UDC 340.1:004.89
DOI https://doi.org/10.32782/2521-6643-2025-2-70.22

Shamov O. A., Intelligent Systems Researcher,
Head of Human Rights Educational Guild
ORCID: 0009-0009-5001-0526

THE ROLE OF GENERATIVE AI IN THE AUTOMATION OF LEGAL DOCUMENT PREPARATION: AN ANALYSIS OF ACCURACY, EFFICIENCY, AND ETHICAL ASPECTS

This article explores the transformative impact of generative artificial intelligence (AI) on legal practice, with a particular focus on the preparation and analysis of legal documents. It examines modern generative AI models, their capabilities, limitations, and the risks associated with their implementation in the legal domain. The aim of the article is to analyze the current state of generative AI applications for automating legal tasks, identify key issues related to the accuracy and reliability of generated content, and propose a hypothetical model to enhance the trustworthiness of AI-generated outcomes.

To achieve this, the article sets out the following objectives: to review recent academic research and publications evaluating the effectiveness and precision of large language models (LLMs) in legal tasks; to identify the unresolved issue of factual verification of generated content; and to formulate and justify a hypothesis for improving baseline research results through the implementation of a hybrid architecture.

The study employs methods of systems analysis, comparative analysis of empirical studies, synthesis, and modeling. It analyzes foundational research on structured approaches to AI-assisted legal document analysis and the outcomes of comparative tests between LLMs and human legal professionals, highlighting the persistent challenge of factual accuracy in AI-generated content

The article proposes a hypothesis that integrating structured prompting techniques with a Retrieval-Augmented Generation (RAG) architecture-leveraging a curated and dynamically updated legal knowledge base- can significantly improve the factual precision and reliability of AI-generated legal documents.

This hypothesis is substantiated with a formalized expression and a potential mechanism for implementation and verification.

The proposed approach is shown to not only minimize the risk of AI "hallucinations" but also to lay the groundwork for the development of more responsible and ethical tools that augment, rather than replace, the professional competencies of legal practitioners.

Key words: *generative artificial intelligence, large language models, legal document automation, Retrieval-Augmented Generation (RAG), AI accuracy, legal technology, AI ethics, data verification.*

Шамов О. А. Роль генеративного ШІ в автоматизації підготовки юридичних документів: аналіз точності, ефективності та етичних аспектів

Стаття присвячена дослідженню трансформаційного впливу генеративного штучного інтелекту (ШІ) на юридичну практику, зокрема на процес підготовки та аналізу юридичних документів. Предметом вивчення є сучасні моделі генеративного ШІ, їхні можливості, обмеження та ризики, пов'язані з їхнім впровадженням у юридичній сфері. Метою статті є аналіз поточного стану застосування генеративного ШІ для автоматизації юридичних завдань, виявлення ключових проблем, пов'язаних з точністю та надійністю згенерованих даних, та розробка гіпотетичної моделі для підвищення достовірності результатів. Для досягнення цієї мети були поставлені такі завдання: проаналізувати останні наукові дослідження та публікації, що оцінюють ефективність та точність великих мовних моделей (LLM) у юридичних завданнях; виокремити невирішену проблему забезпечення фактичної верифікації згенерованого контенту; сформулювати та обґрунтувати гіпотезу щодо покращення результатів базових досліджень шляхом впровадження гібридної архітектури.

Дослідження базується на методах системного аналізу, порівняльного аналізу емпіричних досліджень, синтезу та моделювання.

Проаналізовано базове дослідження щодо структурованого підходу до аналізу юридичних документів за допомогою ШІ, а також результати порівняльних тестів LLM та юристів-людей, виокремлено невирішену проблему забезпечення фактичної верифікації згенерованого контенту.

Запропоновано гіпотезу, що інтеграція методів структурованого промптингу з архітектурою Retrieval-Augmented Generation (RAG), яка спирається на кураторську та динамічно оновлювану базу юридичних знань, може суттєво підвищити фактичну точність та надійність генерованих ШІ юридичних документів.

Обґрунтовано цю гіпотезу, її формалізоване вираження та потенційний механізм імплементації та верифікації.

Розкрито, що запропонований підхід не лише мінімізує ризик «галюцинацій» Ші, але й створює основу для розробки більш відповідальних та етичних інструментів, що посилюють, а не замінюють, професійні компетенції юристів.

Ключові слова: *генеративний штучний інтелект, великі мовні моделі, автоматизація юридичних документів, Retrieval-Augmented Generation (RAG), точність ШІ, юридичні технології, етика штучного інтелекту, верифікація даних.*

Statement of the problem. The penetration of generative artificial intelligence (AI) into various professional fields has become one of the most significant technological trends of recent years. The legal industry, traditionally considered conservative, has not remained aloof from these processes. Large language models (LLMs), such as GPT-4 and its counterparts, demonstrate an impressive ability to generate, analyze, and summarize textual information, opening up unprecedented opportunities for automating routine tasks that previously required significant time and human resources [1]. Drafting contracts, memoranda, statements of claim, analyzing large volumes of evidence, and conducting legal research are just a few of the processes that can potentially be transformed with the help of AI [2].

The connection of this issue with important practical tasks is obvious [3]. According to studies, law firms and corporate legal departments are constantly looking for ways to increase efficiency and reduce costs.

Empirical research shows that LLMs can perform certain legal tasks much faster and more cheaply than junior lawyers [4]. This creates the potential to expand access to justice and provide more affordable legal services [5].

However, alongside the prospects of efficiency, a fundamental scientific and practical problem arises ensuring the accuracy, reliability, and ethical use of these tools. Unlike in many other fields, an error in a legal document can have serious legal and financial consequences for clients. The main risk inherent in modern LLMs is a phenomenon known as «hallucinations», which is the generation of plausible but factually false or inaccurate information [6]. A model might cite non-existent court precedents, misinterpret legal statutes, or create legally incorrect formulations. Thus, a critical task emerges: to develop methods and architectures for AI systems that not only automate but also ensure a high level of credibility and verification of the generated content, consistent with the professional standards of jurisprudence.

Analysis of recent research and publications. The problem of AI application in law is being actively discussed in the scientific community. A significant number of recent publications are devoted to both the potential benefits and the risks.

One of the key areas is the empirical evaluation of LLM performance compared to lawyers. A 2024 study conducted by scholars at Cornell University, titled «Better Call GPT, Comparing Large Language Models Against Lawyers», demonstrated that the GPT-4 model surpassed lawyers in the speed and cost-effectiveness of performing contract analysis and summarization tasks, although it was inferior in accuracy when identifying specific legal risks [4]. Another large-scale study by lawyers from the University of Minnesota and the University of Michigan («AI-Powered Lawyering») showed that lawyers using GPT-4 completed tasks 24% faster and produced 40% higher quality documents. However, the study also found that the model is prone to generating false information [6].

The foundational study for this article was conducted by Wang, Chen, & Li in 2024 and published under the title «Enhancing Legal Document Analysis with Large Language Models: A Structured Approach to Accuracy, Context Preservation, and Risk Mitigation» [7]. The authors proposed a methodology that combines hierarchical document segmentation, Chain-of-Thought prompting, and iterative summarization to improve the accuracy of analyzing complex legal texts.

Their approach showed a significant improvement in preserving context and reducing the risks of misinterpretation compared to standard queries to LLMs [7].

Despite the significant contribution of this research to structuring interaction with AI, it does not solve the main problem-the lack of an external verification mechanism. Structured prompting only optimizes the use of the model's internal «knowledge» but does not guarantee that this knowledge is accurate and up-to-date. This problem is highlighted by numerous works on ethical aspects. Perlman, in his 2025 work, emphasizes that the professional duty of competence requires lawyers not to blindly trust but to thoroughly verify the results of AI's work [8]. The American Bar Association (ABA) also warns of the risks of bias embedded in the data on which the models were trained, which can lead to discriminatory outcomes [9].

To overcome the problem of «hallucinations», researchers are increasingly turning to the Retrieval-Augmented Generation (RAG) architecture. This approach, as noted in the Harvard Journal of Law & Technology, allows for «grounding» LLM responses in data from reliable external sources [10]. Instead of relying solely on its parametric memory, the model first searches for relevant information in a specialized knowledge base (e.g., a database of legislation or case law) and then uses the retrieved data to generate the answer. A 2023 study from Stanford University showed that even RAG-based tools are not completely protected from errors, but they significantly reduce their frequency [11].

Thus, the literature analysis reveals an unresolved part of the general problem: while methods exist to improve the structure of queries to AI (the research of Wang et al.) and methods for external grounding of responses (RAG), there is a lack of a comprehensive approach that would combine both strategies to achieve a synergistic effect of increased reliability. This article is dedicated to addressing this problem.

Task statement. Based on the identified problem, the main goals (tasks) of this article are:

1. To analyze and summarize the methodology and results of the foundational study by Wang, Chen, & Li (2024) on the structured approach to legal document analysis.

2. To identify the limitations of this approach, particularly its vulnerability to generating factually unverified information.

3. To propose and substantiate a hypothesis for improving the accuracy and reliability of generative models by creating a hybrid system that combines structured prompting with the Retrieval-Augmented Generation (RAG) architecture.

4. To develop a conceptual model for the implementation and verification of the proposed hybrid system in the context of automating legal document preparation.

The purpose of this article is to develop and theoretically substantiate a new, hybrid model for the use of generative AI in legal practice, which aims to maximize accuracy and reliability through the synergistic combination of advanced techniques for model interaction and architectural solutions that ensure real-time data verification.

Outline of the main material of the study. The basis for our research is the work of Wang and his colleagues [7], which proposes a systematic approach to using LLMs for analyzing legal contracts. Their methodology consists of three key stages:

1. Hierarchical Segmentation: Instead of analyzing the entire document as a single unit, the authors propose breaking it down into logical parts (sections, articles, clauses). This allows the model to better focus on the local context and avoid losing important details, which often happens when processing long texts.

2. Chain-of-Thought Prompting: For each segmented part, the model is instructed not just to provide a final answer but to demonstrate a step-by-step reasoning process. For example, when analyzing a clause on liability, the model should first identify the parties, then describe the conditions for liability, determine its type and extent, and only then formulate a final conclusion. This makes the process more transparent and allows for easier detection of logical errors.

3. Iterative Summarization: After analyzing individual segments, the model creates short summaries of them. These summaries are then combined and summarized at a higher level, creating a comprehensive yet detailed overview of the entire document.

This structured approach is undoubtedly a step forward. It disciplines the «thinking» of the LLM and reduces the likelihood of gross errors related to context loss. However, its fundamental limitation is that it operates within the «closed world» of the model's own knowledge. If the LLM was trained on outdated data or contains incorrect representations of certain legal norms, structured prompting will not correct this. The model can logically and consistently justify its answer based on a non-existent legal norm [6]. This is where the need to go beyond the model's internal knowledge arises.

The solution to this problem is the Retrieval-Augmented Generation (RAG) architecture. Unlike a standard LLM, a RAG system performs an additional step before generating a response: searching for information in an external database. In a legal context, such a database (or vector database) can contain the full texts of laws, codes, regulations, current case law, and scholarly commentaries [10]. The process of such a system is as follows:

1. Receiving a Query: A lawyer formulates a query, for example, «Draft a clause for a lease agreement on rent indexation in accordance with the current legislation of Ukraine.»

2. Retrieval: The RAG system does not pass the query directly to the LLM. First, it analyzes the query and searches its knowledge base for the most relevant documents: articles of the Civil Code of Ukraine governing rent and indexation, recent Supreme Court decisions on this issue, etc.

3. Augmentation: The found text fragments (e.g., specific articles of law) are added to the user's initial query as additional context.

4. Generation: The modified, augmented query is passed to the LLM with instructions to generate a response based primarily on the provided context.

Thus, RAG forces the model to rely not on its abstract memory but on specific, verified, and current legal sources. This significantly reduces the risk of «hallucinations» and increases the credibility of the result [12].

However, as the Stanford University study shows, even RAG systems can sometimes provide inaccurate information if the retriever selects an irrelevant fragment from the knowledge base [11].

Hypothesis for improvement: Based on the analysis of the strengths and weaknesses of both approaches, we put forward the following hypothesis:

A hybrid model that synergistically combines structured prompting (following the methodology of Wang et al.) with a multi-layered Retrieval-Augmented Generation (RAG) architecture will provide a significantly higher level of factual accuracy, completeness, and reliability of generated legal documents compared to using each of these methods alone.

Rationale for the hypothesis: The rationale for the hypothesis lies in the mutual complementarity and reinforcement of both approaches. Structured Prompting (SP) organizes the model's thought process, forcing it to work methodically and consistently. This creates a «scaffold» for the answer. RAG, in turn, fills this scaffold with factual, verified content.

The synergistic effect arises at several levels:

- SP improves RAG: Applying hierarchical segmentation and chain-of-thought to the query before the retrieval stage in RAG can make the search itself more accurate. Instead of one general query to the knowledge base, the system will perform a series of smaller, more focused queries for each logical step, reducing the risk of retrieving irrelevant documents.
- RAG improves SP: Even the best chain of thought will be flawed if it is based on incorrect initial data. RAG ensures that each step in the chain of thought is based on a current legal norm or judicial precedent, making the entire logical construction credible.

For a scientific rationale, we can introduce a conditional formula for assessing accuracy (A). Let A_{base} be the baseline accuracy of a standard LLM. The use of structured prompting adds a certain improvement factor I_{SP}, and the use of RAG adds a factor I_{RAG}. Our hypothesis is that the effect is not merely additive but contains a synergistic component. This can be expressed by the formula:

$$A_{hybrid} = A_{base} + I_{SP} + I_{RAG} + S(I_{SP}, I_{RAG}) \quad (1)$$

where A_{hybrid} is the accuracy of the hybrid model, and $S(I_{SP}, I_{RAG})$ is a positive synergistic term that depends on the interaction of both methods. This can be more clearly represented as a multiplicative interaction, where RAG enhances the quality of the result already improved by SP:

$$A_{hybrid} = (A_{base} + I_{SP}) \times (1 + k \cdot I_{RAG}) \quad (2)$$

where k is a coefficient reflecting the effectiveness of the integration.

This formula illustrates that RAG does not just add accuracy but multiplies it by acting on an already structured and logically organized process.

Potential implementation and verification. Implementation of the hybrid model requires the creation of two main components:

1. Multi-layered Knowledge Base: This is not just a repository of texts. The base must be structured by levels of reliability and priority:

- Level 1 (Absolute truth): Codified legislation (Constitution, codes, laws).
- Level 2 (Official interpretation): Decisions of the Constitutional and Supreme Courts, regulations of regulatory bodies.
- Level 3 (Doctrine and practice): Scholarly articles, commentaries on legislation, generalized judicial practice of lower courts. The RAG system should be configured for priority search in the higher levels.

2. Hybrid Query Processor: This module will perform a sequence of operations:

- Receives a query from the user.
- Applies hierarchical segmentation and chain-of-thought techniques to break it down into sub-queries.
- For each sub-query, performs a search in the multi-layered knowledge base (RAG).
- Forms augmented prompts for the LLM.
- Generates responses and automatically adds citations to the sources used during the retrieval.

Verification of the hypothesis should be conducted through a controlled experiment. It is necessary to create a set of typical legal tasks (e.g., drafting a contract, risk analysis, writing a statement of claim). Then, three groups should be engaged to perform them:

- Group 1 (Control): Uses a standard LLM (e.g., GPT-4).
- Group 2 (SP): Uses an LLM with the application of structured prompting techniques.
- Group 3 (Hybrid): Uses the proposed hybrid model.

The results of all three groups should be independently evaluated by a committee of expert lawyers according to the following criteria:

- Factual accuracy: Absence of «hallucinations», correct citation of legal norms.
- Completeness: Coverage of all essential aspects of the problem.
- Legal correctness: Correctness of formulations and document structure.
- Presence and correctness of citations: Automatic referencing of sources.

A comparison of the average scores on these criteria between the groups will allow for the confirmation or refutation of the proposed hypothesis.

Conclusions. The conducted research confirms that generative AI has the potential to radically change legal practice by significantly increasing its efficiency. However, the key obstacle to its widespread adoption remains the problem of reliability and accuracy, embodied in the phenomenon of «hallucinations».

The analysis of advanced research has shown that existing approaches, such as structured prompting and the Retrieval-Augmented Generation architecture, while effective, have their limitations and do not solve the problem comprehensively. This article has presented and substantiated a hypothesis that the synergistic integration of these two methods into a single hybrid model can provide a breakthrough in creating reliable legal AI tools.

The proposed model, which combines the logical discipline of structured prompting with the factual grounding of RAG on a curated knowledge base, is capable of not only generating legally sound texts but also confirming their authenticity with citations to primary sources.

The practical significance of the proposed approach lies in creating a roadmap for legal tech developers that will allow them to move from creating «assistants» with unpredictable accuracy to developing reliable tools that enhance the professional capabilities of a lawyer. This, in turn, will promote adherence to ethical standards, particularly the duty of competence and confidentiality.

Prospects for further research in this area are multifaceted:

1. Empirical implementation and testing: The primary task is the practical development and experimental verification of the proposed hybrid model to quantitatively measure its benefits.

2. Optimization of knowledge bases: Researching methods for automatic updating and verification of legal knowledge bases to ensure the RAG system always works with the most current information.

3. Explainable AI (XAI): Integrating XAI mechanisms that would not just show the source of information but also visualize the logical chain that led to a particular conclusion, making the AI's work completely transparent to the lawyer.

4. Adaptation to different fields of law: Studying the specifics of applying the hybrid model in different legal fields (criminal, tax, corporate law), each of which has its unique sources and accuracy requirements.

Ultimately, the future of AI in law is not a race for full autonomy, but the creation of powerful, reliable, and ethical tools that, under the control of a qualified professional, will elevate the quality and accessibility of legal services to a new level.

Bibliography:

1. McCallon, J. Legal Tech ROI: Measuring AI Implementation Success. *CallidusAI Blog*. 2025. URL: https://callidusai.com/legal-tech-roi-measuring-ai-implementation-success/ (дата звернення: 22.06.2025)

2. Terzidou K. Generative AI systems in legal practice: Offering quality legal services while upholding legal ethics. *International Journal of Law in Context*. 2025.

URL: https://www.cambridge.org/core/journals/international-journal-of-law-in-context/article/generative-ai-systems-in-legal-practice-offering-quality-legal-services-while-upholding-legal-ethics/34011A84AA58A2BAB556A406A4653A8D/ (дата звернення: 22.06.2025)

3. Braff D. AI is the future of law, but most legal pros aren't trained for it, a new report says. *ABA Journal*. 2025. URL: https://www.abajournal.com/web/article/ai-is-the-future-of-law-but-most-legal-pros-arent-trained-for-it-a-new-report-says/ (дата звернення: 22.06.2025)

4. Martin L., Whitehouse N., Yiu S., Catterson L. Perera R. Better Call GPT, Comparing Large Language Models Against Lawyers. *arXiv preprint arXiv:2401.16212*. 2024. URL: https://arxiv.org/abs/2401.16212/ (дата звернення: 22.06.2025)

5. Baxter J. AI-driven legal tech trends for 2025. *NetDocuments Blog*. 2025. URL: https://www.netdocuments.com/blog/ai-driven-legal-tech-trends-for-2025/ (дата звернення: 22.06.2025)

6. Schwarcz D., Manning S., Barry P., Cleveland D., Prescott J., Rich B., Deakins O. AI-Powered Lawyering: AI Reasoning Models, Retrieval Augmented Generation, and the Future of Legal Practice. *Learning Ecosystems*. 2025. URL: http://danielschristian.com/learning-ecosystems/2025/05/22/ai-powered-lawyeringby-various-authors-other-legaltech-related-items/ (дата звернення: 22.06.2025)

7. Davenport M. Enhancing legal document analysis with Large Language Models: A structured approach to accuracy, context preservation, and risk mitigation. *Scientific Research an Academic Publisher*. 2025. URL: https://www.scirp.org/journal/paperinformation?paperid=141892/ (дата звернення: 22.06.2025)

8. Perlman A. The Legal Ethics of Generative AI. *Suffolk University Law Review*. 2024. URL: https://bpb-us-e1.wpmucdn.com/sites.suffolk.edu/dist/3/1172/files/ 2025/02/1FINAL_Perlman_Article_WDFF.pdf/ (дата звернення: 22.06.2025)

9. Donald D. Bias in AI Large Language Models: Risks and Remedies. *GPSolo Magazine*. 2025. URL: https://www.americanbar.org/groups/gpsolo/resources/magazine/2025-mar-apr/bias-ai-large-language-models-risks-remedies/ (дата звернення: 22.06.2025)

10. Johnston P. Retrieval-Augmented Generation (RAG): Towards a promising LLM architecture for legal work? *Harvard Journal of Law & Technology: JOLT Digest.* 2025. URL: https://jolt.law.harvard.edu/digest/retrieval-augmented-generation-rag-towards-a-promising-llm-architecture-for-legal-work (дата звернення: 22.06.2025)

11. Magesh V., Surani F., Dahl M., Suzgun M., Manning C., Ho D. Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools. *Journal of Empirical Legal Studies.* 2025. URL: https://dho.stanford.edu/wp-content/uploads/Legal_RAG_Hallucinations.pdf/ (дата звернення: 22.06.2025)

12. Kankatala, V. Retrieval-Augmented Generation (RAG) in the legal industry. *Acroplans Blog.* 2024. URL: https://www.acroplans.com/retrieval-augmented-generation-rag-in-the-legal-industry/ (дата звернення: 22.06.2025)

References:

1. McCallon, J. (2025). Legal Tech ROI: Measuring AI Implementation Success. *CallidusAI Blog*. Retrieved from: https://callidusai.com/legal-tech-roi-measuring-ai-implementation-success/

2. Terzidou, K. (2025). Generative AI systems in legal practice: Offering quality legal services while upholding legal ethics. *International Journal of Law in Context,* First View, 1-22. Retrieved from: https://www.cambridge.org/core/journals/international-journal-of-law-in-context/article/generative-ai-systems-in-legal-practice-offering-quality-legal-services-while-upholding-legal-ethics/34011A84AA58A2BAB556A406A4653A8D

3. Braff, D. (2025). AI is the future of law, but most legal pros aren't trained for it, a new report says. *ABA Journal*. Retrieved from: https://www.abajournal.com/web/article/ai-is-the-future-of-law-but-most-legal-pros-arent-trained-for-it-a-new-report-says

4. Martin, L., Whitehouse, N., Yiu, S., Catterson, L. & Perera, R. (2024). Better Call GPT, Comparing Large Language Models Against Lawyers. *arXiv,* 2401.16212. Retrieved from: https://arxiv.org/abs/2401.16212

5. Baxter, J. (2025). AI-driven legal tech trends for 2025. *NetDocuments Blog*. Retrieved from: https://www.netdocuments.com/blog/ai-driven-legal-tech-trends-for-2025/

6. Schwarcz, D., Manning, S., Barry, P., Cleveland, D., Prescott, J., Rich, B. & Deakins, O. (2025). AI-Powered Lawyering: AI Reasoning Models, Retrieval Augmented Generation, and the Future of Legal Practice. *Learning Ecosystems*. Retrieved from: http://danielschristian.com/learning-ecosystems/2025/05/22/ai-powered-lawyeringby-various-authors-other-legaltech-related-items/

7. Davenport, M. (2025). Enhancing legal document analysis with Large Language Models: A structured approach to accuracy, context preservation, and risk mitigation. *Scientific Research an Academic Publisher,* Vol.15 No.2. Retrieved from: https://www.scirp.org/journal/paperinformation?paperid=141892

8. Perlman, A. (2024). The Legal Ethics of Generative AI. *Suffolk University Law Review,* Vol.LVII:345 2024. Retrieved from: https://bpb-us-e1.wpmucdn.com/sites.suffolk.edu/dist/3/1172/files/2025/02/1FINAL_Perlman_Article_WDFF.pdf

9. Donald, D. C. (2025). Bias in AI Large Language Models: Risks and Remedies. *GPSolo Magazine,* 42:2. Retrieved from: https://www.americanbar.org/groups/gpsolo/resources/magazine/2025-mar-apr/bias-ai-large-language-models-risks-remedies/

10. Johnston, P. (2025). Retrieval-Augmented Generation (RAG): Towards a promising LLM architecture for legal work. *Harvard Journal of Law & Technology: JOLT Digest.* Retrieved from: https://jolt.law.harvard.edu/digest/retrieval-augmented-generation-rag-towards-a-promising-llm-architecture-for-legal-work

11. Magesh, V., Surani, F., Dahl, M., Suzgun, M., Manning, C., & Ho, D. (2025). Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools. *Journal of Empirical Legal Studies,* 1-27. Retrieved from: https://dho.stanford.edu/wp-content/uploads/Legal_RAG_Hallucinations.pdf

12. Kankatala, V. (2024). Retrieval-Augmented Generation (RAG) in the legal industry. *Acroplans Blog.* Retrieved from: https://www.acroplans.com/retrieval-augmented-generation-rag-in-the-legal-industry

Дата надходження статті: 03.10.2025
Дата прийняття статті: 10.11.2025
Опубліковано: 30.12.2025