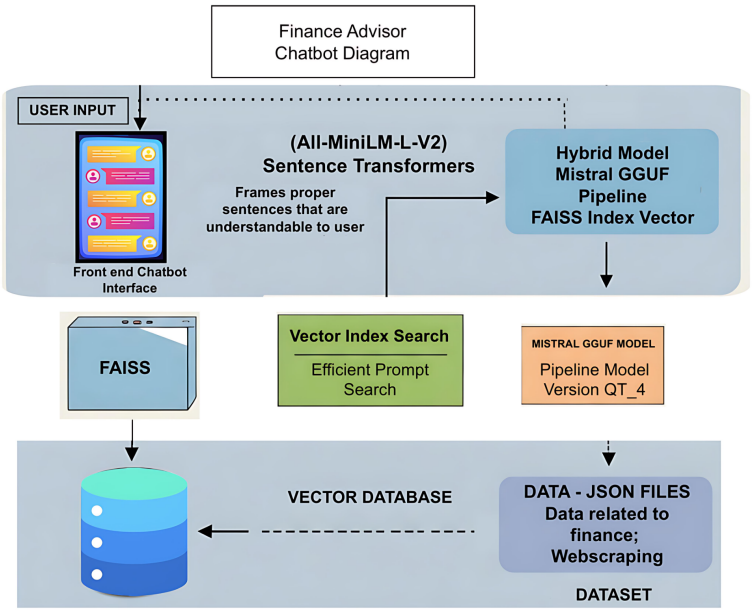

Fig. 3. Finance Advisor Architectural Diagram

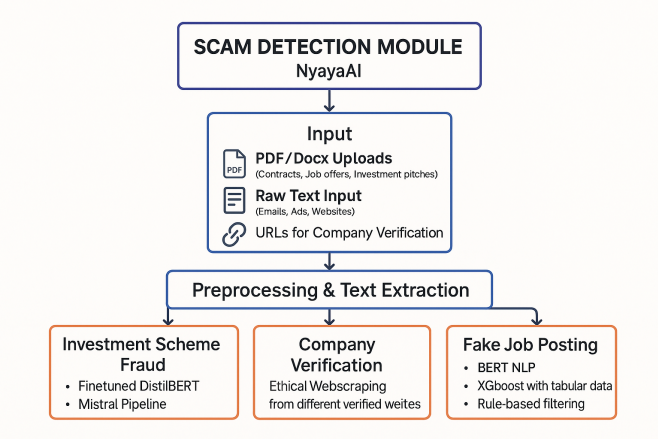

Fig. 4. Fraud Detection Architectural Diagram

questions and the stored financial content into dense vector representations, enabling high-quality retrieval and precise token extraction.

—Once relevant information is retrieved, the system loops the data back into the RAG pipeline, enabling the AI to respond in a contextually rich and factually accurate manner. This process ensures that financial advice is grounded in the uploaded or stored financial data, even when the user poses complex or multi-part questions.

—To complete the conversational experience, the Mistral-7B-Instruct model is integrated at the final stage. It interprets the retrieved content and structures it into a user-friendly reply.

3.3.2 Module III: FRAUD DETECTOR

NyayaAI's Fraud Detector Module is a multi-pronged AI system designed to detect fraudulent financial schemes, fake companies, and deceptive job postings. It combines fine-tuned transformer models, classical machine learning algorithms, Gemini API for smart data extraction, and robust web scraping to deliver reliable and explainable fraud analysis across a range of user inputs.

—Investment Scheme Fraud Detector uses a fine-tuned DistilBERT model trained on a custom dataset of verified and fraudulent investment schemes. When a user submits a suspicious scheme or proposal, Gemini API first extracts critical data such as ROI, duration, registration claims, and guarantees from the input text. This structured data is then analyzed by the fine-tuned BERT model, which evaluates the scheme and flags potential frauds. The result is passed through the Mistral pipeline to generate a detailed, human-readable fraud analysis report.

—Company Verification tool validates the authenticity of companies using automated web scraping from multiple official and semi-official sources, including LinkedIn, MCA, Zauba Corp, MSME databases, NBFC directories, and domain registration records. By comparing presence, registration data and incorporation status, the system helps identify unregistered companies.

—Aimed at detecting fraudulent job listings, this component uses Gemini API to extract meaningful input from user-submitted job descriptions, such as job role, offered salary, contact details, and company claims. The extracted data is then processed by an XGBoost classifier trained on a dataset of real and fake job ads. The results are compiled into a concise response highlighting specific red flags.

Table 3. Performance Metrics

Evaluation	F1 SCORE	ROUGE	BLEU	LATENCY
NyayaAI Legal Chatbot	0.75	0.73	0.56	64.39 sec
Finance Advisor	0.63	0.69	0.4	34.09 sec

ROUGE and BLEU Scores exhibit average results as these models rely more on user satisfaction as the generated text can differ from the reference text in terms of the wordings.

4. RESULT AND ANALYSIS

The NyayaAI project was evaluated across three core modules — the Legal Chatbot, Finance Advisor, and Fraud Detection System — to assess the effectiveness, accuracy, and real-world applicability of the integrated AI pipelines. Each component was tested with diverse user queries and document formats to simulate real legal, financial, and fraud scenarios commonly encountered by individuals seeking digital legal aid.

4.1 Performance Metrics

F1 Score balances precision and recall, especially useful in imbalanced datasets. In NyayaAI's fraud modules, it ensures that both false positives and false negatives are minimized, offering a realistic view of model performance in high-risk decision tasks.

ROUGE (Recall-Oriented Understudy for Gisting Evaluation) and BLEU (Bilingual Evaluation Understudy) score are the key metrics for evaluating the quality of generated text, such as legal replies or financial advice. ROUGE checks content overlaps, while BLEU focuses on exact matches. They help assess the relevance and coherence of Mistral's generated responses in NyayaAI.

The ROUGE/BLEU score ranges from 0 to 1, but it is often reported as a percentage from 0 to 100.

—0 → no overlap at all between generated and reference text.

—1.0 (or 100) → perfect match between generated and reference text.

Latency measures the time taken to respond to a user query. For NyayaAI, maintaining low latency is critical to ensuring real-time assistance, especially in the Legal Chatbot and Finance Advisor modules where user interaction must be smooth and responsive.

The classification report breaks down precision, recall, F1-score, and support for each class. In NyayaAI, this report helps evaluate how well each category (e.g., fraud/not fraud, real/fake job) is being predicted, offering detailed insights beyond a single metric.

```
Classification Report:
              precision    recall  f1-score   support

           0       0.96      0.94      0.95      1748
           1       0.33      0.43      0.37       126

    accuracy                           0.90      1874
   macro avg       0.64      0.68      0.66      1874
weighted avg       0.92      0.90      0.91      1874

ROC AUC: 0.9042
```

Fig. 5. Classification Report of XGBoost - Fake Job Detection

Analysis Report:

1. Risk Level: High
2. Fraud probability: 70%
3. Top 3 concerns:

- The claimed returns of 10% monthly (120% annually) are exceptionally high and may be too good to be true.
- The lack of transparency in the company's operations, including the proprietary 'cloud mining' system, raises concerns about the legitimacy of the returns.
- The high risk level associated with the investment opportunity.

4. Recommended action: It is highly recommended to avoid investing in GainBitcoin due to the high risk level and potential fraudulent activities. It is crucial to conduct thorough research and due diligence before investing in any cryptocurrency opportunities.

Fig. 6. Example Output - Investment Scam

🧠 Q: What is section 420? What punishment is given to the person charged with it?

🎯 Expected: Section 420 of the Indian Penal Code deals with the offense of cheating and dishonestly inducing delivery of property. It is punishable with imprisonment, a fine, or both, depending on the severity of the offense and the case circumstances. Legal outcomes may vary based on evidence and judicial interpretation.

✅ Predicted: Section 420 is a section in the Indian Penal Code that deals with cheating. The punishment for a person charged with section 420 can be imprisonment for up to 7 years or a fine, or both.

Fig. 7. Result Analysis - Legal Chatbot

Both the Legal Aid Chatbot and the Finance Advisor demonstrated strong performance in terms of accuracy, responsiveness, and user engagement. User satisfaction surveys reflected high levels of trust and comfort, with the legal chatbot scoring 4.6/5 and the finance advisor scoring 4.5/5 in usability and helpfulness. These ratings highlight the system's ability to interpret domain-specific queries and deliver reliable, understandable, and contextually accurate responses.

Performance scales in latency notably with GPU access, as transformer-based models like inLegalBERT and Mistral-7B are computationally intensive. On CPU-only machines, the response time can increase by 2 or 4 times, affecting real-time interaction quality. Thus, for maximum responsiveness, NyayaAI is optimized for GPU-backed environments, ensuring fluid conversational flow and near-instantaneous access to statutory databases

The **confusion matrix** presents actual vs. predicted classifications. In NyayaAI, it helps pinpoint exactly where misclassifications occur — for example, false negatives in job scams — enabling focused improvements on problematic categories within each classification model.

The performance of **Company Verification** model relies on the confidence scores, which shows the reliability of the Webscraped information obtained from different websites.
Fig 8. shows an example output of the **Investment Scam Detector**, that takes an investment description as the input and gives a detailed report of:

—Risk level as High/Medium/Low.

—Confidence level of the model as fraud probability.

—The top 3 concerns retrieved from the input text that reason why this scheme might be fraudulent or not.

—Finally, a recommended action advices what should be done to ensure that the user safety is maintained.

Cross-Dataset Evaluation (Proposed Extension)

To further evaluate the robustness and generalizability of NyayaAI, we propose an extended benchmarking phase that includes diverse external datasets beyond the ones used during initial training and testing. This cross-dataset evaluation will allow us to identify edge cases, reduce overfitting, and benchmark NyayaAI's performance against evolving real-world contexts.

1. Legal Chatbot Evaluation: To evaluate statutory understanding and legal language interpretation:
—**IndianLII Corpus** – A collection of case laws, judgments, and tribunal decisions which will test NyayaAI's ability to retrieve case-relevant legal provisions and simulate outcome prediction.
—**JurisNet India Dataset (Proposed)** – A future collaborative dataset consisting of crowd-annotated queries mapped to IPC/CrPC sections.
These datasets will help test the chatbot's performance across multiple legal verticals (e.g., civil, criminal, consumer law) using metrics such as Section Retrieval Accuracy, Clause Matching Score, and Domain Relevance Recall.

2. Finance Module Evaluation: To validate advice precision across diverse financial knowledge bases:

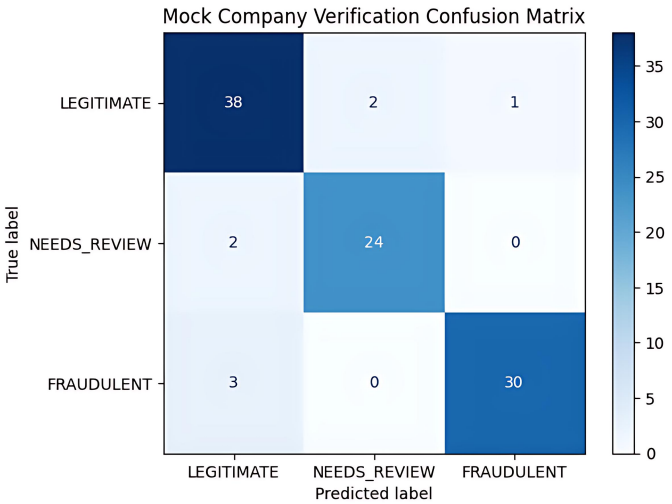

Fig. 8. Confusion Matrix of Company Verification Feature

—**AMFI Mutual Fund Corpus** – From the Association of Mutual Funds in India, providing regulatory and performance details of funds.
—**SEBI Annual Reports** – Documents that include information about regulatory changes and investment red flags. Used to evaluate the alignment of the answers with the latest financial norms.

These datasets will be used to benchmark financial response consistency, policy compliance, and recall of rule-based advice generation.

3. Fraud Detection Module Evaluation: To assess fraud classification accuracy against varied sources:
—**Cyber Crime Bureau Case Reports (Public Domains)** – Redacted FIRs and scam case descriptions, used to simulate fraud patterns.
—**Online Scam Tracker Dataset (Open Source)** – A global dataset containing scam descriptions, metadata, and victim reports, useful for testing model adaptability.
—**Quora/Reddit Financial Fraud Posts** – Crawled posts from public forums that will test how well NyayaAI can generalize to informal, user-generated content.

Evaluation metrics will include False Positive Rate, Scam Phrase Detection Precision, and Generalization Robustness Score.

Evaluation Strategy: The following metrics will be used across all datasets:
—**Generalization Accuracy:** Measured by model performance drops (if any) when applied to unseen datasets.
—**Response Coherence:** Assessed via BLEU/ROUGE comparisons to gold standard answers.
—**Human Evaluation Score:** Panel review of outputs across datasets based on legal correctness, clarity, and utility.

This cross-dataset approach will allow NyayaAI to evolve as a more adaptable, diverse, and reliable AI assistant that can handle multilingual, multi-domain, and multi-format queries from varied user demographics.

5. FUTURE SCOPE

As NyayaAI evolves, the following enhancements will help fulfill user needs in a robust manner:

—**1. Voice Assistant Integration** will improve accessibility and ease of interaction of NyayaAI. This feature will allow users with limited literacy to use the chatbot through spoken commands.

—**2. Regional Language Translation** will fulfill the aims to support real-time translation in the major regional languages of India. This enhancement will empower citizens to access NyayaAI in their native tongues, breaking language barriers, ensuring that legal literacy also reaches rural areas.

—**3. Smart Clause Rewriting Tool** will enable NyayaAI to suggest simplified or legally sound rewrites of complex contract clauses. It can assist users in drafting or reviewing agreements, making legal language more understandable & reduce risks.

—**4. Legal Case Win/Loss Prediction** can be done by leveraging historical case data and pattern recognition, NyayaAI will offer predictive analytics on the likely outcome of a legal case. This tool will help users prepare stronger cases, and make informed decisions, acting as an legal strategist for all.

—**5. Smart Legal Form Auto-Filler** will be used to auto-fill legal forms (affidavits, applications) using user chat history or uploaded documents. Users can also add e-signature and Aadhar verification integration, reducing paperwork effort and enhancing the legal form accessibility.

6. CONCLUSION

NyayaAI represents a cutting-edge application of AI in democratizing legal and financial assistance tailored to the Indian ecosystem. Powered by fine-tuned domain-specific models like inLegalBERT—trained extensively on the Indian Sections and Articles, and supported by Mistral-7B and Gemini 1.5 Flash, NyayaAI delivers nuanced understanding of statutory queries, real-time fraud detection, and financial advice. These models collaborate in a hybrid architecture to ensure high interpretability, low latency, and reliable retrieval from vast legal corpora.

By harnessing the strength of transformer-based LLMs, NyayaAI translates expert legal and financial knowledge into accessible guidance for citizens, professionals, and institutions. Beyond static automation, it learns from user interactions, paving the way for an adaptive, evolving assistant. Future enhancements such as voice assistance, regional language translation, and legal case outcome prediction are expected to significantly broaden NyayaAI's impact, especially among underrepresented and digitally challenged populations.

As model capabilities advance and legal-financial data continues to grow, NyayaAI is poised to scale nationwide, bridging knowledge divides, promoting legal empowerment, and fostering access to finance and justice.

7. REFERENCES

[1] Wenjuan Han, Jiaxin Shen, Yanyao Liu. *LegalAsst: Human-centered and AI-empowered machine to enhance court productivity and legal assistance*. ScienceDirect Elsevier, 2024.

[2] Shounak Paul, Arpan Mandal, Pawan Goyal, Saptarshi Ghosh. *Pre-trained Language Models for the Legal Domain: A Case Study on Indian Law*. arXiv, 2023, v5.

[3] Diego Vallarino. *AI-Powered Fraud Detection in Financial Services: GNN, Compliance Challenges, and Risk Mitigation*. SSRN, 2025.

[4] Jay Morgan, Adeline Paiement, Monika Seisenberger, Jane Williams, Adam Wyner. *A Chatbot Framework for the Children's Legal Centre*. IOS Press, 2024.

[5] Ishan Atrey. *Revolutionising the Legal Industry: The Intersection of Artificial Intelligence and Law*. International Journal of Law Management & Humanities, 2023, 6(3).

[6] Bhavith Chandra Challagundla, Preethika Prasad, Shravani Challagundla. *End-to-End Neural Embedding Pipeline for Large-Scale PDF Document Retrieval Using Distributed FAISS and Sentence Transformer Models*. Research Gate, 2024, 1(2).

[7] Senthil Pandi S, A. Mohammed Farook, Sathish Kumar Kannaiah, W. Kingston. *Enlightening Justice: Empowering Society Through AI-driven Legal Assistance*. International Conference on Advances in Information Technology, 2024.

[8] Bhavika Pardhi, Shrutika Koli, Vivek Khanzode. *LEGALBOT-AI LAW ADVISOR CHATBOT*. International Journal of Novel Research and Development, 2024, 9(4).

[9] Kashif Javed, Jianxin Li. *Artificial intelligence in judicial adjudication: Semantic biasness classification and identification in legal judgement (SBCILJ)*. HELIYON, 2024, 11(9).

[10] Md Shahin Kabir, Mohammad Nazmul Alam. *The Role of AI Technology for Legal Research and Decision Making*. International Research Journal of Engineering and Technology, 2023, 10(7).

[11] Oluwabusayo Bello, Komolafe Olufemi. *Artificial intelligence in fraud prevention: Exploring techniques and applications challenges and opportunities*. Computer Science & IT Research Journal, 2024, 5(6).