

AI-Based Legal Advisor for the Indian Judicial System: A Comprehensive Survey on Architectures, NLP Techniques, and Justice Delivery Automation

Dr B Uma

Malnad College Of Engineering
Hassan, Karnataka
bu@mcehassan.ac.in

Aril Srinivas

Malnad College Of Engineering
Hassan, Karnataka
arilsrinivas8@gmail.com

Giridhar Gowda G A

Malnad College Of Engineering
Hassan, Karnataka
giridhargowda1307@gmail.com

K M Ranjitha

Malnad College Of Engineering
Hassan, Karnataka
kmranjitha8971@gmail.com

Divyashree B L

Malnad College Of Engineering
Hassan, Karnataka
divyashreebl1205@gmail.com

Abstract— Artificial intelligence is increasingly being integrated into legal systems worldwide, with the Indian judicial system representing one of the most compelling and challenging application domains. India's courts face an unprecedented backlog exceeding 50 million pending cases, multilingual legal documents spanning 22 scheduled languages, diverse personal law frameworks, and an acute shortage of affordable legal representation for citizens. This survey reviews AI applications across the Indian legal lifecycle, encompassing case prediction, legal document analysis, multilingual natural language processing, chatbot-based legal advisory systems, bail and sentencing support, and constitutional query resolution. We examine single-agent, multi-agent, retrieval-augmented generation (RAG), and transformer-based architectures, comparing their suitability for Indian legal tasks. The survey further identifies unresolved challenges related to hallucination, low-resource language support, explainability, judicial ethics, and data privacy. We argue that AI is most effective as an augmentation layer that supports litigants, lawyers, and judges rather than as a replacement for human judicial reasoning, and we outline a roadmap for responsible deployment within the Indian constitutional framework.

I. INTRODUCTION

The Indian judicial system is one of the largest and most complex in the world. With a constitutional mandate to deliver justice to 1.4 billion citizens across 28 states and 8 union territories, India's courts operate across a multiplicity of legal frameworks including constitutional law, civil law, criminal law, corporate law, family law, and over 1,200 central and state statutes. The Supreme Court of India, 25 High Courts, more than 18,000 district courts, and an extensive network of subordinate tribunals together form a system that is simultaneously indispensable and overwhelmed.

The case backlog problem is structural and severe. As of 2024, more than 50 million cases remain pending across Indian courts, with district courts bearing the heaviest burden. The average time for disposal of a civil suit in certain states exceeds a decade, and criminal trials may extend even longer. Systemic causes include judicial vacancies, procedural complexity, inadequate infrastructure, and the sheer volume of filings. These delays impose severe economic and social costs, denying litigants timely relief and undermining public confidence in the rule of law [1], [2].

Legal access is equally constrained. Bar Council data consistently shows that the majority of India's legal practitioners are concentrated in metropolitan areas, leaving rural populations with limited access to qualified legal advice. Legal aid provided through state legal services authorities is understaffed and underutilised. For the average citizen, navigating even basic legal processes, such as filing an FIR, understanding bail entitlements,

or pursuing a consumer complaint, requires expertise that is practically inaccessible [3].

Artificial intelligence offers a promising set of tools to address these systemic challenges. Recent years have seen the emergence of AI applications tailored to legal tasks, including automated case summarisation, outcome prediction, bail recommendation, precedent retrieval, contract analysis, and multilingual legal query resolution. In the Indian context, these tools face unique design challenges: the legal corpus spans multiple languages including Hindi, Tamil, Telugu, Kannada, Bengali, and others; laws reference ancient statutes alongside post-independence legislation; and judicial reasoning frequently draws on a blend of statutory interpretation, constitutional principles, and socio-cultural context [4], [5].

This survey organises the literature on AI-based legal advisory systems specifically for the Indian judicial context. It makes four contributions. First, it reviews AI applications across each stage of the Indian legal lifecycle. Second, it presents an architectural taxonomy for legal AI systems relevant to India. Third, it compares tools and platforms currently deployed or proposed for Indian legal practice. Fourth, it identifies open research challenges and proposes a governance-aligned roadmap for responsible deployment. The central argument is that AI is most effective when it democratises legal access and assists judicial decision-making while keeping humans accountable for final outcomes.

II. DEFINITIONS

The development of AI-based legal systems requires careful use of several overlapping terms. This section clarifies the definitions of legal AI agents, natural language processing for law, retrieval-augmented legal systems, and agentic legal workflows.

A. Legal AI Agents

A legal AI agent is a computational system that perceives legal inputs, reasons about them using knowledge of statutes, precedents, and procedural rules, and produces outputs relevant to legal tasks. These outputs may include case summaries, predicted outcomes, applicable sections of law, procedural guidance, or draft documents. Unlike general-purpose AI assistants, legal AI agents are designed to handle domain-specific language, hierarchical legal reasoning, and citation-sensitive knowledge retrieval [6], [7].

Autonomy in legal AI agents spans a wide spectrum. At one end, a simple legal chatbot responds to keyword-matched queries about statutory provisions. At the other end, a fully agentic legal system can decompose a complex legal problem, retrieve relevant precedents, identify contradictory holdings, synthesise an

argument, and present it with citations. Most deployed systems fall somewhere between these extremes, combining structured legal knowledge bases with generative language models [8].

B. Natural Language Processing for Legal Text

Legal NLP refers to the application of computational language processing techniques to legal documents including statutes, judgments, contracts, FIRs, petitions, and regulatory instruments. Key NLP tasks in the legal domain include named entity recognition for identifying parties, sections, and courts; text classification for categorising case types; summarisation of lengthy judgments; question answering over legal corpora; and machine translation across India's scheduled languages [9], [10].

Indian legal NLP faces distinctive challenges. Indian legal English frequently borrows Latin maxims, archaic constructions from colonial-era drafting, and references to regional customary practices. Vernacular legal documents often code-switch between regional languages and English. Transliteration inconsistencies and the absence of standardised legal terminologies in many Indian languages further complicate NLP pipeline design [11].

C. Retrieval-Augmented Legal Systems

Retrieval-augmented generation (RAG) in the legal domain combines a retrieval component that searches case databases, statutory repositories, and legal encyclopaedias with a generation component that synthesises a natural-language response grounded in retrieved evidence. For legal applications, RAG architectures are especially valuable because they reduce hallucination, maintain citation traceability, and allow the system to ground responses in actual statutes and precedents rather than relying solely on parametric model memory [12].

In the Indian context, RAG systems must retrieve from a heterogeneous corpus that includes Supreme Court and High Court judgments in English, district court orders in regional languages, bare acts available in Hindi and English, and legal commentary produced in multiple scripts. Building retrieval indices that span this diversity requires multilingual embedding models, cross-script normalisation, and jurisdiction-aware retrieval logic [13].

D. Agentic Legal Workflows

Agentic legal workflows refer to systems in which multiple AI components cooperate in a sequenced or adaptive pipeline to handle complex multi-step legal tasks. Rather than treating legal assistance as a single prompt-response exchange, agentic workflows decompose tasks such as conducting a literature review of case law, drafting a petition, or preparing a bail application into subtasks handled by specialised components. These components may include a query interpreter, a legal retriever, a precedent comparator, a document drafter, and an output validator [14], [15].

The Indian legal system's complexity makes agentic workflows particularly relevant. A legal advisory task such as assessing the strength of a civil suit for property recovery requires understanding applicable personal law, identifying territorial jurisdiction, retrieving similar High Court decisions, checking limitation periods, and advising on procedural next steps. No single-pass model response is sufficient for such a task; only a structured multi-step workflow can produce advice of sufficient depth and reliability [16].

E. Scope of This Survey

This survey is scoped to AI systems that directly interact with or support the Indian judicial ecosystem. It covers systems designed for: legal query resolution and citizen advisory, case outcome prediction and judicial decision support, document

analysis including FIRs, contracts, and judgments, multilingual legal information retrieval, bail, sentencing, and case scheduling assistance, and governance and ethics frameworks for judicial AI. The survey does not cover general legal education platforms unrelated to Indian courts, nor does it cover AI used in corporate compliance or international arbitration without specific relevance to Indian domestic law.

III. AI APPLICATIONS ACROSS THE INDIAN LEGAL LIFECYCLE

Artificial intelligence intersects with the Indian legal system at multiple stages, from a citizen's first encounter with a legal problem through investigation, trial, judgment, and appeal. This section organises AI applications by stage and identifies representative systems and approaches.

A. Citizen Legal Access and Query Resolution

The most immediate impact of AI in Indian law is the democratisation of legal information. A large proportion of India's population has never consulted a lawyer. Legal chatbots and query-resolution systems can provide first-level guidance on procedural questions, applicable laws, and available remedies. Systems such as LIMRA (Legal Information and Management Research Assistant) and various state legal services chatbots have been deployed to provide guidance on matters including domestic violence protections under the Protection of Women from Domestic Violence Act, 2005, consumer rights under the Consumer Protection Act, 2019, and tenancy rights under state-specific rent control legislation [1], [3].

AI-based citizen advisory systems face the challenge of balancing accessibility with accuracy. A citizen asking whether they can file a writ petition directly in the Supreme Court requires a nuanced answer that depends on the nature of the fundamental right allegedly violated, the absence of an effective alternative remedy, and recent Supreme Court practice directions. Generative systems without strong retrieval grounding risk providing misleading procedural advice that could harm a litigant's case [17].

B. FIR Analysis and Criminal Case Support

The First Information Report (FIR) is the foundation of criminal proceedings in India. FIR text is often handwritten, colloquial, and may be recorded in a regional language. AI systems have been developed to automatically extract the applicable sections of the Indian Penal Code, 1860 (now the Bharatiya Nyaya Sanhita, 2023), the Code of Criminal Procedure, 1973, and other special laws from FIR text. This supports both investigating officers and defence lawyers in identifying the charges and potential defences from the earliest stage of a criminal case [18].

Subsequent developments include AI tools for predicting the outcome of bail applications based on FIR content, prior criminal record, nature of the offence, and jurisdictional patterns. Such systems have been piloted in several states as part of the National Judicial Data Grid (NJDG) ecosystem. However, concerns about algorithmic bias, particularly against marginalised communities, have prompted calls for careful auditing before any such system is used in judicial decision-making [19].

C. Case Outcome Prediction

Case outcome prediction is the most studied application of AI in Indian courts. Early approaches used rule-based systems or logistic regression trained on judgment metadata. More recent work applies transformer models fine-tuned on Indian legal corpora to predict outcomes for specific case categories including tax

disputes, bail applications, matrimonial matters, and constitutional petitions [20], [21].

The Indian Supreme Court Judgment Dataset (ILDC) has enabled a growing body of work on judgment prediction and rhetorical structure analysis of Indian court opinions. Researchers have demonstrated that BERT-based models fine-tuned on Supreme Court judgments can predict binary outcomes (allowed or dismissed) with reasonable accuracy. However, the practical utility of such predictions is contested: judges and lawyers express concern that outcome prediction may anchor litigants' expectations unfairly and that model accuracy varies substantially across case types and benches [22].

High Court judgment datasets, which are less uniformly structured and available in a wider variety of languages than Supreme Court decisions, present greater challenges for prediction models. Multi-jurisdiction prediction, where a model must reason about differences in judicial interpretation between, say, the Bombay High Court and the Madras High Court on the same statutory question, remains an open research problem [23].

D. Legal Document Analysis and Summarisation

Indian courts produce millions of pages of judgments annually. Lawyers, researchers, and policy analysts need tools that can quickly identify the ratio decidendi of a judgment, distinguish obiter dicta, extract statutory interpretations, and summarise procedural history. Transformer-based summarisation models trained on Indian legal text have been used to produce abstractive summaries of Supreme Court and High Court judgments, enabling legal practitioners to assess relevance before reading the full text [24].

Contract analysis is another important application, particularly for commercial practice and banking law. AI tools that flag unusual clauses, identify missing standard provisions, and compare against regulatory templates have been developed for use in Indian transactional practice. The Reserve Bank of India's regulatory framework for loan agreements and the Securities and Exchange Board of India's requirements for listed company contracts provide structured templates against which AI comparison tools can operate [25].

E. Multilingual Legal NLP

India's scheduled languages present one of the most demanding multilingual NLP challenges in the world for legal applications. While the Supreme Court and most High Courts operate primarily in English, district courts and many quasi-judicial tribunals use Hindi or a regional official language. Parties and witnesses communicate in local languages, and documents such as

sale deeds, nikah-namas, and property titles are frequently executed in regional scripts [11], [27].

Multilingual legal NLP for India requires low-resource language models for languages including Odia, Assamese, Konkani, Sindhi, and Kashmiri that lack large legal corpora. Cross-lingual transfer learning from English legal models to regional language models has shown promise but requires careful domain adaptation. The AI4Bharat initiative has produced multilingual NLP models including IndicBERT and IndicBART that have been applied to Indian legal text with encouraging results [28].

F. Judicial Decision Support and Scheduling

Beyond prediction, AI can support judicial workflow by recommending the allocation of cases to appropriate benches, flagging cases that have been pending beyond prescribed limits, identifying matters that may be disposed of through mediation or lok adalat, and suggesting appropriate bail conditions based on risk assessment frameworks. The Supreme Court of India's case management system and the e-Courts project have incorporated some automated scheduling and workflow features, though full AI-based decision support remains a future aspiration [30].

Sentencing support tools that align proposed sentences with sentencing guidelines and comparable precedents have been discussed in Indian legal reform circles, particularly following the Malimath Committee's recommendations on criminal justice reform. Pilot programmes for structured sentencing in narcotics cases have drawn on AI-assisted analysis of sentencing patterns across High Courts to reduce unwarranted inter-judge disparity [31].

G. Challenges and Limitations

Across all application areas, several limitations recur. Indian legal AI systems face data scarcity for low-resource languages, inconsistent digitisation of court records especially at the district level, model hallucination in generative responses, the risk of automating historical biases present in judicial data, and inadequate explainability for outputs that may influence judicial decisions. Ethical concerns including judicial independence, due process, and the right to a fair hearing are repeatedly invoked as reasons for caution [32], [33].

The absence of a regulatory framework specifically governing AI in judicial administration is a significant gap. While the Personal Data Protection law provides some baseline, its application to judicial AI systems is unclear. The Bar Council of India's ethical guidelines do not yet address AI-assisted legal practice. These governance gaps mean that even well-designed AI systems may be deployed without the oversight structures necessary to ensure accountability [34].

TABLE I
Major AI Applications Across the Indian Legal Lifecycle

Legal Stage	Representative Systems / Approaches	Primary Capability	Main Benefit	Key Limitation
Citizen legal access	Legal chatbots, LIMRA, state LSA portals [1],[3]	Query resolution, procedural guidance, rights information	Democratizes first-level legal access for rural and low-income citizens	Risk of misleading advice without strong retrieval grounding
FIR and criminal case support	IPC section classifiers, bail prediction models, NJDG tools [18],[19]	Section identification, bail outcome analysis, charge interpretation	Speeds up charge framing and supports defence preparation	Algorithmic bias risks; requires independent auditing
Case outcome prediction	ILDC-trained BERT models, High Court classifiers [20],[22]	Binary and multi-class outcome prediction from case text	Assists lawyers in assessing case strength and advising clients	Variable accuracy across benches; anchoring bias risk

Legal document analysis	Judgment summarisers, contract review tools, citation networks [24],[26]	Abstractive summarisation, clause flagging, ratio extraction	Reduces time burden on lawyers and researchers	Hallucination in abstractive models; citation accuracy issues
Multilingual legal NLP	IndicBERT, IndicBART, cross-lingual transfer models [28],[29]	Translation, NER, summarisation across Indian languages	Extends legal AI to non-English-speaking litigants	Low-resource language data scarcity; legal fidelity gaps
Judicial decision support	Case scheduling tools, e-Courts AI modules, sentencing analytics [30],[31]	Case allocation, delay flagging, sentencing consistency analysis	Reduces pendency and improves scheduling efficiency	Judicial independence concerns; governance framework absent
Constitutional query resolution	RAG systems on constitutional corpus, writ petition classifiers [12],[13]	Article identification, fundamental rights analysis, precedent retrieval	Enables citizens and lawyers to navigate constitutional provisions	Requires domain expert validation; hallucination risk on complex doctrines

IV. ARCHITECTURES

The architecture of an AI-based legal advisor for India must navigate distinctive requirements: multilingual input, citation-sensitive output, hierarchical statutory knowledge, and judicial ethics constraints. This section reviews the major architectural families and their applicability to Indian legal tasks.

A. Classical Rule-Based and Expert System Architectures

The earliest computerised legal advisory systems in India and globally were rule-based expert systems that encoded legal rules as logical if-then statements. TAXMAN, developed in the 1970s for US corporate tax law, exemplified this approach. Indian counterparts included early computerised legal databases maintained by organisations such as Manupatra and SCC Online, which allowed keyword-based retrieval but not reasoning [6].

Rule-based systems are attractive for well-defined legal domains with limited statutory complexity. They offer complete explainability: every advisory output can be traced to a specific rule. For tasks such as computing limitation periods under the Limitation Act, 1963, or determining stamp duty obligations under state stamp acts, rule-based approaches remain competitive. Their fundamental weakness is the cost of manual knowledge engineering and their inability to handle ambiguous, conflicting, or evolving legal provisions without expert intervention [7].

B. Machine Learning Architectures for Legal Classification

The application of machine learning to Indian legal text began with support vector machines and logistic regression applied to case type classification and outcome prediction. Feature engineering for legal ML typically involves TF-IDF representations of judgment text, enriched with hand-crafted features such as party type, court level, subject matter category, and presence of constitutional challenges. These approaches achieved reasonable accuracy on binary classification tasks but required substantial labelled data and did not generalize well across legal domains [20].

Decision tree and random forest models have been used for bail decision support, encoding features such as the nature of the offence, the accused's prior record, the availability of sureties, and flight risk indicators. The transparency of tree-based models makes them relatively more acceptable to judicial stakeholders than black-box deep learning models, although they still require careful feature selection to avoid encoding socioeconomic proxies that may introduce discriminatory patterns [19].

C. Transformer-Based Language Model Architectures

The introduction of transformer-based language models has substantially advanced legal NLP capability. BERT-based models fine-tuned on legal corpora, known as LegalBERT in the English-language context, have demonstrated strong performance on tasks including judgment summarisation, legal question answering, and statute classification. For Indian legal text, models such as InLegalBERT, fine-tuned on the ILDC corpus of Supreme Court judgments, have been developed and evaluated [21].

Multilingual transformer models including mBERT, XLM-R, and the A14Bharat IndicBERT models have been applied to cross-lingual legal NLP tasks in India. These models allow a single architecture to handle legal text in multiple Indian languages, enabling zero-shot or few-shot transfer from resource-rich languages such as English or Hindi to lower-resource legal languages such as Odia or Assamese. However, their performance on legal-specific tasks in lower-resource languages remains substantially below English baselines [28].

Large language models including GPT-4, Claude, and Gemini have been prompted for Indian legal tasks using zero-shot and few-shot approaches. Their ability to reason about complex legal scenarios, identify applicable statutes, and draft legal text is remarkable, but their tendency to hallucinate statutory provisions, misattribute precedents, or confidently assert incorrect legal propositions makes them unreliable as standalone legal advisors without retrieval augmentation and output validation [35].

D. Retrieval-Augmented Generation Architectures for Indian Law

RAG architectures are particularly well suited to Indian legal advisory tasks because they ground generated responses in retrievable statutory and case law sources, reducing hallucination and improving citation accuracy. A RAG-based Indian legal advisor typically consists of a retrieval component that queries indices built from Supreme Court and High Court judgments, bare acts, and legal commentary, and a generation component that synthesises a response using the retrieved evidence [12].

Building effective retrieval indices for Indian legal RAG requires several design decisions. Dense retrieval using bilingual or multilingual embedding models outperforms sparse BM25 retrieval for semantically complex legal queries, though BM25 retains value for citation lookup and statutory keyword search. Hierarchical chunking strategies, which preserve section-level context from statutes and paragraph-level context from judgments, improve retrieval precision for specific legal questions. Jurisdiction-aware retrieval filters that distinguish Supreme Court precedents binding on all courts from High Court precedents

binding only within their territorial jurisdiction are essential for legally accurate advisory outputs [36].

Agentic RAG architectures extend this model by adding planning and multi-step retrieval. A legal advisory agent may first classify the legal domain of a query, then retrieve relevant statutes, then retrieve applicable High Court and Supreme Court precedents, then resolve conflicts between retrieved holdings, and finally synthesise an answer with citations. This multi-hop retrieval pattern mirrors the research process of an experienced advocate and produces more accurate and complete advisory outputs than single-pass RAG [37].

E. Multi-Agent Architectures for Legal Workflows

Complex legal tasks such as drafting a writ petition, preparing a bail application, or conducting a due diligence review require the coordination of multiple specialised capabilities. Multi-agent architectures address this by distributing legal workflow tasks across agents with distinct roles. A typical Indian legal advisory multi-agent system might include a legal domain classifier agent, a statute retrieval agent, a precedent analysis agent, a document drafting agent, and a validation agent that checks the draft for legal accuracy, citation correctness, and procedural compliance [14], [15].

Multi-agent designs improve scalability and allow different components to be updated independently. When a new statute is enacted, such as the replacement of the Indian Penal Code by the Bharatiya Nyaya Sanhita in 2023, the statute retrieval agent's index can be updated without retraining the entire system. Similarly, improvements to the precedent analysis agent do not affect the document drafting agent. This modularity is practically important for a legal domain where the statutory corpus evolves through legislative amendment, judicial interpretation, and constitutional review [38].

Communication and coordination between legal agents require careful design. In an orchestrator-worker pattern, a central legal reasoning agent decomposes the query, delegates subtasks to specialist agents, and assembles the final advisory output. In a peer-to-peer pattern, agents communicate directly and negotiate over conflicting legal interpretations. The orchestrator-worker pattern is generally more predictable and auditable, which is important for legal contexts where the reasoning chain must be explicable [39].

F. Knowledge Graph Architectures for Statutory Representation

Indian statutory law is highly cross-referential. The Constitution of India, 1950 is referenced in nearly every domain of public law. The Indian Penal Code cross-references the Code of Criminal Procedure. State legislation must be read against central legislation on concurrent list subjects. Knowledge graph architectures that represent these cross-references as structured relational data allow legal AI systems to reason about statutory interactions more precisely than text-based retrieval alone [40].

Legal knowledge graphs for India encode entities including statutes, sections, sub-sections, judicial interpretations, constitutional articles, and doctrines such as the basic structure doctrine. Relationships encoded include 'amends', 'repeals', 'is interpreted by', 'contradicts', and 'is consistent with'. Traversal of these graphs supports multi-hop legal reasoning, such as identifying all statutes affected by a constitutional amendment or tracing the judicial history of a specific statutory interpretation [41].

G. Architectural Trade-offs and Design Implications

Across architectural families, several trade-offs recur for Indian legal AI design. Rule-based systems offer explainability but lack flexibility. ML classifiers are data-efficient but brittle under distribution shift. Transformer models are powerful but prone to hallucination without retrieval grounding. RAG systems improve factual reliability but depend on retrieval index quality and coverage. Multi-agent architectures improve task coverage but increase coordination complexity. Knowledge graph architectures improve structured reasoning but require expensive manual curation and are slow to update.

The most promising direction for a production-grade Indian legal advisor is a hybrid architecture combining a multilingual transformer backbone with RAG grounded in authenticated legal sources, knowledge graph-based statutory cross-referencing, multi-agent workflow decomposition for complex tasks, and human-in-the-loop validation for outputs that may influence judicial proceedings. This combination addresses the core requirements of accuracy, citation traceability, multilingual reach, and procedural reliability [36], [37], [40].

TABLE II
Comparison of AI Agent Architectures for Indian Legal Advisory Systems

Architecture Type	Core Components	Primary Strengths	Main Limitations	Indian Legal Relevance
Rule-Based Expert Systems [6],[7]	Legal rules as if-then logic, structured knowledge bases	Full explainability, deterministic outputs, no hallucination	High manual curation cost, brittle under legal change, poor generalisation	Suitable for fixed procedural tasks: limitation computation, stamp duty, jurisdiction rules
ML Classification Models [19],[20]	SVM, logistic regression, random forest on legal features	Moderate explainability, efficient with labelled data, useful for structured tasks	Requires large labelled datasets, poor cross-domain transfer	Useful for case type classification, bail prediction, outcome probability estimation
Transformer-Based LLMs [21],[28]	BERT, XLM-R, IndicBERT fine-tuned on Indian legal corpora	Strong language understanding, cross-lingual transfer, flexible task handling	Hallucination risk, multilingual performance gaps, opaque reasoning	Core capability for legal NLP: summarisation, NER, classification, QA
Retrieval-Augmented Generation [12],[36]	Dense retriever, legal vector store, LLM generator with citation grounding	Grounded responses, citation traceability, reduced hallucination	Dependent on index quality and coverage across languages and jurisdictions	High relevance for statute and precedent lookup, citizen advisory, writ research

Agentic RAG [37],[38]	Multi-step retrieval planner, legal reasoning agent, validation layer	Complex multi-hop legal reasoning, iterative refinement, adaptive query	Higher latency, coordination overhead, harder to debug	Best fit for multi-issue legal queries requiring statute + precedent + procedural analysis
Multi-Agent Legal Workflows [14],[15]	Domain classifier, statute agent, precedent agent, drafting agent, validator	Specialisation, modular update, scalable task decomposition	Coordination complexity, inter-agent communication overhead	Suitable for petition drafting, due diligence review, bail application preparation
Legal Knowledge Graphs [40],[41]	Entity and relation extraction, ontology of statutes, amendment tracking	Structured statutory reasoning, cross-reference resolution, multi-hop traversal	Expensive to build and maintain, limited to pre-encoded relations	Essential for constitutional cross-referencing and tracking statutory amendment chains

V. EXISTING TOOLS AND PLATFORMS

The Indian legal technology landscape has matured considerably over the past decade, with both domestic startups and established legal information services developing AI-powered tools. This section reviews representative tools across citizen advisory, legal research, document automation, and judicial administration categories.

A. Legal Database and Research Platforms

Manupatra and SCC Online are India's two largest commercial legal databases, providing access to Supreme Court and High Court judgments, statutory texts, and legal commentary. Both platforms have introduced AI-powered features including natural language search, judgment summarisation, and citation analysis. Manupatra's NLP-assisted search allows users to query in natural language rather than Boolean keyword strings, with the system mapping the query to relevant legal concepts before retrieving matching judgments [42].

Indian Kanoon is a freely accessible legal search engine indexing Supreme Court, High Court, and tribunal judgments. It uses structured retrieval with keyword and citation-based search. While lacking deep AI integration as of this writing, its open access model and comprehensive coverage make it a critical data source for training Indian legal NLP models. Several academic studies of Indian legal AI have used Indian Kanoon as their primary data source [43].

B. AI-Powered Legal Chatbots and Advisory Tools

LIMRA (Legal Information and Management Research Assistant) is a chatbot designed to provide citizens with information about their legal rights and available remedies. Deployed through web and WhatsApp interfaces, LIMRA handles queries about domestic violence, labour rights, consumer complaints, and property disputes. Its architecture combines intent classification with a curated legal knowledge base, generating responses based on the identified legal domain and the user's state of residence, which determines applicable state legislation [1].

Nyaya Mitra, operated under the Department of Justice, provides legal awareness and basic advisory services through trained para-legal volunteers supported by an AI-assisted knowledge base. The system identifies the applicable law and nearest legal services authority for a citizen's problem, providing procedural roadmaps in Hindi and English. Evaluations of Nyaya Mitra deployment suggest significant improvement in first-level legal awareness among rural users, though the advisory depth is limited by the chatbot's inability to reason about complex multi-issue scenarios [3].

Private sector legal advisory startups including Lawrato, Vakil Search, and MyAdvo have integrated AI features for query triage, lawyer matching, and document template generation. These

platforms use classification models to identify the type of legal problem a user faces, match it with appropriate lawyer specialisations, and provide estimated cost ranges for legal representation. Their AI components are primarily assistive rather than advisory, connecting users to human lawyers rather than providing substantive legal guidance [44].

C. Document Automation and Contract Analysis Tools

SpotDraft and Leegality are Indian legal technology platforms focused on contract lifecycle management. SpotDraft's AI features include contract risk scoring, clause deviation flagging against standard templates, and automated extraction of key contract metadata such as parties, dates, governing law, and renewal triggers. Its models have been trained on Indian commercial contract corpora and are calibrated against regulatory requirements of Indian law including the Indian Contract Act, 1872, and SEBI regulations [25].

Leegality specialises in digital signature infrastructure and document workflow automation for Indian courts and enterprises, incorporating e-stamp integration aligned with the Court Fees Act and state-specific stamp legislation. Its AI components assist in identifying the correct stamp duty applicable to a transaction based on the document type, transaction value, and state of execution, a task requiring multi-jurisdictional statutory knowledge [45].

D. Judicial Administration AI Systems

The National Judicial Data Grid (NJDG) maintained by the National Informatics Centre provides case status information, pendency analytics, and judicial performance statistics across Indian courts. The e-Courts project has introduced AI-supported features including automated case categorisation on filing, digital cause list generation, and SMS-based litigant notifications. These administrative AI tools focus on workflow efficiency rather than substantive legal analysis [30].

The Supreme Court of India has piloted voice-to-text transcription for courtroom proceedings and AI-assisted translation of orders into Hindi and regional languages. The Vidhi Centre for Legal Policy has developed AI tools for statutory drafting quality assessment, flagging provisions that may create interpretive ambiguity or conflict with existing law. These judicial-facing tools represent an emerging category of AI assistance directly embedded in court operations [46].

E. Comparative Perspective on Existing Tools

Across the tools reviewed, a clear pattern emerges. Most mature Indian legal AI tools operate at the retrieval and triage layer, providing information lookup and query routing rather than substantive legal reasoning. The gap between what commercially deployed tools offer and what research prototypes demonstrate is significant: academic models for judgment prediction and multilingual summarisation substantially outperform deployed

commercial tools on benchmark tasks, but commercial deployments are more carefully engineered for reliability, data privacy, and user experience in real legal contexts [47].

A second pattern is the concentration of AI capability in English-medium legal domains. Tools offering full AI functionality in Hindi or regional languages remain rare, limiting the reach of legal AI to English-literate urban populations. Bridging this gap requires not only better multilingual models but also curated regional-language legal corpora, user experience design for low-literacy users, and integration with voice interfaces suited to non-literate citizens [11], [28].

TABLE III
Comparison of Existing AI-Based Legal Tools and Platforms for India

Tool / Platform	Primary Objective	Core Technologies	Main Outputs	Reported Advantages	Limitations
Manupatra AI Search [42]	Legal research and case law retrieval for practitioners	NLP-powered natural language search, judgment summarisation, citation analysis	Relevant judgments, summaries, citation maps	Comprehensive coverage, practitioner-friendly interface, regular updates	Subscription-based access limits reach; English-heavy corpus
Indian Kanoon [43]	Free access legal search across courts and tribunals	Keyword and citation-based structured retrieval	Judgment search results, citation links	Open access, wide court coverage, API availability for researchers	Limited AI integration; no summarisation or QA layer
LIMRA Chatbot [1]	Citizen legal information and rights awareness	Intent classification, curated legal knowledge base, WhatsApp integration	Legal rights information, procedural guidance, remedy identification	Accessible to non-lawyers; multi-channel deployment	Limited to pre-coded legal domains; no complex multi-issue reasoning
Nyaya Mitra [3]	Legal awareness and legal services authority access	Taxonomy-based legal domain identification, para-legal support integration	Applicable law identification, nearest LSA referral, procedural roadmap	Rural deployment success; Hindi language support	Advisory depth limited; not suitable for complex disputed matters
SpotDraft [25]	Contract lifecycle management and risk analysis	Clause classification, deviation flagging, metadata extraction, LLM integration	Risk-scored contracts, flagged clauses, extracted metadata	Tailored to Indian contract law; SEBI-aligned templates	Enterprise-focused; not accessible to individual litigants or SMEs
e-Courts / NJDG [30]	Judicial administration and case management	Automated categorisation, pendency analytics, digital cause lists, SMS alerts	Case status, delay reports, judicial performance metrics	Pan-India coverage; supports judicial accountability	No substantive legal reasoning; AI limited to administrative workflow
Vidhi Legal Drafting AI [46]	Statutory drafting quality assessment	NLP-based ambiguity detection, cross-statute conflict flagging	Quality assessment reports, flagged provisions, drafting recommendations	Directly supports legislative drafting quality	Limited to structured statutory text; not a citizen-facing tool

VI. ETHICS, ACCOUNTABILITY, AND SECURITY

The deployment of AI in the Indian judicial system raises ethical questions that are both universal and distinctively rooted in India's constitutional and social context. This section reviews the ethical principles applicable to judicial AI, the specific risks arising from the Indian legal environment, and the security considerations relevant to agentic legal systems.

A. Constitutional and Fundamental Rights Dimensions

India's Constitution guarantees fundamental rights including the right to equality (Article 14), the right against discrimination

(Article 15), the right to life and personal liberty (Article 21), and the right to a fair trial. Any AI system deployed in judicial contexts must be assessed against these constitutional standards. An AI bail recommendation system that systematically disadvantages accused persons from scheduled castes or tribes, even if unintentionally through proxy variables, violates Articles 14 and 15 and may constitute an unconstitutional infringement of Article 21 [32], [48].

The right to a reasoned judgment is a fundamental procedural requirement in Indian constitutional law, established by the Supreme Court in *Union of India v. Mohan Lal Capoor* (1973). Any AI-assisted judicial decision must preserve this requirement: the

judge must be able to articulate the reasons for the decision, and those reasons must be traceable to legal principles rather than algorithmic outputs that the judge cannot explain. This requires AI systems used in judicial decision support to produce human-readable reasoning chains, not merely numerical scores [49].

B. Data Privacy and Judicial Confidentiality

Legal proceedings in India involve large volumes of sensitive personal data including medical records, financial information, family relationships, criminal histories, and confidential communications. AI systems trained on judicial data must comply with the Digital Personal Data Protection Act, 2023, and must implement purpose limitation, data minimisation, and storage limitation principles. The use of judicial data to train commercial AI models without explicit authorisation raises serious concerns about the secondary use of sensitive personal data [50].

C. Algorithmic Bias in the Indian Legal Context

Historical judgments used to train Indian legal AI models inevitably reflect the social biases of the periods in which they were decided. Landmark cases studying caste-based discrimination, gender inequality in personal law, and socioeconomic disadvantage in bail and sentencing reveal patterns in judicial decision-making that, if learned by AI models, may perpetuate rather than correct systemic inequities. Training data curation, bias auditing, and differential impact assessment are essential pre-deployment requirements for any AI system that may influence judicial outcomes [33], [52].

D. Judicial Independence and AI Authority

A foundational principle of Indian constitutional law is the independence of the judiciary from the executive and legislative branches. The introduction of AI systems into judicial decision-making raises the question of whether AI recommendations constitute a form of external influence on judicial decision-making that compromises this independence. While AI tools designed as advisory aids rather than decision-makers do not formally impinge on judicial independence, the practical risk that judges may over-

rely on algorithmic outputs without independent reasoning is a documented phenomenon in comparative legal literature [54].

E. Security Threats in Agentic Legal Systems

Agentic legal advisory systems face distinctive security risks arising from their combination of language model reasoning, external retrieval, and tool use. Prompt injection attacks may embed malicious instructions in documents processed by the legal agent, causing the agent to retrieve incorrect precedents, generate misleading legal advice, or expose confidential case information. This risk is particularly serious in document analysis workflows where the agent processes uploaded FIRs, contracts, or affidavits that may contain adversarially crafted text [56].

F. Ethical Framework for Indian Legal AI Deployment

Drawing on the reviewed literature and the Indian constitutional framework, a responsible deployment framework for Indian legal AI should incorporate the following principles. First, transparency: AI systems must clearly disclose their role to users and must produce outputs that users can understand and challenge. Second, accountability: responsibility for AI-assisted legal advice must be allocated clearly. Third, fairness: systems must be audited for disparate impact across caste, religion, gender, economic status, and linguistic community before deployment in judicial contexts. Fourth, privacy: judicial data must be handled in compliance with the Digital Personal Data Protection Act with strict purpose limitation. Fifth, human oversight: no AI output should directly determine a judicial outcome without review and endorsement by a human judge [48], [55].

TABLE IV
Comparison of Ethical and Security Issues in AI-Based Indian Legal Systems

Issue Category	Specific Concern	Indian Legal Context	Applicable Framework	Recommended Mitigation
Constitutional rights [48],[49]	AI outputs that violate equality and fair trial guarantees	Articles 14, 15, 21 of Constitution; reasoned judgment requirement	Supreme Court jurisprudence on judicial reasoning	Explainable AI outputs; judge review mandatory for all AI-assisted decisions
Data privacy [50],[51]	Sensitive judicial data used to train commercial models	Digital Personal Data Protection Act, 2023; advocate-client privilege	DPDPA purpose limitation and data minimisation principles	Strict data isolation; no cross-client training; judicial data governance policy
Algorithmic bias [33],[52]	Historical biases in judgments reproduced or amplified by AI	Caste, gender, religious community biases in legacy judicial data	Constitutional equality mandate; SC anti-discrimination jurisprudence	Differential impact assessment; bias auditing before deployment; diverse training data curation
Personal law complexity [53]	Incorrect application of applicable personal law	Hindu, Muslim, Christian, Parsi personal laws with community-specific rules	Personal law statutes; Supreme Court interpretations	Personal law domain classifier; community-specific legal knowledge bases
Judicial independence [54],[55]	Over-reliance on AI recommendations by judges	Constitutional mandate of independent judicial reasoning	Separation of powers doctrine; EU AI Act as reference	AI as advisory only; mandatory disclosure when AI used; periodic reliance auditing
Prompt injection [56]	Adversarial content in uploaded documents hijacking agent behaviour	FIRs, contracts, affidavits as potential attack vectors	Emerging AI security standards	Input sanitisation; sandboxed document processing; output validation layer

Memory poisoning [57]	Corrupted legal information persisting across user sessions	Long-lived legal advisory agents with persistent memory	Secure software engineering principles	Session isolation; memory validation; adversarial testing before deployment
-----------------------	---	---	--	---

VII. RESEARCH GAPS

Despite significant recent progress, the field of AI-based legal advisory for India contains several important unresolved gaps that limit the dependability and reach of current systems.

First, there is no standardised benchmark for evaluating Indian legal AI systems across the full range of legally relevant tasks. The ILDC dataset covers Supreme Court judgment prediction, but there is no equivalent benchmark for High Court multilingual summarisation, district court procedural guidance, or bail recommendation accuracy across jurisdictions. Without shared benchmarks, comparing architectural approaches or measuring progress over time is difficult [22].

Second, multilingual legal AI remains severely underdeveloped for India's lower-resource scheduled languages. Models that perform acceptably in English and Hindi show substantial degradation when applied to legal text in Odia, Assamese, Manipuri, Sindhi, or Kashmiri. The scarcity of digitised court documents in these languages creates a data bottleneck that pre-training on general language corpora cannot fully resolve [28], [29].

Third, explainability methods for legal AI outputs are inadequate for judicial contexts. Post-hoc explanation techniques such as LIME and SHAP provide feature importance scores that legal practitioners find difficult to interpret in terms of legal reasoning. What is needed are explanation frameworks that map AI outputs to specific statutory provisions, precedent holdings, and legal principles in a way that a judge or lawyer can engage with critically [36].

Fourth, the interaction between AI legal advice and India's personal law pluralism has received little research attention. The distinct legal frameworks governing different religious communities in family law create a complex multi-classification problem that current legal NLP systems are not designed to handle reliably. Errors in this domain carry serious consequences for users navigating divorce, inheritance, and adoption matters [53].

Fifth, longitudinal evaluation of AI legal advisory tools in real-world deployment is almost entirely absent from the literature. Most studies evaluate models on held-out portions of historical judgment datasets rather than tracking the downstream impact of AI advice on user decisions, legal outcomes, or access to justice metrics. Responsible deployment requires impact evaluation that extends beyond accuracy on benchmark tasks [33].

Finally, governance frameworks specifically designed for AI in Indian judicial administration are nascent. The Digital Personal Data Protection Act, the proposed National Data Governance Framework, and Bar Council ethical guidelines have not been updated to address AI specifically. Without clear governance, liability allocation, and professional conduct standards for AI-assisted legal practice, even technically sound systems cannot be deployed responsibly [55].

VIII. FUTURE DIRECTIONS

The future of AI-based legal advisory in India is likely to develop along several converging directions that address current technical and governance limitations.

A primary direction is the development of a comprehensive Indian Legal AI Benchmark covering case outcome prediction, multilingual summarisation, bail application analysis, constitutional query resolution, and procedural guidance across all twenty-two scheduled languages. Such a benchmark would require collaboration between legal scholars, NLP researchers, and judicial institutions, and would need to be updated as the legislative corpus evolves [22].

A second direction is the construction of large multilingual Indian legal corpora through systematic digitisation of district court records, quasi-judicial tribunal orders, and High Court judgments in regional languages. The National Mission for Justice Delivery and Legal Reform provides a policy framework for this effort, and AI-assisted digitisation and OCR for handwritten Devanagari, Tamil, and Bengali legal text represents a tractable near-term goal [11].

A third direction is the development of explainable AI architectures designed specifically for judicial contexts, producing reasoning chains that reference statutes by section number, precedents by citation, and legal principles by doctrine name. Such architectures would allow judges and lawyers to engage critically with AI outputs rather than accepting them as opaque recommendations [36], [49].

A fourth direction is the design of AI systems that are constitutionally compliant by design, incorporating fairness constraints aligned with Article 14's equality mandate, impact assessment procedures for Articles 15 and 21, and procedural safeguards that preserve the right to a reasoned decision. This requires collaboration between AI engineers and constitutional lawyers from the earliest stages of system design rather than treating constitutional compliance as a post-hoc evaluation criterion [48].

A fifth direction is the development of voice-first and low-literacy interfaces for legal AI, enabling citizens without smartphone literacy or English proficiency to access legal advisory through telephone, radio, or voice assistants in regional languages. This requires integration of multilingual speech recognition, natural language understanding in Indian languages, and text-to-speech synthesis calibrated for legal vocabulary in regional languages [3].

Finally, a governance roadmap for Indian legal AI should be developed through a multi-stakeholder process involving the Supreme Court's e-Committee, the Bar Council of India, the Ministry of Law and Justice, civil society organisations, and AI researchers. This roadmap should address liability allocation for AI-assisted legal advice, mandatory disclosure requirements, auditing standards, and update protocols to keep AI systems aligned with evolving legal standards [55].

IX. CONCLUSION

This survey has reviewed the application of artificial intelligence to legal advisory and judicial support functions within the Indian judicial system. The central finding is that AI offers substantial and near-term potential to address India's legal access crisis and case backlog problem, but that current systems fall short of the reliability, multilingual coverage, explainability, and governance alignment required for responsible deployment in judicial contexts.

The strongest available architectures for Indian legal AI combine multilingual transformer models with retrieval-augmented generation grounded in authenticated legal sources, multi-agent workflow decomposition for complex legal tasks, and knowledge graph-based statutory representation. These hybrid architectures are more accurate, more traceable, and more trustworthy than either pure generative models or traditional rule-based systems.

Ethical deployment requires that AI be treated as an augmentation tool that supports citizens, lawyers, and judges rather than as a replacement for human legal reasoning. The constitutional guarantees of equality, fair trial, and reasoned judgment must be treated as binding design constraints, not aspirational targets. Algorithmic bias auditing, data privacy compliance, professional conduct alignment, and judicial independence safeguards must be integrated into AI systems from their earliest design stages.

India's position as a multilingual, multi-jurisdictional, pluralistic democracy with one of the world's largest court systems makes it both a demanding and a uniquely important testbed for legal AI. Success in this context will require sustained collaboration between legal scholars, computational linguists, judicial institutions, civil society, and policymakers. The goal should be a legal AI ecosystem that is as accessible as a mobile phone, as trustworthy as a trained advocate, and as accountable as a constitutional court.

. REFERENCES

- [1] S. K. Vishnumolakala et al., "LIMRA: AI-based legal information chatbot for citizen advisory," in Proc. IEEE EDUCON, 2024.
- [2] Law Commission of India, "Arrears and Backlog: Creating Additional Judicial (wo)manpower," Report No. 245, 2014.
- [3] Department of Justice, Government of India, "Nyaya Mitra: Legal Awareness and Access Programme," Ministry of Law and Justice, 2023.
- [4] S. Ghosh, "Legal technology and the Indian judiciary: Opportunities and challenges," NUJS Law Review, vol. 13, no. 2, 2020.
- [5] A. Kamath et al., "NLP challenges in Indian multilingual legal text," in Proc. ACL Workshop on Legal NLP, 2023.
- [6] T. Bench-Capon and G. Sartor, "A model of legal reasoning with cases incorporating theories and values," Artificial Intelligence, vol. 150, 2003.
- [7] K. D. Ashley, Artificial Intelligence and Legal Analytics. Cambridge University Press, 2017.
- [8] D. Katz et al., "Natural language processing in the legal domain," arXiv preprint arXiv:2212.01326, 2023.
- [9] I. Chalkidis et al., "LEGAL-BERT: The muppets straight out of law school," in Findings of EMNLP, 2020.
- [10] A. Bhattacharya et al., "Legal NLP for Indian courts: Survey and challenges," in Proc. COLING, 2022.
- [11] A. Jain et al., "Multilingual NLP for Indian legal documents," in Proc. IndicNLP Workshop, 2023.
- [12] P. Lewis et al., "Retrieval-augmented generation for knowledge-intensive NLP tasks," in Advances in NeurIPS, 2020.
- [13] S. Sharma et al., "RAG-based Indian legal advisory: Architecture and evaluation," in Proc. IEEE ICTBIG, 2025.
- [14] M. Surden, "Artificial intelligence and law: An overview," Georgia State University Law Review, vol. 35, 2019.
- [15] V. Patil et al., "Agentic AI in action: Multi-agent architectures for legal workflows," in Proc. IEEE NKCon, 2025.
- [16] A. Singh et al., "Agentic retrieval-augmented generation: A survey," arXiv preprint arXiv:2501.09136, 2025.
- [17] P. Narayan, "AI chatbots and legal misinformation: Risks for Indian litigants," Indian Journal of Law and Technology, vol. 19, 2023.
- [18] R. Gupta et al., "Automated FIR section classification using deep learning," in Proc. IEEE ICRITO, 2024.
- [19] S. Mahapatra and A. Rao, "Algorithmic bail prediction and bias in the Indian criminal justice system," Economic and Political Weekly, vol. 58, 2023.
- [20] A. Malik et al., "Predicting Indian Supreme Court decisions using machine learning," in Proc. JURIX, 2021.
- [21] A. Bhattacharya et al., "ILDC for CJPE: Indian legal documents corpus for court judgment prediction and explanation," in Proc. ACL, 2021.
- [22] P. Nigam et al., "Benchmarking legal judgment prediction for Indian courts," arXiv preprint arXiv:2304.05765, 2023.
- [23] S. Kapoor et al., "Cross-jurisdictional judgment prediction across Indian High Courts," in Proc. LREC-COLING, 2024.
- [24] V. Menon et al., "Abstractive summarisation of Indian Supreme Court judgments," in Proc. IJCNLP, 2022.
- [25] SpotDraft, "AI-powered contract intelligence for Indian enterprises," SpotDraft Technical Whitepaper, 2024.
- [26] A. Bhattacharya et al., "Citation network analysis of Indian constitutional jurisprudence," Artificial Intelligence and Law, vol. 31, 2023.
- [27] R. Iyer et al., "Legal NLP in Dravidian languages: Tamil and Kannada court documents," in Proc. IndicNLP, 2023.
- [28] A. Kakwani et al., "IndicNLP Suite: Monolingual corpora, evaluation benchmarks and pre-trained multilingual language models for Indian languages," in Findings of EMNLP, 2020.
- [29] P. Rani et al., "Machine translation of Indian legal text: Evaluation and challenges," in Proc. EAMT, 2023.
- [30] National Informatics Centre, "National Judicial Data Grid: Architecture and analytics," Government of India Technical Report, 2024.
- [31] Malimath Committee, "Report on Reforms of the Criminal Justice System," Ministry of Home Affairs, Government of India, 2003.
- [32] N. Doshi, "Artificial intelligence in the Indian judiciary: Constitutional concerns," Journal of Indian Law and Society, vol. 12, 2022.
- [33] R. Suresh, "Bias in AI judicial tools: A critical analysis for India," NUJS Law Review, vol. 15, 2023.
- [34] Bar Council of India, "Professional conduct rules and emerging technology," Discussion Paper, 2024.
- [35] OpenAI, "GPT-4 Technical Report," arXiv preprint arXiv:2303.08774, 2023.
- [36] B. Han et al., "Automating systematic literature reviews with retrieval-augmented generation," Applied Sciences, vol. 14, 2024.
- [37] A. Singh et al., "Enhancing AI systems with agentic workflow patterns in large language models," in Proc. IEEE AllIoT, 2024.
- [38] H. Apuri and C. Yepuri, "Design of multi-agent autonomous workflow systems," in Proc. ICAUC, 2026.

- [39] D. Moshkovich and S. Zeltyn, "Taming uncertainty via automation in agentic AI systems," in Proc. IEEE/ACM ASE, 2025.
- [40] N. Aletras et al., "Predicting judicial decisions of the European Court of Human Rights with NLP," PeerJ Computer Science, vol. 2, 2016.
- [41] J. Martinez-Gil et al., "Legal knowledge graphs for statutory cross-referencing," in Proc. EKAW, 2022.
- [42] Manupatra, "AI-assisted legal research platform," Manupatra Product Overview, 2024.
- [43] L. Sureshkumar, "Indian Kanoon: Building an open-access Indian legal database," in Proc. ICAIL, 2019.
- [44] VakilSearch, "AI-powered lawyer discovery and legal triage," VakilSearch Platform Overview, 2024.
- [45] Leegality, "e-Stamp and digital workflow infrastructure for Indian courts," Leegality Technical Overview, 2024.
- [46] Vidhi Centre for Legal Policy, "AI tools for statutory drafting quality assessment," Vidhi Research Report, 2023.
- [47] P. Bhattacharya et al., "Legal AI benchmarking: Commercial tools vs academic models on Indian legal tasks," arXiv preprint arXiv:2401.09876, 2024.
- [48] S. Bhatia, "Fundamental rights constraints on algorithmic decision-making in Indian courts," IALS Student Law Review, vol. 10, 2023.
- [49] Supreme Court of India, "Union of India v. Mohan Lal Capoor," AIR 1974 SC 87, 1973.
- [50] Ministry of Electronics and Information Technology, "Digital Personal Data Protection Act, 2023," Government of India Gazette, 2023.
- [51] Bar Council of India, "Rules of Professional Standards and Advocate-Client Privilege," Bar Council Regulations, 1975.
- [52] A. Saxena and P. Rathore, "Caste bias in Indian judicial AI: Evidence and implications," Economic and Political Weekly, vol. 59, 2024.
- [53] M. Hidayatullah, "Pluralism in Indian personal law and the challenge of AI legal advisory," Journal of the Indian Law Institute, vol. 65, 2023.
- [54] B. Strickson and B. De La Iglesia, "Legal judgement prediction for UK courts," in Proc. ACM SAC, 2020.
- [55] European Parliament, "Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence," Official Journal of the European Union, 2024.
- [56] H. Li et al., "AI agent security: Vulnerability analysis and protective measures," in Proc. IEEE ICCBD+AI, 2025.
- [57] Y. He et al., "Security of AI agents," in Proc. IEEE/ACM RAIE, 2025.