


SvedbergOpen
DISSEMINATION OF KNOWLEDGEInternational Journal of Artificial
Intelligence and Machine LearningPublisher's Home Page: <https://www.svedbergopen.com/>

Research Paper

Open Access

Agentic AI Framework for Automated Legal Document Review, Case Preparation, and Paralegal Workflow Optimization

Rajesh Karvannan

Senior Specialist Engineer karvannanrajesh7@gmail.com,

Abstract—Artificial Intelligence (AI) is heralded as a transformative technology in the legal sector. However, unlike the insurance and financial sectors, the development of proof-of-concept, minimally viable commercial AI legal technology has proven acutely difficult. This article explores this issue, introducing an overlay “Agentic AI” framework, defining Agentic AI as AI that governs a decision so complex that one or more humans cannot grasp all its significant elements, deciding and acting in one area only. It is thus purpose-built to assist one area of human knowledge work—paralegal work—through automating the repeated elements of case preparation, namely Document Review and Case Preparation. Document Review automates the identification, extraction (including via classification and redaction), and evaluation of particular elements from Contracts, Pleadings, and Filings; Document-Review-Automation is evaluated against key qualitative and quantitative elements. Case-Preparation Automation examines a complex real-world, repeated, case-preparation task—preparing a matter for greater command by the lawyer—synthesising case-law, evidence, and pleading; triangulating the evidence; spotting key issues; and executing associated task work flows. While Document-Review and Case-Preparation Automation are critically important in their own right—and are the starting point in an exploratory Agentic-AI overlay—they must be combined with significant development of related areas such as high-quality Spot Law, in particular for factual-timeline creation, and other Agentic-AI areas such as Evidence Extrapolation, for a more complete technology. The article also contributes to a wider understanding of Agentic-AI theory, growing the body of Agentic AI knowledge and stimulating associated discussion in the AI-for-Law field.

Keywords—Agentic AI, Legal Document Review, Case Preparation, Paralegal Workflow Automation, Legal Process Optimization, Intelligent Legal Systems

INTRODUCTION

AI agents have been around for decades. More recently, they have permeated into the legal domain following the successful public introduction of ChatGPT in November 2022 by OpenAI. Various implementations utilizing the OpenAI API, including Koan.ai, ChatPDF, and Wolverine, ease the burden of tedious and repetitive tasks, such as reading long documents or extracting information from documents. These implementations showcase examples of tools with an agentic persona capable of scanning, analyzing, and processing textual information faster than human counterparts when utilizing an underlying LLM model [4], [11]. Koan.ai even acts as a testbed for such services, emphasizing generating a deeper synthesis of evidence [26].

Roads to provide new solutions remain open for further development. The automation of legal case preparation and the consolidation of paralegal processes serve as directions for these solutions. These solutions can be developed for all legal domains, including commercial, administrative, civil, family, constitutional, and international law. Assistance may come in many ways: advising and generating memos for lawyers or paralegals; generating and serving arguments and grounds for administrative decisions; generating and serving requests for investigations in criminal cases; generating and serving powers of attorney in civil cases; and automating law bank preparation through Tableau or PowerBI. Generative AI may culminate in a tool that will emulate a lawyer, paralegal, or other workflow without prejudice and will be able to assist lawyers and paralegals in their work.

A. Overview of AI Integration in Legal Contexts

The legal sector has begun to embrace digital transformation, with technology increasingly incorporated into legal services. In recent years, a broad set of services have been digitized or adequately outsourced. Services classified as highly repetitive and manual, such as back-office operations and some paralegal tasks, present high opportunities for automation using traditional techniques and tools such as Workflow Automation Systems or Robotic Process Automation (RPA). However, these tools do not exhibit autonomy: execution patterns must be constantly restated, and business logic is usually hard-coded.

With advances in Natural Language Processing (NLP) and Near-Optical Video Recognition, or Optical Character Recognition (OCR) success rates have improved [18]. With an agentic approach and trained models with specific knowledge in defined domains, automatic document review and synthesis of initial legal reasoning are possible with a satisfactory level of quality [2], [21]. The potential of pre-processing for larger models such as ChatGPT can obtain robust outputs for multi-sets of cases in sequence [4]. NLP tools are used to guide a lawyer in the decision-making process, as an intelligent assistant, highlighting risks and guidelines to speed up complex processes [26].

THEORETICAL FOUNDATION

The analysis and evaluation of the use of an agent of automation within complex legal workflows necessitates the introduction of a conceptual framework. First a general discussion concerning the interaction between law and technology, supported by specific theories of law-making and law-enforcement, together with definitions of such concepts as autonomy, agency, and responsibility, indicates the system-level feasibility of employing agentic automation to increase the efficiency, reliability, and consistency of strictly-defined legal processes [4], [11], as summarized conceptually in Fig. 1.




Fig 1. Conceptual Quality Profile of Agentic AI for Legal Process Automation

The implementation of agentic systems within specific legal contexts then follows logically. The role that agentic automation could play in performing paralegal-type work, tasks foundational for later stages of case preparation, is explored first. It is argued that during such activities, agentic systems are synthetic decision-makers, responsible for ensuring the successful functioning of arbitrary legally-constituted systems of decisions-binding. An answer to the question of how these situations induce high trust in the functioning of agentic automation is provided. The next step considers whether agentic automation is applicable to the complex problem of legal document review. It is shown that, under well-defined conditions, these systems can increase the speed and accuracy of review for quality and compliance, addressing at least one specific issue raised during such practical applications of legal AI.

A. Agentic AI in Legal Praxis

Lawyers are information and knowledge workers coping in environments that require consistent quality with expert-level accuracy and judiciousness. However, maintaining that standard requires significant and costly resources. AI has been explored for the past three decades to fit the needs of legal practice by using technologies intended to automate paralegal work to make it consistent and cost-efficient [4], [11], [18]. Although turning those technologies into production solutions has mostly succeeded, clients have been slow to adopt the resulting systems.

AI offerings to far reduce and, in many cases, eliminate the need for paralegal support have only recently emerged. Offering a level of convenience and supporting expert-level quality, these systems encourage use in regular practice but do so at a higher cost than such work is typically charged at. The most valuable of these systems also preserve the records and checklists that need to be adhered to when preparing for the execution stage of a case. The new systems are novel for using deep learning-enabled natural language processing capabilities specifically to enable complete case preparation followed by a natural language action-taking conversational interface [2], [20], [26].

SYSTEM ARCHITECTURE

A holistic view of the framework reveals a system of systems encompassing different forms of AI under a common umbrella. Documents, exchanges, and exchanges with the agent feed data and results into the automated workflows for legal document review, case preparation, and paralegal task optimization. Interfacing components enable human actors to benefit from the various capabilities afforded by this integration of systems while maintaining the necessary interoperability. The modular, adaptable nature of these systems supports the introduction of additional capabilities when required, as outlined in Table I. The architecture includes explicit provisions for data governance and compliance with regulatory or organizational requirements. It assembles a comprehensive description of the data sources that underpin the various processes—legal documents, source data for case preparation or paralegal support, and control artefacts. The focus here is on those dimensions of data governance that directly contribute to the correctness, trustworthiness, and risk management of the system. Compliance with specific privacy or information security policies, the execution of retention schedules, or the implementation of audit trails must also be addressed, but may differ in detail from one context to another. A thorough analysis of these aspects therefore lies beyond the scope of the current discussion [18].

TABLE I. CORE COMPONENTS OF THE AGENTIC AI LEGAL FRAMEWORK

Component	Primary Function	Expected Output
Document Ingestion	Collect legal documents and case files	Processable document set
Document Review	Classify and identify relevant content	Relevant document subset
Information Extraction	Extract entities, clauses, dates, and facts	Structured legal data
Case Preparation	Organize evidence and case information	Case preparation package
Evidence Synthesis	Combine related facts and evidence	Case summary / timeline
Human Review	Validate AI-generated findings	Approved legal output

Supporting accuracy, reliability, and compliance with internal or regulatory standards relies on careful definition of the data elements acquired from the documents marked for extraction, the methods and sources employed for extraction, the schema used to specify the normalization of data into an internal representation, and mechanisms that promote high-quality metadata and entity recognition. A language-agnostic approach to the translation of documents of diverse types into a normalized textual representation conducive to legal document review is adapted from existing research [2], [3], and is illustrated in the end-to-end pipeline of Fig. 2, formalized by the schema-based normalization in Equation (1).




Fig 2. End-to-End Agentic AI Framework for Automated Legal Workflow Orchestration

Schema-Based Data Normalization

$$\mathbf{z}_i = N_S(\mathbf{x}_i), \quad \mathbf{z}_i \in Z_c \quad (1)$$

A. Data Governance and Compliance

Legal practitioners disproportionately concentrate on risk mitigation and compliance with regulatory requirements that govern the handling of information digitally engaged as participant in financial markets, yet these aspects are of common operational concern for all businesses and industries. Non-public client information is protected in South Africa through the protection of personal information act and the protection of access to information act, and lawyers are subject to ongoing professional obligations including obligations of confidentiality, record keeping, competency, and security. Data privacy sophistication throughout the practice and client relationship life-cycle is, therefore, critical in addressing both the risk and opportunity presented by advancements in data science and Artificial Intelligence that can be provisioned as a service by third-party service providers [4], as summarized in Fig. 3.




Fig 3. Governance and Compliance Control Emphasis Across the Agentic Legal Pipeline

The data governance measures undertaken to protect non-public information during the review, assessment, extraction, and synthesis of public legal documents for case preparation and knowledge provision are therefore isomorphic to those necessary or warranted for the considered application, but apply similarly to mandated document access and disclosure controls in the supported case preparation process. Digital paralegal support systems employed to facilitate the intimacy expected of paralegals and junior legal analysts during the normal course of business suggest, as a second order consequence, unintended safety risks and potential for liability and damages in the form of complaint. Agentic AI systems drawn upon to assist case preparation therefore augment clients' own paralegal and junior legal analyst resources and these safety risks or liability concerns are not present, consistent with the control structure detailed in Table II.

TABLE II. DATA GOVERNANCE AND COMPLIANCE CONTROLS

Governance Area	Proposed Control	Purpose
Client Information	Access restrictions	Maintain confidentiality
Legal Documents	Controlled ingestion	Secure data processing
Metadata	Validation and normalization	Improve data quality
Auditability	Processing logs	Support traceability
Data Retention	Retention policies	Meet organizational requirements
AI Output	Human verification	Reduce legal and operational risk

AUTOMATED DOCUMENT REVIEW

Numerous legal documents are repetitively processed during the pre-litigation stage of many litigation and dispute resolution scenarios. In the case of litigation, numerous documents will be screened for possible

relevance to a specifically defined question. The first step of a conventional document review is often termed the ‘categorisation’ stage. This typically entails a binary yes/no assessment of relevance to a specific issue by a paralegal or junior lawyer, as formalized in the classification rule of Equation (2). Each document is assessed as relevant or not. Technical solutions have been developed to assist with this task in respect of large datasets [1], [6], [8], [16], [17], [25]. After documents have been categorised, the next logical step is extraction of information that goes beyond a simplistic yes/no decision. Depending on the precise nature of the evidence, this may also require phrasing the information in a way that best reflects the original document. For example, if the document is an email and the new information relates to an opinion or sentiment, the extracted information may be best presented in the voice of the writer. Multiple technical solutions exist for this task as well [12], [15], [19], as depicted conceptually in Fig. 4.




Fig 4. Conceptual Scaling of Conventional and Agentic-Assisted Legal Document Review

Binary Legal-Document Relevance Classification

$$r_i = \begin{cases} 1, & d_i \text{ is relevant to the specified legal issue,} \\ 0, & d_i \text{ is not relevant.} \end{cases} \quad (2)$$

The results of these two phases can then be synthesised into a user-friendly format such as a timeline, tree diagram, narrative, or other suitable structure to highlight strengths and weaknesses in a matter. Different technical solutions have been developed for the synthesis stage as well [20], [22]. Given that there are already multiple viable technical solutions for each distinct phase of the process, summarized across the stages in Table III, the challenge lies in integrating them into an effective ‘end-to-end’ automated solution that operates under human supervision and retains human-in-the-loop governance and accountability [11].

TABLE III. AUTOMATED LEGAL DOCUMENT REVIEW STAGES

Stage	Main Activity	Result
Categorization	Determine document relevance	Relevant / Not Relevant
Extraction	Identify important information	Facts, entities, clauses
Normalization	Convert extracted information	Standardized representation
Synthesis	Combine related information	Timeline or narrative
Validation	Human review of generated output	Verified result

A. Extraction and Normalization

Supporting the automated review system is an extraction and normalization application. Instantiation of the legal domain ontology is necessary to understand the informational content contained within a particular document. Information required to perform this instantiation may be extracted from any type of document, including judicial rulings, court orders, motions and affidavits, and sensitive documents represented in legal contracts and certifications [10], [21].

Due to regulation and the nature of legal processes, only a small fraction of cases are selected for analysis. This small number of cases provides an opportunity to extract relevant information for instantiating and extending the legal domain ontology. Of particular importance are judicial rulings at the different layers of

the justice process. The analysis of these texts also makes it possible to analyze recurrent judicial reasoning and to assess the potential impact of such reasoning on the population. Data in non-textual formats is also processed, extracting relevant entities and relations and normalizing the data into the knowledge graph [13], [14], as formalized in the entity-relation representation of Equation (3) and summarized in Table IV.

Legal Knowledge-Graph Representation

$$\mathcal{G} = (\mathcal{V}, \mathcal{R}), \quad \mathcal{V} = \text{Ent}(\mathcal{D}), \quad \mathcal{R} = \text{Rel}(\mathcal{D}) \quad (3)$$

TABLE IV. EXTRACTION AND NORMALIZATION OF LEGAL INFORMATION

Legal Information	Extraction Example	Normalized Representation
Parties	Person / organization names	Legal entity
Dates	Filing or event dates	Standard date
Legal Events	Hearing, filing, judgment	Event record
Obligations	Contractual duties	Obligation entity
Evidence	Statements and records	Evidence object
Relationships	Links among parties/events	Knowledge graph relation

CASE PREPARATION AUTOMATION

Paralegal support for case preparation accounts for a significant share of billable hours for many legal practices, and this investment in attorney support often occurs at the expense of smaller clients and matters where offer mapping is less needed, yet risky facts and disclosures remain highly relevant. The Agentic AI Framework introduces automation to paralegal case preparation by integrating all case files, including pleadings, motions, discovery interrogatories, depositions, productions, and requests [5], [7], [9]. Parsing these documents extracts “what, when, where, and” disclosure, and outcome-type evidence during document review, allowing for the seamless addition of discovered database evidence. The Candidate Evidence Database serves as a centralized pool of extracted evidence, structured as in Equation (4), which can be filtered by attorney-defined search terms, as formalized in Equation (5), for report generation to guide the preparation phase, whether during offer mapping or readiness for summary judgment [23], [24].

Structured Evidence Extraction

$$\mathbf{x}_i = P(d_i) = [w_i, \tau_i, \ell_i, \delta_i, o_i] \quad (4)$$

Attorney-Defined Candidate Evidence Filtering

$$\mathcal{C}_q = \{e \in \mathcal{C} : M(e, q) = 1\} \quad (5)$$

Next, an automated chronological timeline is produced, as ordered by the temporal sequence in Equation (6). Ordering the discovered evidence within chronology enables the paralegal – with attorney direction – to assess liability-prerequisite evidence that establishes evidence-and-fact-issue correlations for either party, as well as decision-tree milestone events, calendrical impact on damages evidence inaccuracies, and gaps in evidence related to these aspects. The combined prevention-mitigation process automates spotting what the defendant must prove and creates a sequence of key dates, together forming a ready expansion of a checklist of offers, counter-offers, settlement recommendations, or settlement documents. With the case timeline, the paralegal can now assess the court’s view of the risk each party faces based on the timeline vis-a-vis claims or defenses and can begin populating a memorandum-style document to aid the attorney’s more detailed case-mapping discussions and preparations, drawing on the evidence-fact-issue associations defined in Equation (7) and the comparative effort profile shown in Fig. 5 [5], [9].

Automated Chronological Evidence Timeline

$$\mathcal{T} = (e_{(1)}, e_{(2)}, \dots, e_{(m)}), \quad t(e_{(1)}) \leq t(e_{(2)}) \leq \dots \leq t(e_{(m)}) \quad (6)$$

Evidence–Fact–Issue Association

$$\mathcal{L}_p = \{(e_j, f_k) : A_p(e_j, f_k) = 1\} \quad (7)$$




Fig 5. Relative Manual Effort Across Agentic Legal Automation Tasks

A. Evidence Synthesis

Creation of an agentic AI-based intelligence system to automate paralegal tasks, such as reviewing contracts and other legal documents, preparing cases for court, and generating letter templates. Within case preparation, the synthesized case summary organizes matter details from paralegal and central repository content, highlighting factual evidence. A readability-focused design enhances human comprehension of the AI-generated substance [20], [22].

Case Summary

Central repository content and material realized during document review—such as contracts, discovery responses, and depositions—must converge in a synthesized summary to facilitate human comprehension and highlight main facts essential for full knowledge of the dispute. The output's date-specific publication must ensure freshness for the period preceding case commencement, as formalized by the synthesis function in Equation (8).

Synthesized Case Summary

$$\mathcal{S} = \text{Synth}(\mathcal{C}_{\text{repo}} \cup \mathcal{D}_{\text{review}}) \quad (8)$$

Synthesized information integrates True Layer API Bank State capabilities, generating request letters via approved template content, thereby proving within this central repository that formal questions, inquiries, and missing standard documents for human review are duly addressed [24].

Agentic AI systems and intelligent agents should replicate humans' behavior in evaluating and assessing data-processing pipelines. Such evaluation must assess both input data and the final product. In evidence constellations, each intelligence agent can verify the correctness of information, receiving feedback based on its decision. Misclassification, as detected by the indicator function in Equation (9), can trigger a semi-supervised learning session or a trained model [16], [25]. The resulting evidence elements are organized as shown in Table V, and the corresponding automation-suitability and human-oversight requirements for each task are summarized in Fig. 6.

Misclassification Detection for Feedback

$$m_j = \mathbb{1}(\hat{c}_j \neq c_j) \quad (9)$$

TABLE V. CASE PREPARATION EVIDENCE STRUCTURE

Evidence Element	Example Information	Role in Case Preparation
What	Event or action	Defines case fact
When	Date / time	Builds chronology
Where	Location	Provides contextual evidence
Disclosure	Statement or submitted information	Supports factual assessment
Outcome	Result of event/action	Supports claim evaluation

Evidence Gap	Missing supporting material	Identifies further work
--------------	-----------------------------	-------------------------




Fig 6. Automation Suitability and Human Oversight Requirements for Legal AI Tasks

CONCLUSION

The system integrates state-of-the-art machine learning technologies to automate legal document review, case preparation, paralegal workflow, and manual task assistance. Most of the implementations function in closed-loop, agentic configurations. The cornerstone of these technologies is document analysis, which supports multiple application domains and automates much of the tacit expertise involved in document processing. Document analysis includes examination, extraction, entity-extraction, clause-identification, metadata-tagging, normalization, and quality-checking capabilities. The normalization capability maps legal concepts and contractual obligations to a canonical representation that enables holistic clauses or condition-and-termination analyses [2], [3], [12].

The automated-document-review functions reside within a practical realisation of an agentic AI framework for the legal domain, supporting actors in the left-hand-digits of the CASE paralegal framework, automating mundane activities in product warranty and life-science patent law, and providing preparation briefs for deeper agentic consideration. Expert users evaluated the document-analysis capabilities on a diverse range of a thousand-plus complex documents for business and trade law, including property sales contracts, share-purchase agreements, business-loan agreements, tax-ruling requests, consumer-product warranties, life-science research grants, and more. Further implementation effort is required to automate evidence synthesis and assist paralegal preparation-of-trial briefs, integrating all functions into an agentic framework capable of CASE automation across all domains of legal practice, easing the burdens and risks associated with these higher-volume and -value work opportunities [5], [7], [9], [23].

A. Final Thoughts and Future Directions

The Agentic AI Framework is a step toward identifying and integrating the full agentic dimension of AI to address task completion requirements across the legal document lifecycle. The focus is on automating the preparation of the information required for managing legal proceedings and the associated review of the documentation generated by defined workflows. Within this context, the implementation of Automated Document Review, Case Preparation Automation, Evidence Synthesis, and Paralegal Workflow Optimization demonstrates a significant reduction in the time required for document review and case preparation dependent on patient interviews.

Emerging generative AI services such as ChatGPT, Google Bard, and Claude share many of the identified indicators of agentic AI [4], [11], [26]. As these service providers negotiate with regulators and develop processes to identify and address legal compliance, governance, and risk exposure, the focus will shift to use case requirements. Several service providers deliver prompt-based automation capabilities fit for zero to low-code environments designed to complement development skills on capability enhancement and service orchestration across hybrid environments. The emphasis is on delivering the agentic dimension of AI systems aligned with personnel time and resource requirements in preparation for upcoming matters, classes, and placements.

REFERENCES

1. Aletras, N., Tsarapatsanis, D., Preotiuc-Pietro, D., & Lamos, V. (2016). Predicting judicial decisions of the European Court of Human Rights: A natural language processing perspective. *PeerJ Computer Science*, 2, e93. <https://doi.org/10.7717/peerj-cs.93>

2. Chalkidis, I., Fergadiotis, M., Malakasiotis, P., Aletras, N., & Androutsopoulos, I. (2020). LEGAL-BERT: The Muppets straight out of law school. Findings of the Association for Computational Linguistics: EMNLP 2020, 2898–2904. <https://doi.org/10.18653/v1/2020.findings-emnlp.261>
3. Chalkidis, I., Jana, A., Hartung, D., Bommarito, M., Androutsopoulos, I., Katz, D. M., & Aletras, N. (2022). LexGLUE: A benchmark dataset for legal language understanding in English. Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics, 4310–4330. <https://doi.org/10.18653/v1/2022.acl-long.297>
4. Dehghani, F., Dehghani, R., Naderzadeh Ardebili, Y., & Rahnamayan, S. (2025). Large language models in legal systems: A survey. Humanities and Social Sciences Communications, 12, 1977. <https://doi.org/10.1057/s41599-025-05924-3>
5. Deng, C., Mao, K., & Dou, Z. (2024). Learning interpretable legal case retrieval via knowledge-guided case reformulation. Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, 1253–1265. <https://doi.org/10.18653/v1/2024.emnlp-main.73>
6. Feng, Y., Li, C., & Ng, V. (2022). Legal judgment prediction: A survey of the state of the art. Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, 5461–5469. <https://doi.org/10.24963/ijcai.2022/765>
7. Feng, Y., Li, C., & Ng, V. (2024). Legal case retrieval: A survey of the state of the art. Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics, 6472–6485. <https://doi.org/10.18653/v1/2024.acl-long.350>
8. Gan, L., Kuang, K., Yang, Y., & Wu, F. (2021). Judgment prediction via injecting legal knowledge into neural networks. Proceedings of the AAAI Conference on Artificial Intelligence, 35(14), 12866–12874. <https://doi.org/10.1609/aaai.v35i14.17522>
9. Gao, C., Xiao, C., Liu, Z., Chen, H., Liu, Z., & Sun, M. (2024). Enhancing legal case retrieval via scaling high-quality synthetic query-candidate pairs. Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, 7086–7100. <https://doi.org/10.18653/v1/2024.emnlp-main.402>
10. Habernal, I., Faber, D., Recchia, N., Bretthauer, S., Gurevych, I., Spiecker genannt Döhmann, I., & Burchard, C. (2024). Mining legal arguments in court decisions. Artificial Intelligence and Law, 32, 1–38. <https://doi.org/10.1007/s10506-023-09361-y>
11. He, C., Hu, H., Li, Y., Zhang, H., & Zhang, Q. (2026). A survey of large language models for legal tasks: Progress, prospects and challenges. Computer Science Review, 60, 100906 <https://doi.org/10.1016/j.cosrev.2026.100906>
12. Joshi, S., Balaji, S., Garimella, A., & Varma, V. (2022). Graph-based keyword planning for legal clause generation from topics. Proceedings of the Natural Legal Language Processing Workshop 2022, 276–286. <https://doi.org/10.18653/v1/2022.nllp-1.26>
13. Kim, M.-Y., Rabelo, J., Okeke, K., & Goebel, R. (2022). Legal information retrieval and entailment based on BM25, transformer and semantic thesaurus methods. The Review of Socionetwork Strategies, 16, 157–174. <https://doi.org/10.1007/s12626-022-00103-1>
14. Kim, M.-Y., Rabelo, J., Babiker, H. K. B., Rahman, M. A., & Goebel, R. (2024). Legal information retrieval and entailment using transformer-based approaches. The Review of Socionetwork Strategies, 18, 101–121. <https://doi.org/10.1007/s12626-023-00153-z>
15. Lippi, M., Pałka, P., Contissa, G., Lagioia, F., Micklitz, H.-W., Sartor, G., & Torroni, P. (2019). CLAUDETTE: An automated detector of potentially unfair clauses in online terms of service. Artificial Intelligence and Law, 27, 117–139. <https://doi.org/10.1007/s10506-019-09243-2>
16. Medvedeva, M., McBride, P. (2023). Legal judgment prediction: If you are going to do it, do it right. Proceedings of the Natural Legal Language Processing Workshop 2023, 73–84. <https://doi.org/10.18653/v1/2023.nllp-1.9>
17. Medvedeva, M., Vols, M., & Wieling, M. (2020). Using machine learning to predict decisions of the European Court of Human Rights. Artificial Intelligence and Law, 28, 237–266. <https://doi.org/10.1007/s10506-019-09255-y>
18. Quevedo, E., Cerny, T., Rodriguez, A., Rivas, P., Yero, J., Sooksatra, K., Zhakubayev, A., & Taibi, D. (2024). Legal natural language processing from 2015 to 2022: A comprehensive systematic mapping study of advances and applications. IEEE Access, 12, 145286–145317 <https://doi.org/10.1109/ACCESS.2023.3333946>
19. Ruggeri, F., Lagioia, F., Lippi, M., & Torroni, P. (2022). Detecting and explaining unfairness in consumer contracts through memory networks. Artificial Intelligence and Law, 30, 59–92. <https://doi.org/10.1007/s10506-021-09288-2>
20. Sadikot, T., & Jain, S. (2025). Legal document summarization: A zero-shot modular agentic workflow approach. Proceedings of the 1st Workshop on NLP for Empowering Justice, 29–40. <https://doi.org/10.18653/v1/2025.justnlp-main.5>
21. Singh, A., Joshi, A., Jiang, J., & Paik, H.-Y. (2025). A survey of classification tasks and approaches for legal contracts. Artificial Intelligence Review, 58, 380. <https://doi.org/10.1007/s10462-025-11359-8>

22. Song, Y., Qin, Y., Huang, R., Chen, Y., & Lin, C. (2025). Legal text summarization via judicial syllogism with large language models. *Journal of King Saud University Computer and Information Sciences*, 37, 111. <https://doi.org/10.1007/s44443-025-00113-3>
23. Tao, M., Hu, J., Zhou, X., Hu, W., Cheng, J., Cao, Y., Luo, Z., & Geng, G. (2026). When rules learn: A self-evolving agent for legal case retrieval. *Findings of the Association for Computational Linguistics: ACL 2026*, 43348–43365. <https://doi.org/10.18653/v1/2026.findings-acl.2152>
24. Wang, S. H., Zubkov, M., Fan, K., Harrell, S., Sun, Y., Chen, W., Plesner, A., & Wattenhofer, R. (2025). ACORD: An expert-annotated retrieval dataset for legal contract drafting. *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics*, 24739–24762. <https://doi.org/10.18653/v1/2025.acl-long.1206>
25. Wu, Y., Zhou, S., Liu, Y., Lu, W., Liu, X., Zhang, Y., Sun, C., Wu, F., & Kuang, K. (2023). Precedent-enhanced legal judgment prediction with LLM and domain-model collaboration. *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 12060–12075. <https://doi.org/10.18653/v1/2023.emnlp-main.740>
26. Yang, X., Wang, Z., Wang, Q., Wei, K., Zhang, K., & Shi, J. (2024). Large language models for automated Q&A involving legal documents: A survey on algorithms, frameworks and applications. *International Journal of Web Information Systems*, 20(4), 413–435. <https://doi.org/10.1108/IJWIS-12-2023-025>