

GENERATIVE AI IN SUPPORT OF SELECTED LEGAL AND ADMINISTRATIVE ACTIVITIES

Maria Mikołajewska¹, Monika Łaganowska², Dariusz Mikołajewski³

¹Wydział Prawa i Administracji, Uniwersytet Gdański, ul. Bażyńskiego 8, 80-309 Gdańsk

²Politechnika Bydgoska, ul. Kaliskiego 7, 85-796 Bydgoszcz

³Wydział Informatyki, Uniwersytet Kazimierza Wielkiego, ul. Mikołaja Kopernika 1, 85-074 Bydgoszcz
e-mail: dariusz.mikolajewski@ukw.edu.pl

Abstract: *Generative AI (Gen AI) transforms legal and administrative work by helping to rapidly draft contracts, pleadings, and routine correspondence, freeing professionals' time to focus on more demanding tasks, such as valuable analysis and strategy. Accelerates legal research through natural language queries and summaries, revealing precedents and regulations that match nuanced fact patterns in seconds. In administrative contexts, generative models automate form generation, policy templates, and multilingual communication, reducing administrative errors and turnaround times. When combined with ingest-enhanced generation and audit trails, these systems enable transparent sourcing, version control, and compliance monitoring, meeting evidentiary and procedural requirements. The result is a hybrid workflow where human expertise guides judgmental decisions while AI enables scalable, cost-effective document development, research, and management.*

Słowa kluczowe: *Computer science, artificial intelligence, generative AI, legal applications, administrative applications.*

GENERATYWNA SZTUCZNA INTELIGENCJA WE WSPARCIU WYBRANYCH DZIAŁAŃ PRAWNYCH I ADMINISTRACYJNYCH

Streszczenie: *Generatywna sztuczna inteligencja (GenAI) zmienia oblicze pracy prawnej i administracyjnej, pomagając szybko opracować umowy, pisma procesowe i rutynową korespondencję, uwalniając czas profesjonalistów, aby mogli skupić się na bardziej wymagających zadaniach, np. wartościowej analizie i strategii. Przyspiesza badania prawne poprzez zapytania i podsumowania w języku naturalnym, ujawniając precedensy i przepisy, które pasują do niuansów wzorców faktów w ciągu kilku sekund. W kontekstach administracyjnych modele generatywne automatyzują generowanie formularzy, szablony zasad i komunikację wielojęzyczną, redukując błędy administracyjne i czas realizacji. W połączeniu z generowaniem rozszerzonym o pobieranie i elementami audytu systemy te umożliwiają przejrzyste pozyskiwanie, kontrolę wersji i monitorowanie zgodności, spełniając wymogi dowodowe i proceduralne. Rezultatem jest hybrydowy przepływ pracy, w którym ludzka wiedza specjalistyczna kieruje decyzjami wymagającymi osądu, podczas gdy sztuczna inteligencja zapewnia skalowalne, ekonomiczne opracowywanie, badania i zarządzanie dokumentacją.*

Keywords: *Informatyka, sztuczna inteligencja, generatywna sztuczna inteligencja, zastosowania prawne, zastosowania w administracji.*

1. Introduction

The origins of generativeAI (genAI) for legal and administrative work date back to the 1980s, when rule-based expert systems like Mycin (a 1970s medical expert system) inspired early document authoring tools that could create simple contracts. With the release of word vector models like word2vec in 2013, lawyers began experimenting with semantic similarity searches to find

precedents more flexibly than keyword queries. The large-scale debut of transformer architectures in 2018, starting with Google's BERT (bidirectional encoder representations from transformers), showed that pre-trained language models can capture legal nuances, improving the classification of contract clauses and review of electronic disclosures. In 2020, OpenAI's GPT-3 demonstrated that a single pretrained model could produce, summarize, and translate legal text with minimal tuning, prompting legal technology vendors to embed "AI authoring assistants" into

their platforms. Around 2021, specialized legal domain corpora were used to tune GPT-style models (e.g., CaseLaw-GPT) that generate case law summaries and suggest jurisdiction-specific arguments. Regulatory agencies soon adopted pilots of generative AI: the U.S. Securities and Exchange Commission (SEC) used language models to flag invalid language in disclosures, while the U.K. Cabinet Office tested automated transcription of plain language briefing notes. By 2022, generative contract lifecycle management (CLM) systems could identify redlines in real time, negotiating a Non-Disclosure Agreement (NDA) or Data Protection Addendum (DPA), by comparing the model’s output to the organization’s handbook. Courts experimented with AI-generated plain-language explanations of judgments to improve public accessibility, an early form of “explainable justice.” In 2023–24, search-assisted generation (RAG) became the standard, combining vector search in secure knowledge bases with on-the-fly design to reduce hallucinations—critical to legal credibility. The state-of-the-art 2025 frontier is “automated policy,” where multi-agent generative systems create, simulate, and iterate administrative rules, shortening consultation cycles while keeping humans in the approval loop. Development constraints may be not only technical, but also legal (AI Act, General Data Protection Regulation - GDPR) [1,2].

The aim of this article is to analyze to what extent genAI can currently and in the future support selected legal and administrative activities.

1.1. Identified scientific gaps

Despite rapid adoption, research on generative AI for legal and administrative applications still has significant gaps in the science. Rigorous benchmarks specific to legal reasoning (beyond superficial textual similarity) are rare, making it difficult to quantify model competence across jurisdictions and procedural stages. Methods to ensure factual fidelity when models draft or summarize statutes lag behind practice, as search-augmented generation reduces hallucinations but does not offer formal proofs of correctness. Data scarcity persists for many administrative domains: confidential documents, minority language laws, and local government regulations are underrepresented in training corpora, undermining equity and coverage. There is limited understanding of how prompts or tuning choices propagate implicit biases to downstream legal outcomes, especially in the context of sentencing guidelines or benefits entitlement. Explainability techniques remain descriptive

(e.g., attention heatmaps) rather than normative, offering little practical information about why a model favored one legal argument over another [3,4]. Current privacy-preserving approaches (e.g., differential privacy, secure enclaves) have not been comprehensively tested against cross-document inference attacks common in litigation discovery. Human-AI interaction studies rarely examine long-term cognitive effects on lawyers (e.g., overreliance) beyond short-term task metrics. Interdisciplinary frameworks for assessing the constitutionality or administrative law of AI-generated rules are in their infancy, leaving policy-generation tools on uncertain legal footing. There is also a gap in socio-technical evaluation: there is a lack of long-term field studies that measure how GenAI changes access to justice, public trust, and systemic justice across populations [5,6].

1.2. Current state of the art and challenges

Cutting-edge generative AI platforms now produce contracts, policy summaries, and regulatory guidance in seconds by connecting multi-hundred-billion-dollar LLMs to secure, searchable knowledge bases. Advanced legal models (e.g., Harvey, CaseLaw-GPT, BloombergGPT) achieve near-human accuracy in classifying clauses and citing key passages when summarizing case law, enabling faster due diligence and opinion letters. Integrated contract lifecycle management packages redline clauses in real time against organizational manuals, automatically suggest counterlanguage, and generate audit-proof logs of negotiation rationales. Regulators are testing AI assistants that convert dense statutes into plain-language FAQs and multilingual versions, increasing accessibility while reducing staff workload. Experimental “AI-authored” English-language explanations of judgments are available, and some e-filing portals automatically fill out forms based on conversational messages. However, factual robustness remains a fundamental challenge: search alone cannot fully prevent hallucinations or subtle misinterpretations of statutory cross-references. Robust validation processes combining symbolic rule checking, citation tracking, and human review are still heavily manual, limiting end-to-end automation. Data management obstacles persist: privileged documents, local laws, and minority-language texts are rare or isolated, leading to gaps in coverage and potential bias [7,8]. Explainability tools reveal patterns of attention but rarely provide legally meaningful justifications, making it difficult for practitioners to trust AI-suggested arguments in high-stakes situations. Regulatory clarity lags behind

capabilities; Future AI accountability and transparency regulations will force vendors to embed audit logs, consent mechanisms, and origin tracking before wide-scale deployment is permitted [9,10].

2. Cases of use

In everyday life, also in a highly digital one, people encounter situations where they need to understand their rights, both in general and in relation to specific issues, even if they are not lawyers or do not have a lawyer in place. This is where the potential of genAI can be exploited, providing better accessibility to legal information, taking into account its limitations in different use cases in the legal field.

2.1. Transparent interaction

Pre-training of pre-trained models for specialist issues is faster, cheaper and increasingly widely used. The performance of genAI compared to legal decision support tools may even be better than that of classic systems designed specifically for legal advice, because genAI is better at understanding everyday language. Therefore, all users will be able to describe their situation, their needs, emerging legal problem more easily and obtain more precise answers. However, genAI may offer worse support in terms of accuracy, completeness and reliability of answers (AI estimates probability, not calculates, so the most probable answer does not have to be correct or literal) [11].

2.2. Representational fidelity

GenAI enables the editing of documents (e.g. payment demands, lawsuits, argumentative frameworks, or legal advice), but at a higher level of complexity it can summarize and summarise documents, identify legal strategies used in them, or suggest questions for cross-examination. However, genAI's results in law tests are on average at the level of a student taking the test, but with a low grade [12]. The way in which the questions are formulated is important here: genAI performs better in questions based on (e.g. essay), defining the general legal framework, etc., but performs worse in questions concerning specific legal problems, in which it must apply specific scenarios [13]. In such cases, training genAI using specific cases characteristic of a given group of legal

problems is of great importance. It is crucial to provide genAI with accurate information, for which the systems must be up-to-date, and the source of legal information - verified and authentic [11]. Otherwise, genAI may not comply with best practices, may provide false information, or even omit facts. This often allows for the detection of use by staff and students [14]. Poorly prepared training data can not only reduce the quality and accuracy of genAI, but also introduce bias, violate accountability, responsibility, confidentiality and privacy. This can lead to the erosion of skills in the long term and, with over-reliance, to a reduced sense of professional responsibility, as well as new educational challenges, changes in competitive advantages and the emergence of new forms of using legal services.

2.3. Informed action

Modified versions of solutions such as ChatGPT are already used by law firms to automate the drafting of legal documents and analyze contracts. Law students use them for educational purposes [15]. In addition, such tools facilitate interactions with clients (also: those who do not speak the local language well) and have significant potential in document review, updating the status of legal acts, collecting information and generating news summaries [16]. The lack of clarity regarding the use of genAI by specialists, transparency of rules in this area and full mutual understanding can lead to distrust between lawyers and clients, especially when lawyers and clients independently rely too much on genAI and check each other. Additionally, more advanced versions of genAI are more expensive, but more often trained, and therefore more accurate and up-to-date, but not everyone can afford them. This creates a completely new area of uneven access to information: with similar access to reasoning, the client may have significantly worse access to current data than the law firm, and therefore will achieve a worse effect. This will lead to new rates of using genAI, advanced prompt engineering, etc. On simpler, cheaper genAI systems it is better not to perform complex searches and legal analyses, and the results should always be checked by human experts for reliability of answers and their sources.

2.4. Learning and adaptation

There is a need to intensify training in the use of GenAI tools by lawyers [17]. In addition, it is important to build

awareness that the potential of genAI is enhanced by integration with other AI tools, especially those specifically designed to solve legal problems. This increases the accuracy and completeness of searches/answers while maintaining high quality of interaction. It is enough to mention here the Mixture of Experts (MoE) technique, which is increasingly used in genAI [11]. In some cases, the problem of incomplete answers can be solved by appropriate shaping of hints.

3. Discussion

The risk of data privacy breaches, unauthorized access, and misuse of personal data can be mitigated by anonymizing personal data (including edge computing), obtaining informed consent, using diverse data sets, and increasing transparency. The existing regulatory framework for genAI is already considered insufficient. New regulations should focus on data protection (including online consumer behavioral data), cybersecurity, and intellectual property. Oversight is needed to ensure stakeholder interests are maintained, rights and responsibilities are assigned, and a multi-layered approach to control is necessary, especially for high-risk applications. This is because society is already wary of AI's excessive impact on everyday life. Regulatory efforts should therefore be made to ensure transparency, risk management, and non-discrimination.

3.1. Limitations

GenAI models are expensive to train and produce, and only large companies, especially global corporations, can afford to fully exploit them. In order to balance the market opportunities, governments should invest in the development of such models and then make them available in the public domain for free or subsidized. Otherwise, there will be an imbalance of opportunity in the technology race. Up to now, genAI models have been largely based on European languages, using text/data available on the Internet in these languages, in accordance with their culture. On the one hand, emphasis must be placed on the development of national and regional systems (e.g. legal systems), and on the other hand, it is necessary to be aware of the trends that dominate, for example, in Chinese AI, also in the area of healthcare and law, where the need for specialization is important and exclusion (e.g. lack of access to up-to-date local legal information) has serious consequences. It seems reasonable to postulate that genAI tools should become public goods, available to everyone in a transparent manner with respect to the data and IT

methods/techniques used (not only AI algorithms, but also the strength of security against data leaks). A separate issue is the regulation of genAI tools, not only taking into account the specifics of genAI, but also defined for (foreseeable) new contexts and cases of their use. GenAI still generates hallucinations (i.e. fabricated facts or quotes) that can misinterpret legal arguments or official records from documents. Worse still, the training data for GenAI, as it is based on pre-trained systems, may contain outdated statutes or jurisdiction-specific nuances, so the results may even contradict current law or guidelines. The models rarely encode permission rules, risking inadvertently disclosing confidential or client-sensitive information when prompts or results are shared. They also struggle with precise statutory interpretations, such as disambiguating “must” versus “may,” which can have critical legal consequences. Because most systems are probabilistic (they estimate with a certain probability, not calculate), they cannot guarantee deterministic repetition of already approved clauses, which complicates version control in regulated documents. Prompt inputs can leak into future model responses (data retention or “prompt leakage”), raising concerns about GDPR and professional secrecy. Training data biases (including racial, gender, socioeconomic) can propagate to conviction records, enrollment lists, or benefits decisions, exposing institutions to discrimination claims. Current large language models (LLMs) do not have real-time access to sealed records or proprietary compliance databases, resulting in incomplete study coverage. Their explanations of reasoning steps are post hoc and can obscure logical errors, making it difficult for courts and agencies to require transparent reasoning. Licensing of model outputs is often ambiguous, so copyright or database rights violations may apply to generated forms or legal opinions. Adapting GenAI based on large pre-trained systems to a narrow set of regulations is expensive and fragile: changing regulations can invalidate months of training of a GenAI system on specific domains. For these reasons, relying on AI projects without legal oversight can violate professional conduct rules, but also fails to ensure the safety and due diligence of litigants. Fringe fact patterns (common in litigation) can fall outside the model's training distribution, leading to shallow or generalized guidance without analyzing the specifics of a particular case in detail. Latency and bandwidth constraints mean that large numbers of queries or the generation of mass notifications can be a bottleneck, especially when encryption or local hosting is required. Liability remains diffuse: when an AI recommendation causes a procedural error or damage, the allocation of liability among vendors,

lawyers, is legally uncertain. For these reasons, it seems reasonable to assume that decisions in this area should always be made by a human lawyer with the competence to assess the correctness of GenAI's advice.

3.2. Directions of further research

Future research should prioritize hybrid architectures that tightly couple LLMs with authoritative legal databases, ensuring that results remain anchored in current statutes, case law, and guidelines. Developing robust hallucination detection layers capable of flagging unsupported claims or missing citations in real time will be essential to the reliability of courts and regulations. eXplainable AI (XAI) techniques must evolve from post hoc summaries to traceable reasoning graphs, allowing practitioners to examine every textual leap from precedent to conclusion without a black box effect. Researchers should explore privacy-preserving tuning methods, such as federated learning, to ensure that confidential case files never leave institutional boundaries. Domain-adaptive reinforcement learning could teach models to follow jurisdiction-specific editorial conventions, such as citation rules or legislative formats, without constant manual correction. Benchmarks that measure the accuracy of statutory interpretation, not just surface-level semantic similarity, are needed to reflect true legal competence. Multimodal systems that combine text with scanned evidence, audio testimony, and video evidence can improve discovery but require new matching strategies to maintain fidelity to cross-modal inferences. Workflows involving human intervention merit empirical research testing how attorneys, clerks, or benefits officers best intervene to correct, approve, or override AI suggestions. Bias mitigation research should target "second-order" harms (such as disparate effects of official letters) by auditing both language and decision metrics. Power-efficient and device-based variants of LLM will be critical for government agencies that struggle with tight budgets or operate on secure, siloed networks. Generating synthetic legal data, carefully curated, could supplement rare training examples of rare applications or niche administrative rulings. Smart contracts and AI co-design could enable self-execution compliance checks, where generative models develop clauses that software agents later enforce. Cross-linguistic transfer techniques should improve, ensuring equal quality of service for minority and immigrant litigants navigating administrative systems. Adaptive risk assessment

frameworks could dynamically adjust the level of model autonomy—requiring more human oversight for high-stakes applications and less for routine forms. Interdisciplinary research combining computer science, jurisprudence, ethics, and public administration will be crucial to creating standards that balance innovation with due process protections.

The need for harmonized actions requires cooperation between policymakers, scientists and lawyers to streamline and control the dynamic landscape of genAI. On the one hand, this will ensure fuller use of the potential of genAI, but on the other hand, slowly preserve the stable intellectual property system. It will also take into account ethical concerns among users, creators and regulators, focusing on key issues of data origin, data privacy, usage rights, threats of unauthorized access or use, and broad new applications and malleability of genAI. Comprehensive education, frameworks and implementation strategies will be necessary for responsible integration of genAI into practice, protecting against potential societal risks.

4. Conclusions

GenAI has evolved from experimental text completion to a robust co-pilot that accelerates drafting, reviewing, and public communication in legal and administrative contexts. Its value lies in combining LLM creativity with domain-specific search and close human oversight, delivering speed without sacrificing compliance. Early concerns about "robot lawyers" have shifted to a pragmatic governance framework that treats AI as a support layer rather than a decision-maker. Regulators are now focused on transparency requirements (audit logs, citation tracking, bias testing) that make AI results questionable in court or in public inquiries. Meanwhile, cost savings and access to justice benefits are driving governments and companies to embed generative tools into everyday workflows. The next challenge is to harmonize international standards so that AI-enabled law and policy remain interoperable, accountable, and ultimately human-centric.

References

1. Dwivedi, D., De, R. Potential for GenAI in the Public Domain: A Review of Transportation, Healthcare, Agriculture, and Law. *Digit. Gov. Res. Pract.* 6(1): 3:1-3:11 (2025)

2. Dawood Al Lawati, S.F. I Am Not a Caveman: An Eye-tracking Study of How Users are Influenced to Search in the Era of GenAI. *WWW (Companion Volume)* 2025: 681-684.
3. Samuelson P. Generative AI meets copyright. *Science*. 2023 Jul 14;381(6654):158-161. doi: 10.1126/science.adi0656.
4. Pedersen I. The rise of generative AI and enculturating AI writing in postsecondary education. *Front Artif Intell*. 2023 Aug 10;6:1259407. doi: 10.3389/frai.2023.1259407.
5. Poddar A, Rao SR. Evolving intellectual property landscape for AI-driven innovations in the biomedical sector: opportunities in stable IP regime for shared success. *Front Artif Intell*. 2024 Sep 17;7:1372161. doi: 10.3389/frai.2024.137216.
6. Tang BL. The underappreciated wrong of AIgiarism - bypass plagiarism that risks propagation of erroneous and bias content. *EXCLI J*. 2023 Aug 26;22:907-910. doi: 10.17179/excli2023-6435.
7. Bockting CL, van Dis EAM, van Rooij R, Zuidema W, Bollen J. Living guidelines for generative AI - why scientists must oversee its use. *Nature*. 2023 Oct;622(7984):693-696. doi: 10.1038/d41586-023-03266-1.
8. Gao J, Wang D. Quantifying the use and potential benefits of artificial intelligence in scientific research. *Nat Hum Behav*. 2024 Dec;8(12):2281-2292. doi: 10.1038/s41562-024-02020-5.
9. Ong JCL, Chang SY, William W, Butte AJ, Shah NH, Chew LST, Liu N, Doshi-Velez F, Lu W, Savulescu J, Ting DSW. Ethical and regulatory challenges of large language models in medicine. *Lancet Digit Health*. 2024 Jun;6(6):e428-e432. doi: 10.1016/S2589-7500(24)00061-X.
10. Prokopowicz, P., Mikołajewski, D., Mikołajewska, E., Kotlarz, P. (2017). Fuzzy System as an Assessment Tool for Analysis of the Health-Related Quality of Life for the People After Stroke. In: Rutkowski, L., Korytkowski, M., Scherer, R., Tadeusiewicz, R., Zadeh, L., Zurada, J. (eds) *Artificial Intelligence and Soft Computing. ICAISC 2017. Lecture Notes in Computer Science()*, vol 10245. Springer, Cham. https://doi.org/10.1007/978-3-319-59063-9_64.
11. Jinzhe Tan, Hannes Westermann, and Karim Benyekhlef. 2023. ChatGPT as an artificial lawyer? Artificial intelligence for access to justice. In *Proceedings of the Workshop on Artificial Intelligence for Access to Justice (AI4AJ '23)*. 1–8.
12. Jonathan H. Choi, Kristin E. Hickman, Amy B. Monahan, and Daniel Schwarcz. 2021. ChatGPT goes to law school. *Journal of Legal Education* 71 (2021), 387.
13. Stuart Hargreaves. 2023. *Words Are Flowing Out Like Endless Rain into a Paper Cup: ChatGPT & Law School Assessments*. Research Paper No. 2203-03. The Chinese University of Hong Kong Faculty of Law.
14. Marjan Ajevski, Kim Barker, Andrew Gilbert, Liz Hardie, and Francine Ryan. 2023. ChatGPT and the future of legal education and practice. *Law Teacher* 57, 3 (2023), 352–364.
15. Sakshi Aggarwal. 2023. A review of ChatGPT and its impact in different domains. *International Journal of Applied Engineering Research* 18, 2 (2023), 119–123.
16. Julie Ayre, Olivia Mac, Kirsten McCaffery, Brad R. McKay, Mingyi Liu, Yi Shi, Atria Rezwan, and Adam G. Dunn. 2024. New frontiers in health literacy: Using ChatGPT to simplify health information for people in the community. *Journal of General Internal Medicine* 39, 4 (2024), 573–577.
17. Som S. Biswas. 2023. Role of Chat GPT in public health. *Annals of Biomedical Engineering* 51, 5 (2023), 868–869.
18. Szu-Wei Cheng, Chung-Wen Chang, Wan-Jung Chang, Hao-Wei Wang, Chih-Sung Liang, Taishiro Kishimoto, Jane Pei-Chen Chang, John S. Kuo, and Kuan-Pin Su. 2023. The now and future of ChatGPT and GPT in psychiatry. *Psychiatry and Clinical Neurosciences* 77, 11 (2023), 592–596.
19. Weiqiang Jin, Biao Zhao, and Guizhong Liu. 2024. Exploring the capability of ChatGPT for cross-linguistic agricultural document classification: Investigation and evaluation. In *Neural Information Processing*, Biao Luo, Long Cheng, Zheng-Guang Wu, Hongyi Li, and Chaojie Li (Eds.). Communications in Computer and Information Science. Springer Nature, Singapore, 220–237.