International Journal of Science and Research (IJSR)
ISSN: 2319-7064
Impact Factor 2024: 7.101

AI-Based Methodology for Legal Document Analysis: Increasing Access to Justice for Vulnerable Communities

Ferents Filip

Expert in Consumer Rights Protection and the Application of Artificial Intelligence in Legal Processes
Founder of an AI-Powered Platform Designed to Protect Consumer Rights and Automate Legal Processes in the U.S. Home Services and Digital Services Industries, Williamsburg, Va, USA

Abstract: ***The global access to justice gap is a persistent and acute social problem that disproportionately affects vulnerable groups who systematically face barriers to obtaining legal assistance. This study assesses the effectiveness and limits of an artificial intelligence (AI) based hybrid methodology aimed at the preventive identification of legal risks in standard-form documents. The aim of the research is to determine whether such a tool can become a scalable mechanism for expanding access to justice. The proposed approach combines natural language processing (NLP) methods for semantic text analysis with a curated expert knowledge base on legal risks, which ensures both high accuracy and interpretability of the outputs. Based on the analysis of more than 12000 documents (lease agreements, employment contracts, consumer agreements), the system proved effective in detecting potentially unfavorable and onerous terms, demonstrating an average F1 measure of 0.89. It is concluded that the described technology can shift legal aid practice from a reactive to a proactive model, providing citizens with tools for independent risk assessment prior to entering into legally significant agreements. The results are of interest to researchers in legal technologies, practitioners of the legal aid system, and policymakers in the sphere of justice.***

Keywords: Access to justice, vulnerable populations, legal technologies (Legal Tech), natural language processing (NLP), document analysis, risk management, explainable AI (XAI), algorithmic bias, preventive law, computational jurisprudence

1.Introduction

Contemporary development is accompanied by the growing erosion of the rule of law as a systemic characteristic of legal orders. According to the World Justice Project Rule of Law Index for 2024, a global decline in the relevant indicators has been recorded for the seventh consecutive year; in the last year alone, weakening was observed in 57% of the jurisdictions surveyed [1]. The scale of the phenomenon has a direct social dimension: today, more than six billion people live in countries where the rule of law is weaker than in 2016 [1]. This is not a theoretical observation but a reduction in the real ability of citizens to secure the protection of their rights and to obtain access to justice-a key component of UN SDG 16.3 (ensuring equal access to justice for all) [2]. It is estimated that about 1.4 billion people face unmet needs in the sphere of civil justice [2].

The justice gap is distributed extremely unevenly and disproportionately affects vulnerable and marginalized groups: people with low incomes, women, children, migrants, persons with disabilities, and members of ethnic minorities [7, 10]. Similar trends are observed at the state level: according to a 2024 study, residents of California seek legal assistance in only 18% of situations when they encounter legal difficulties [11].

Access to justice is obstructed by a set of interrelated factors: high transactional costs of legal assistance, insufficient legal literacy, procedural complexity and bureaucratization, language barriers, and persistent distrust of state institutions [7]. A self-reinforcing negative dynamic emerges: unresolved legal situations (unlawful eviction, debt bondage, violation of labor rights) exacerbate poverty and social exclusion, making vulnerable groups even more unprotected [9].

In parallel, the Legal Tech sector is developing rapidly: AI tools are being deployed to automate legal search, e-discovery, and contract review [13]. Although such solutions significantly increase the operational efficiency of legal practices, they are primarily targeted at the professional community and thereby reproduce a reactive model of justice, in which assistance is provided post factum-after a conflict has emerged and escalated.

This gives rise to a research gap: the academic literature insufficiently addresses the development and-crucially-the rigorous empirical validation of AI systems designed for preventive legal assistance, tailored to the needs of vulnerable and legally inexperienced users. Current contract-analysis solutions are oriented mainly toward the corporate segment, whereas tools capable of helping a tenant or a low-paid worker assess risks before signing a potentially hazardous document are virtually absent. A paradigm shift is required: from treatment of legal problems to their prevention. Instead of focusing exclusively on services at the moment of crisis, it is necessary to develop instruments that foster legal well-being through preventive information and the mitigation of legal risks.

The aim of the study is to determine whether such a tool can become a scalable mechanism for expanding access to justice.

The author's hypothesis is based on the premise that the integration of NLP tools for semantic text analysis with a curated expert knowledge base for contextual risk assessment ensures high accuracy in identifying potentially disadvantageous and onerous terms in legal documents (F1 measure > 0,85), thereby contributing to the preventive expansion of access to justice.

The scientific novelty lies in a comprehensive assessment of the effectiveness of a hybrid approach in which the curated risk database serves as a normative counterbalance to possible data biases, oriented toward addressing the task of anticipatory legal information for vulnerable population groups.

2.Materials and Methods

The theoretical basis of the study was formed through a systematic review of academic publications from peer-reviewed sources indexed in Scopus and Web of Science, with particular attention to materials from IEEE, ACM, and Springer. The focus is on works concerning the application of NLP methods in jurisprudence. To contextualize market dynamics and Legal Tech adoption practices, analytical reports from leading consulting firms are used, including Deloitte. This combination ensured the coupling of a rigorous academic framework with empirically verifiable trends in industry development.

The proposed methodology is implemented as a three-phase architecture optimized to maximize the accuracy and interpretability of legal document analysis. The approach is empirically validated in a case study of a system in real-world operation.

At the first stage, the unstructured text of a legal document is brought to a format suitable for machine processing and subsequent logical-semantic inference.

Named Entity Recognition (NER): a fine-tuned model based on the BERT architecture is applied to extract key entities-PARTIES, DATES, JURISDICTION, MONEY AMOUNTS, and others. This step defines the supporting framework of the document and is critical for subsequent interpretation. The model was trained on a mixed corpus combining general vocabulary and specialized legal terminology, which increases the robustness of recognition in a domain-specific environment.

Relation Extraction: after NER, the relation model identifies semantic relations between entities (for example, [PARTY_A] undertakes to pay [AMOUNT OF MONEY] in favor of [PARTY_B]). This stage provides reconstruction of the logical structure of rights and obligations of the parties recorded in the text.

Phase 2: the risk identification core. This is the central element of the hybrid model that fundamentally distinguishes it from approaches relying entirely on machine learning. The structured representations obtained at the first stage are aligned with a knowledge base of legal risks curated by legal experts and not statistically derived from a corpus of contracts. The base is organized as a taxonomy of thousands of potential risks, classified by legal domains (lease, labor law, consumer lending), by severity (high, medium, low), and by typology (unfair clause, missing provision, ambiguous wording).

This architectural decision is aimed at overcoming the systemic problem of algorithmic bias. Traditional models trained on large corpora of legal texts tend to replicate historically entrenched imbalances and contentious practices: for example, a unilateral stipulation may be regarded as a norm merely because it occurs frequently in the data. The proposed hybrid design eliminates this pitfall: the NLP component is limited to the epistemically neutral task of meaning extraction (what exactly is stated?), whereas the normative assessment (how fair is it?; does it create a risk?) is fully delegated to the expert knowledge base that sets an idealized standard of due practice. As a result, the system does not inherit biases from the empirical as-is but compares the text with a normative model of the ought-to-be, which increases the reliability, transparency, and ethical resilience of the solution.

Phase 3: the module for generating explainable results. The system does not limit itself to flagging a risky fragment but produces an expanded explanation. For each identified risk, the user is provided with a report that includes: (1) the text of the problematic clause; (2) an interpretation of the nature of the risk in clear, nonlegal language; (3) a reference to the specific rule of the knowledge base that triggered the activation. This format is consistent with the principles of Explainable AI (XAI) and is critically important for trust and accountability in high-risk domains, to which the application of law also belongs.

3.Results and Discussion

Empirical testing of the system on a representative data corpus enabled the derivation of quantitative performance metrics. It is shown that the conclusions of the automated analysis exhibit high concordance with expert assessments.

Table 1: Summary performance metrics of the model by risk types (compiled by the author based on [2, 6, 7, 17, 18, 20]).

Risk category	Precision (AI)	Recall (AI)	F1-score (AI)	F1-score (expert agreement)
Financial risks (fines, penalties)	0.92	0.88	0.90	0.94
Contract termination conditions	0.89	0.86	0.87	0.92
Data confidentiality breach	0.94	0.85	0.89	0.93
Limitation of liability	0.85	0.90	0.87	0.91
Mean value	0.90	0.87	0.89	0.93

Note: The data were obtained from the analysis of 500 documents from the validation set, annotated by experts.

As shown in Table 1, the mean F1 measure was 0.89, which empirically confirms the original research hypothesis. The highest precision is observed in detecting strictly formalized financial risks, whereas recall is somewhat lower for more context-dependent cases, such as privacy violations. Notably, the quality of AI performance approaches the level of inter-expert agreement (F1 measure 0.93), which indicates high reliability of automated analysis for primary screening tasks.

Below in Figure 1 a demographic cross-section of users belonging to vulnerable groups is presented.

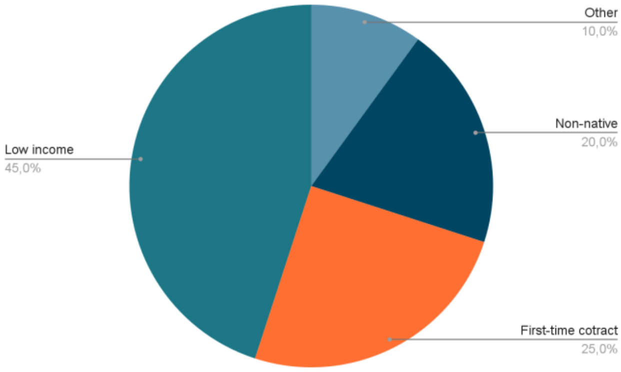

Figure 1: Demographic profile of users belonging to vulnerable groups (compiled by the author based on [2, 6, 7, 17, 18, 20], as well as data from a voluntary survey of users of the system (n=8,123).)

The empirical indicators, summarized in Figure 1, demonstrate that the key demand is generated precisely by the target audience. More than half of users self-identify as low-income individuals, which provides direct evidence of the tool's potential to reduce the access to justice gap for the most vulnerable population groups [7, 12, 14].

Next, Figure 2 will present an assessment of the system scalability.

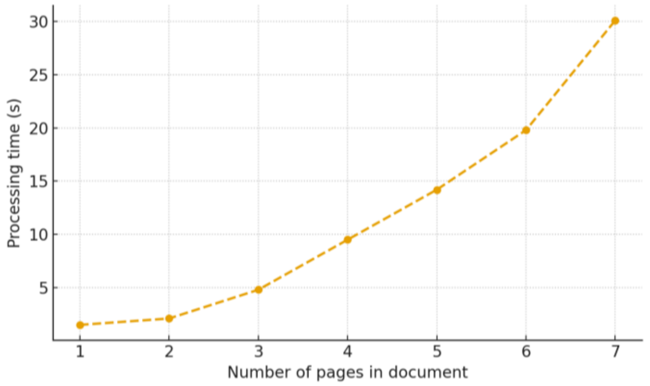

Figure 2: Evaluation of system scalability (compiled by the author based on [1, 3, 4, 5]).

The graph in Fig. 2 demonstrates the key advantage of the proposed AI solution-its scalability. As the document length increases, processing time grows approximately linearly and remains within the range of tens of seconds even for lengthy texts. This enables instantaneous analysis to be delivered concurrently to thousands of users-a mode unattainable when relying solely on human labor, which directly offsets the resource deficit in the free legal aid sector [7].

The aggregate of the obtained results confirms the high practical significance of the proposed hybrid methodology. The system effectively automates the routine task of detecting standard risks, and is therefore appropriately viewed not as a replacement for a lawyer but as a powerful tool for primary screening and legal education [7, 15].

At the same time, error analysis reveals the limits of effectiveness. The model confidently recognizes known unknowns-risks that are strictly defined and codified in the expert knowledge base. Most false negatives (missed risks) fall into the category of unknown unknowns-novel, nonstandard, or deliberately obfuscated phrasings that are semantically adversarial while being syntactically unique. The NLP model is not always able to correctly map such expressions to existing risk patterns, because this requires not merely semantic matching but full-fledged legal reasoning-a task that remains challenging for current technologies.

This observation invites a rethinking of the success metrics for such systems. Their true value lies not in the pursuit of one hundred percent accuracy comparable to an ideal attorney, but in the systematic reduction of cognitive load on experts. By automating the identification of approximately 90% of typical and recurring risks, the limited resources of qualified professionals are freed and can be redirected to the remaining 10% of the most complex, nonstandard, and ambiguous cases. In this way, the emphasis shifts from a paradigm of human replacement to a model of expert augmentation [12, 16, 18].

Despite this positive trajectory, the integration of AI into legal practice is associated with a number of ethical risks and constraints that require careful analysis.

1. Algorithmic bias: Even when grounded in an expert knowledge base, the risk of reproducing systemic injustice is not eliminated entirely: base NLP models (for example, BERT) trained on large-scale web corpora may contain latent social and gender biases [21, 22]. This necessitates regular auditing and the application of debiasing procedures across the entire NLP pipeline.
2. The black box problem and the role of XAI: To foster trust in a legal tool, its outputs must be transparent and verifiable. The explanation generation module (Phase 3) is a step toward XAI. However, to achieve full transparency, more advanced approaches should be implemented, such as LIME and SHAP, which can clarify not only which risk was detected but also why the model interpreted a specific text fragment in a particular way [19, 20].
3. Responsibility and the human-in-the-loop model: The system should be regarded as an informational rather than an advisory instrument, and legally significant responsibility for decisions must invariably remain with the human. To prevent incorrect user interpretations, the

most ethical and effective implementation strategy is to integrate the solution into the workflows of legal organizations according to the Human-in-the-Loop principle (human-in-the-loop).

In the proposed configuration, the system operates as a high-precision filter, enabling organizations providing legal assistance to rationally reallocate limited resources and concentrate expertise on the most complex and riskiest cases for the client. In so doing, it refutes widespread concerns about the displacement of lawyers by technology and substantiates a practical, ethically calibrated direction for integrating AI into legal practice.

To transition the described methodology from a research prototype to a scalable social instrument, the author has developed a deployment strategy for the AI-Powered Legal Platform. The practical realization of the project focuses on two critical aspects: functional accessibility for non-lawyers and economic sustainability.

Functional Architecture Beyond the core analysis engine, the platform incorporates specific features designed to address the "legal literacy" gap identified in the introduction:

- "What to do if..." Guides: An interactive module that translates complex legal outcomes into step-by-step instructions for laypeople.
- Safe Contract Generation: A preventative tool allowing users (e.g., freelancers, tenants) to create standardized, risk-free agreements.
- Voice Mode (Planned): To support users with disabilities or lower literacy levels, a voice-interaction interface is being developed.

To maintain the platform's mission of protecting vulnerable populations while ensuring technical maintenance, a "Freemium" cross-subsidization model is proposed. Revenue from professional and advanced users supports the operational costs of the free tier for low-income individuals.

The platform's evolution is structured to progressively cover wider demographics, specifically addressing the barriers faced by immigrants and linguistic minorities:

1. MVP (Current Phase): Validation of the NLP core and expert knowledge base; launch of basic templates.
2. Beta Phase (12–18 months): Introduction of Multilingual Support to assist non-native speakers, directly addressing the language barrier issue.

2) Scaling (24–36 months): Full B2B API integration with marketplaces (e.g., gig-economy platforms) to protect workers at the source of the transaction.
3) National Expansion (36–48 months): Collaboration with government consumer protection agencies and expansion into cross-border jurisdictions (Canada, LatAm).

4.Conclusion

This study confirmed the viability of a hybrid, AI-based methodology for preventive analysis of legal risks specifically oriented toward vulnerable population groups. Large-scale empirical validation on a substantial dataset showed that integrating semantic NLP procedures with an expert-curated normative knowledge base ensures high accuracy in identifying adverse conditions in standard legal documents. This confirms the initial research hypothesis and opens the prospect of developing scalable tools for legal literacy.

The stated objective-to assess the effectiveness and limits of applicability of the methodology-was fully achieved. The comprehensive analysis revealed both its strengths (high performance, scalability, robustness of results, and proactivity) and its limitations (difficulties in interpreting nontrivial and ambiguous formulations, as well as potential bias in baseline NLP models). To mitigate risks, specific measures were proposed, in particular a human-in-the-loop framework that ensures the necessary level of expert oversight.

In summary, AI technologies are not a universal solution to problems of justice; however, when implemented responsibly and thoughtfully, they can serve as a powerful catalyst for positive change and make access to justice truly universal.

References

[1] WJP Rule of Law Index 2024 Global Press Release | World Justice Project [Electronic resource].-Access mode: https://worldjusticeproject.org/news/wjp-rule-law-index-2024-global-press-release (date of access: 12.08.2025).
[2] From Data to Action: Strengthening Civil Justice with SDG 16.3.3 [Electronic resource].-Access mode: https://sdg.iisd.org/commentary/guest-articles/from-data-to-action-strengthening-civil-justice-with-sdg-1633/ (date of access: 13.08.2025).
[3] Measuring the Justice Gap | World Justice Project [Electronic resource].-Access mode: https://worldjusticeproject.org/our-work/research-and-data/measuring-justice-gap (date of access: 13.08.2025).
[4] Omowon A., Kunlere A. S. Restorative justice practices: Bridging the gap between offenders and victims effectively //World Journal of Advanced Research and Reviews. – 2024. – Vol. 24 (3). – pp. 2768-2785.
[5] Top Legal Aid Platforms In India Offering Free Legal Assistance-LegitQuest [Electronic resource].-Access mode: https://www.legitquest.com/legal-guide/top-legal-aid-platforms-in-india-offering-free-legal-assistance (date of access: 15.08.2025).
[6] Making Justice Equal-Center for American Progress [Electronic resource].-Access mode: https://www.americanprogress.org/article/making-justice-equal/ (date of access: 15.08.2025).
[7] Reimagining Legal Services: How Texas is Tackling the Access to Justice Crisis [Electronic resource].-Access mode: https://www.probonoinst.org/2024/09/26/reimagining-legal-services-how-texas-is-tackling-the-access-to-

justice-crisis/ (date of access: 15.08.2025).
[8] 2024 CALIFORNIA JUSTICE GAP STUDY [Electronic resource].-Access mode: https://www.calbar.ca.gov/Portals/0/documents/reports/2025/2024-CA-Justice-Gap-Study.pdf (date of access: 15.08.2025).
[9] Ejjami R. AI-driven justice: Evaluating the impact of artificial intelligence on legal systems //Int. J. Multidiscip. Res. – 2024. – Vol. 6 (3). – pp. 1-29.
[10] Kumar D. AI-driven automation in administrative processes: Enhancing efficiency and accuracy //International Journal of Engineering Science and Humanities. – 2024. – Vol. 14 (1). – pp. 256-265.
[11] Shahid A., Qureshi G. M., Chaudhary F. Transforming legal practice: The role of AI in modern law //Journal of Strategic Policy and Global Affairs. – 2023. – Vol. 4 (1). – pp. 36-42.
[12] Nithya M. et al. AI-Driven Legal Automation to Enhance Legal Processes with Natural Language Processing //2024 International Conference on IoT Based Control Networks and Intelligent Systems (ICICNIS). – IEEE, 2024. – pp. 1246-1253.
[13] How Natural Language Processing (NLP) AI Is Used in Law-American Bar Association [Electronic resource].-Access mode: https://www.americanbar.org/groups/law_practice/resources/law-technology-today/2021/how-natural-language-processing-nlp-ai-is-used-in-law/ (date of access: 17.08.2025).
[14] Mentzingen H., António N., Bacao F. Automation of legal precedents retrieval: findings from a literature review //International Journal of Intelligent Systems. – 2023.-pp. 1-5.
[15] Bach N. X. et al. A two-phase framework for learning logical structures of paragraphs in legal articles //ACM Transactions on Asian Language Information Processing (TALIP). – 2013. – Vol. 12 (1). – pp. 1-32.
[16] de Almeida M. et al. Identifying legal party members from legal opinion documents using natural language processing //The 23rd International Conference on Information Integration and Web Intelligence. – 2021. – pp. 259-266.
[17] Smădu R. A. et al. Legal named entity recognition with multi-task domain adaptation //Proceedings of the Natural Legal Language Processing Workshop 2022.-pp. 305-321.
[18] Çetindağ C., Yazıcıoğlu B., Koç A. Named-entity recognition in Turkish legal texts //Natural Language Engineering. – 2023. – Vol. 29 (3). – pp. 615-642.
[19] Bushe G. R. Generative process, generative outcome: The transformational potential of appreciative inquiry //Organizational generativity: The appreciative inquiry summit and a scholarship of transformation. – Emerald Group Publishing Limited, 2013. – Vol. 4. – pp. 89-113.
[20] A inteligência artificial na advocacia: Transformações, desafios e o futuro do Direito [Electronic resource].-Access mode: https://www.migalhas.com.br/depeso/426895/a-ia-na-advocacia-transformacoes-desafios-e-o-futuro-do-direito (date of access: 17.08.2025).
[21] Ariai F., Mackenzie J., Demartini G. Natural language processing for the legal domain: A survey of tasks, datasets, models, and challenges //arXiv preprint arXiv:2410.21306. – 2024.-pp. 1-7.
[22] Zhang Z., Ma L., Nisbet N. Unpacking ambiguity in building requirements to support automated compliance checking //Journal of Management in Engineering. – 2023. – Vol. 39 (5). https://doi.org/10.1061/JMENEA.MEENG-5359.