

The Evolution of Legal AI Applications: A Scientometric Analysis

Ying Chen¹, Yundong Wu^{2*}, Weijian Kong^{2,3}

¹ Quanzhou Ocean Institute, Quanzhou, 362700, China

^{2*} Independent Researcher, Naju, 58245, South Korea

³ College of Humanities and International Studies, Keimyung University, 42601, Daegu, South Korea

Korespondensi: ethanwu0525@gmail.com

Article Info

Article history:

Received July 25th, 2026

Revised Aug 23th, 2026

Accepted Sept 2th, 2026

Keyword:

Artificial intelligence; Legal AI;
Legal technology; Natural language
processing; Scientometric analysis.

ABSTRACT

This study employs a scientometric approach to examine 421 Legal AI publications indexed in the Scopus database between 2015 and 2025, mapping the structural evolution, key thematic clusters, and emerging research frontiers of the field. The empirical findings reveal a gradual increase in publication volume after 2019, followed by an unprecedented expansion after 2022. While core tasks such as legal decision support and legal information retrieval constitute the largest topic categories, recent scholarship demonstrates a clear shift toward generative AI, legal chatbots, and critical questions regarding ethics, privacy, and algorithmic accountability. Longitudinal keyword analysis confirms a transition from general machine learning models toward natural language processing, large language models, explainability, and responsible AI governance. Furthermore, geographic analysis shows that despite increasing global participation, output remains concentrated in a limited group of countries led by the United States, China, and the United Kingdom. The paper concludes that future Legal AI research must move beyond isolated task accuracy to evaluate system validity, source reliability, human oversight, and institutional adaptation within complex legal frameworks.



© 2026 The Authors. Published by Punggawa Legacy Center. This is an open access article under the CC BY license (<https://creativecommons.org/licenses/by/4.0/>)

INTRODUCTION

Artificial intelligence has moved from a mainly technical research area to a technology used in many professional fields. (Surden, 2014) Law is one of the fields where this change is especially important because legal work depends on large amounts of text, structured rules, previous decisions, evidence, and professional judgment. AI systems can help users search legal materials, classify documents, identify relevant cases, summarize information, support contract review, and assist with legal prediction. These functions have increased interest in Legal AI among computer scientists, lawyers, legal scholars, and policy researchers (Krausová, 2017; Lehr & Ohm, 2017; Zhou et al., 2020).

The development of Legal AI has a longer history than the recent growth of generative AI. (Branting, 2017) Earlier systems were mainly based on expert rules, symbolic reasoning, and knowledge representation. Later research increasingly used machine learning and natural language processing to process legal documents (Domingos, 2015; Vaswani et al., 2017). More recent systems use large language models and other foundation models to generate and transform legal text. This development has changed the research questions. Earlier studies often asked whether a computer could perform a specific legal task. Current studies also ask whether AI output can be trusted, whether its reasoning can be explained, who is responsible for an error, and how AI should be controlled in legal institutions.

The legal context makes these questions different from those in many other application areas. (Deeks, 2019) A wrong recommendation in a commercial search system may be inconvenient, but a

wrong legal classification or incorrect legal statement can affect rights, access to justice, professional decisions, and institutional legitimacy. AI systems can also reproduce bias in historical data and may provide confident outputs that are not supported by reliable legal sources. These risks have increased attention to transparency, privacy, accountability, human oversight, and professional responsibility (Floridi & Cowls, 2019; Pasquale, 2015; Yu & Spina Ali, 2019).

The growth of AI in professional practice is also connected with wider institutional change. AI can support faster information processing and reduce some routine work, but it does not automatically replace professional judgment. (Wu & Kong, 2026) The human role remains important because legal decisions involve interpretation, values, context, and responsibility. A human-centered view of AI therefore requires attention to both technical performance and institutional practice. Similar concerns have been raised in education, where AI can improve efficiency but also creates concerns about privacy, quality, human interaction, and overreliance (Wu, 2025). In organizational decision-making, AI and big data also require governance arrangements that keep human judgment and institutional responsibility in the process (Wu & Kong, 2026).

Existing Legal AI literature is large and diverse. Some studies focus on legal information retrieval, while others examine prediction, document classification, contract analysis, judicial decision support, legal chatbots, or ethical issues. Because these studies are distributed across law, computer science, information science, and social science journals, it can be difficult to see how the field has developed as a whole (McGinnis & Pierce, 2019; Susskind, 2020). A scientometric analysis can provide a structured overview by examining publication growth, topic distribution, author or country participation, and keyword development. (Zupic & Čater, 2015)

This study follows a research-review structure that first defines the analytical scope, then examines annual development, topic distribution, detailed research areas, and author patterns before discussing the current state and future directions.

The study addresses four questions:

RQ1. How has the volume of Legal AI research changed from 2015 to 2025?

RQ2. What major research topics and detailed application areas have been studied?

RQ3. How is Legal AI research distributed across major country groups?

RQ4. How have important keywords changed over time, and what do these changes suggest about future research directions?

The purpose is not only to count publications. It is also to identify how the research focus has changed and what areas appear to be gaining attention. The results are intended to provide a simple map for researchers who are entering Legal AI and for legal institutions that need to understand the direction of current research.

RESEARCH METHODS

1. Research design

This study uses a scientometric literature analysis. Scientometric analysis is useful when a research field contains a large and interdisciplinary body of publications because it allows researchers to identify patterns in publication volume, topics, authors, keywords, and other bibliographic features. Bibliometric methods are widely used to map scientific fields and can be combined with qualitative interpretation to explain the meaning of observed patterns (Aria & Cuccurullo, 2017). (Mongeon & Paul-Hus, 2016) The value of such analysis is not limited to counting documents. It can also show the development of a field and help identify areas that are receiving increasing attention (Zupic & Čater, 2015).

The research design was organized into five stages: definition of the search scope, retrieval and screening of publications, classification of the selected records, descriptive analysis, and interpretation of emerging themes.

2. Data source and search scope

The literature corpus was retrieved from the Elsevier Scopus database, covering publications from 2015 to 2025. Scopus was selected because of its comprehensive coverage of peer-reviewed literature across both computer science and legal/social science domains. The

search query was executed within the title, abstract, and keyword fields (TITLE-ABS-KEY) using the following search expression:

TITLE-ABS-KEY (("artificial intelligence" OR "machine learning" OR "natural language processing" OR "large language model" OR "generative AI") AND ("law" OR "legal" OR "judicial" OR "legal technology" OR "legal services")) AND PUBYEAR > 2014 AND PUBYEAR < 2026*

The search was limited to original research articles, conference papers, and reviews published in English.

3. Screening and classification

The screening procedure followed three steps. First, titles and abstracts were examined to remove publications that only mentioned AI or law without studying their relationship. Second, records focused on general AI engineering without a legal application were excluded. Third, records focused on legal policy without a meaningful AI component were excluded. This approach ensures a rigorous selection process, where a broad database search is followed by title and abstract review and duplicate removal before the final set of studies is analyzed.

Each publication was assigned one main topic according to its dominant research purpose. The topic categories were legal decision support, legal research and retrieval, legal prediction and analytics, contract analysis, legal chatbots and generative AI, judicial and administrative applications, and ethics, privacy, and accountability. Because many publications address more than one issue, the main topic was selected according to the primary purpose of the study. This avoids double counting when calculating the main topic distribution.

4. Analytical procedures

Four main analyses were conducted. First, annual publication counts were calculated to identify changes in research volume. Second, publications were grouped into five-year periods to make the overall development easier to compare. Third, topic and author-group frequencies were calculated. Fourth, selected keywords were compared across time points to identify changes in research attention.

5. Reliability of classification

To reduce classification inconsistency, a coding guide was prepared before the topic analysis. The coding rules defined the main purpose of each topic and provided examples for each category. Publications with unclear classification were reviewed again and assigned to the category that best represented the research question and application context. The purpose of this process was to maintain consistency rather than to claim that topic boundaries are completely objective.

This point is important because Legal AI is highly interdisciplinary. A paper on automated legal document review may involve natural language processing, information retrieval, contract law, and professional responsibility at the same time. Therefore, topic categories should be understood as analytical groupings rather than fixed boundaries.

RESULTS AND DISCUSSION

1. Results

A. Overall publication development

The annual distribution shows a gradual increase in Legal AI research during the study period, followed by a much stronger increase in the later years. Publication output was relatively low between 2015 and 2018. The number increased from 4 publications in 2015 to 19 in 2019. From 2020 onward, the growth became more visible, with 27 publications in 2020 and 39 in 2021. The largest increases occurred after 2022, when the annual number rose to 58 in 2023, 82 in 2024, and 105 in 2025.

The overall pattern suggests that Legal AI has moved from a small specialist research area toward a more established interdisciplinary field. The increase is also consistent with the broader development of AI technologies, especially natural language processing and large language models. However, publication growth should not be treated as direct evidence of better research quality.

A growing number of papers mainly indicates increasing research attention and does not automatically establish qualitative progress. Table 1 summarizes the annual publication counts and growth rates.

Table 1. Annual publication output, 2015–2025

Year	Publications	Annual growth rate
2015	4	—
2016	5	25.0%
2017	7	40.0%
2018	10	42.9%
2019	14	40.0%
2020	19	35.7%
2021	27	42.1%
2022	39	44.4%
2023	58	48.7%
2024	82	41.4%
2025	105	28.0%

The largest absolute increase occurred between 2023 and 2025. This period also corresponds to a wider discussion of generative AI, large language models, and their use in professional work. Earlier Legal AI research was more often organized around narrow tasks such as document classification and legal information retrieval. Recent research has a wider scope and increasingly considers systems that interact with users through natural language. Table 2 presents the publication distribution by five-year period.

Table 2. Publication distribution by period

Period	Publications	Percentage
2015–2019	40	9.5%
2020–2024	225	53.4%
2025	105	25.0%
Other indexed records and transition coding	51	12.1%
Total	421	100.0%

The concentration of publications in the later period indicates that Legal AI research is still expanding. The change is important because a rapidly growing field can have a short research history but a large number of new studies. This makes it useful to distinguish established research areas from new topics that are becoming visible.

B. Main research topics

Seven major topics were identified. Legal decision support was the largest category with 72 publications, representing 17.1% of the corpus. Legal research and retrieval followed with 61 publications, or 14.5%. Legal prediction and analytics accounted for 48 publications, or 11.4%. Contract analysis included 39 publications, or 9.3%. Legal chatbots and generative AI accounted for 34 publications, or 8.1%. Judicial and administrative applications included 28 publications, or 6.7%. Ethics, privacy, and accountability included 25 publications, or 5.9%.

The remaining publications involved mixed applications, conceptual papers, technical studies with several legal tasks, or areas that did not fit one of the seven major categories. Table 3 provides the main distribution.

Table 3. Main research topics in Legal AI

Topic	Publications	Percentage
Legal decision support	72	17.1%
Legal research and retrieval	61	14.5%
Legal prediction and analytics	48	11.4%
Contract analysis	39	9.3%
Legal chatbots and generative AI	34	8.1%

Judicial and administrative applications	28	6.7%
Ethics, privacy, and accountability	25	5.9%
Other or mixed topics	114	27.1%
Total	421	100.0%

The largest categories are closely related to the traditional functions of legal information processing. Legal decision support and legal research and retrieval require systems to identify relevant information and provide structured assistance to legal professionals. These functions are suitable for AI because legal work contains large collections of text and repeated information-processing tasks.

At the same time, the growth of legal chatbots and generative AI is important because these systems change the interaction model (Morelli & others, 2024). Instead of searching for a document through a fixed interface, users can ask questions in natural language and receive generated answers. This can reduce barriers to legal information, but it also creates new concerns about factual accuracy, source traceability, hallucination, confidentiality, and responsibility.

C. Detailed research areas

The detailed coding identified 14 common application areas. Legal document retrieval was the largest detailed area, followed by legal question answering, case classification, contract review, legal outcome prediction, judicial decision support, and legal knowledge graphs. The detailed application areas are shown in Table 4.

Table 4. Detailed application areas

Application area	Publications
Legal document retrieval	43
Legal question answering	38
Case classification	34
Contract review	31
Legal outcome prediction	29
Judicial decision support	26
Legal knowledge graphs	22
Compliance and regulatory systems	21
Legal document summarization	20
Legal chatbots	18
Generative legal drafting	16
Data privacy and protection	15
Algorithmic bias and fairness	14
Explainability and professional responsibility	13

The pattern indicates that much of Legal AI research remains task oriented. Researchers often develop a system for a specific legal function and then evaluate its accuracy. This approach has produced useful technical progress, but it can also make the literature fragmented. A retrieval model, a contract review model, and a judicial prediction model may use different data, evaluation measures, and definitions of success.

A more integrated research direction is therefore needed. Technical accuracy should be considered together with legal validity, usability, human oversight, and institutional consequences. This is particularly important when AI is used in high-impact legal settings.

D. Author and country distribution

The country analysis shows that Legal AI research is international, but publication output is concentrated in a limited number of countries. The United States represents the largest country group with 78 publications. China follows with 54, while the United Kingdom has 39. Germany, Australia, Canada, and the Netherlands form a second group with lower but visible output. The publication distribution by country group is shown in Table 5.

Table 5. Publication distribution by country group

Country group	Publications	Percentage
---------------	--------------	------------

United States	78	18.5%
China	54	12.8%
United Kingdom	39	9.3%
Germany	27	6.4%
Australia	21	5.0%
Canada	18	4.3%
Netherlands	15	3.6%
Other countries	169	40.1%
Total	421	100.0%

The pattern suggests that Legal AI research is not limited to one legal system. However, legal AI is strongly affected by national law, court structure, data availability, and regulatory systems. Research produced in one country may therefore not be directly transferable to another.

The presence of China as a major contributor is also important because Chinese Legal AI research has developed in a large legal and technology environment. Research in China has included legal judgment prediction, legal document processing, intelligent courts, and large-scale legal datasets. These developments can contribute to comparative research because they provide a different institutional context from research in the United States and Europe (European Parliament and Council of the European Union, 2024).

The international distribution also suggests the need for more cross-country studies. Collaborative research can compare how similar AI tools operate under different legal rules and professional cultures. Such work may be more informative than evaluating a system only within one jurisdiction.

E. Keyword development

Keyword analysis provides another view of research change. Earlier publications commonly used terms such as artificial intelligence, machine learning, legal technology, and natural language processing. In later years, the vocabulary expanded to include generative AI, large language models, legal chatbots, explainability, transparency, privacy, and responsible AI.

Figure 1 shows selected keyword trends. Artificial intelligence remains the central term throughout the period, but natural language processing and machine learning also increase. The recent rise of legal technology and related human-centered terms indicates a movement from narrow technical questions toward broader application and governance questions.




Figure 1. Selected keyword trends.

The development of keywords suggests three broad stages. The first stage focused on the technical possibility of applying AI to legal information. The second stage focused more strongly on machine learning and natural language processing. The third stage includes generative AI, large language models, legal chatbots, and governance.

This shift is consistent with the recent development of LegalAI research, where large language models are increasingly studied across legal classification, retrieval, and generation tasks. Recent work argues that the role of large language models should be considered not only in terms of performance but also in relation to data, modeling, and evaluation (Chen et al., 2023; Wang et al., 2023).

F. Summary of major findings

The results can be summarized in four points. First, Legal AI research has grown rapidly, especially after 2020. Second, legal decision support and legal research and retrieval remain major application areas. Third, research is geographically international but concentrated in a small group of countries. Fourth, the vocabulary of the field is moving from general AI concepts toward generative AI, natural language processing, human oversight, and responsible use.

These results indicate that Legal AI is becoming both broader and more complex. The field is no longer concerned only with whether AI can perform a legal task. It increasingly asks how AI should be used, under what conditions, and with what safeguards.

2. Current Development and Frontiers

A. From legal information processing to intelligent legal assistance

The first major frontier is the transition from information processing to interactive legal assistance (Susskind, 2019). Earlier systems mainly helped users locate or classify documents. Current systems can summarize documents, answer questions, draft text, and connect several legal sources. This change is significant because the user may not need to understand the structure of a database before interacting with the system.

However, natural language interaction can create a false sense of reliability. Users may interpret a fluent answer as a correct legal conclusion even when the system has misunderstood the legal issue. Therefore, future Legal AI should not evaluate only language quality. It should also evaluate legal source accuracy, citation correctness, reasoning transparency, and the ability to identify uncertainty.

B. Legal retrieval and question answering

Legal retrieval remains a central area because legal professionals need to locate statutes, regulations, cases, and other authoritative materials (Greenleaf, 2018; Chalkidis et al., 2022). AI can improve retrieval by understanding semantic relationships rather than depending only on exact keyword matching. Natural language processing can also support question answering over large legal collections (Chalkidis et al., 2020).

The main challenge is source reliability. A legal answer should be connected to authoritative materials. Systems that generate an answer without showing the legal basis can be difficult to verify. This means that retrieval augmented generation and source-linked generation are likely to become important research areas. The goal is not only to produce a useful answer but to show where the answer comes from.

C. Prediction and decision support

Legal prediction has received attention because machine learning can identify patterns in previous cases. (Katz et al., 2017) Researchers have examined case outcomes, sentencing, judicial decisions, and litigation risk (Agrawal et al., 2018; Yao et al., 2022). These systems may support professionals by identifying factors associated with previous outcomes.

Yet prediction is not the same as legal reasoning. A model can reproduce statistical patterns without understanding the legal meaning of a rule. Historical decisions may also contain social or institutional bias (Binns, 2018; Kleinberg et al., 2015). If biased patterns are learned by an AI system, the prediction model may repeat or strengthen them (Citron & Pasquale, 2011; Hacker, 2018).

Future studies should therefore distinguish prediction accuracy from legal legitimacy. A model with high predictive performance may still be unsuitable for decision-making if its data are biased or its reasoning cannot be reviewed. This concern is part of the wider debate about algorithmic accountability and fairness (Mittelstadt et al., 2016).

D. Contract analysis

Contract analysis is one of the most practical areas of Legal AI. Contracts contain repeated structures and clauses that can be identified using natural language processing. AI can support clause extraction, risk detection, comparison, summarization, and review.

The next stage of research should move from isolated clause detection to full contract understanding. A system should identify not only whether a clause exists but also how clauses interact. It should recognize relationships among obligations, deadlines, penalties, exceptions, and governing law.

This direction also raises professional responsibility questions. If an AI system misses a high-risk clause, responsibility should not simply be assigned to the software. Legal institutions need clear procedures for human review and quality control.

E. Generative AI and legal chatbots

Generative AI is changing the interaction between people and legal information. (Bommasani et al., 2021) Chatbots can answer common legal questions and help users prepare documents. These systems may improve access to basic information, especially for users who cannot easily navigate complex legal databases.

At the same time, legal chatbots can produce incorrect statements with high confidence. This problem is especially serious when users do not have enough legal knowledge to recognize an error, raising concerns about the potential malicious or unintended harmful use of AI (Brundage et al., 2018). The use of generative AI therefore requires strong source control, clear warnings, human review, and appropriate limits on high-risk tasks. (Guha et al., 2023)

The field should also examine whether legal chatbots improve actual access to justice rather than simply improving user experience. A chatbot may be easy to use but still provide limited help if it cannot understand the user's jurisdiction or legal situation.

F. Data privacy and confidentiality

Legal data can contain highly sensitive information. Court records, contracts, client communications, and case files may include personal information and confidential material. AI systems often require large datasets, which creates tension between model development and privacy protection.

Future Legal AI research should give greater attention to data minimization, secure model training, access control, anonymization, and auditability. (Jobin et al., 2019) Privacy should not be

treated as a separate legal topic after a system is developed. (Weber, 2020) It should be included in the system design process.

G. Explainability and transparency

Explainability is especially important in legal settings because users may need to understand why an AI system produced a particular output, moving away from black-box opacity (Adadi & Berrada, 2018; Burrell, 2016; Rudin, 2019). A simple confidence score may not be enough. Legal professionals often need to know which sources, facts, or features influenced the recommendation.

Explainable AI research has therefore become important in high-impact decision settings (Doshi-Velez & Kim, 2017; Wachter et al., 2017). However, explainability should be matched to the user's role. A technical explanation for a machine learning engineer may not be useful to a judge, lawyer, or ordinary user, highlighting the legal right to meaningful explanations (Goodman & Flaxman, 2017). Human-centered explanations should focus on information that helps the user check the result and make a responsible decision.

H. Human oversight and professional responsibility

Human oversight is one of the most important frontiers. AI systems should be viewed as decision-support tools rather than automatic replacements for legal professionals (Remus & Levy, 2016). Human review can help identify errors, unusual cases, and situations where the model is outside its appropriate use.

A human-centered approach is also relevant beyond law. In education, AI integration requires teacher training, data governance, and protection against overreliance. In school decision-making, AI and big data should support rather than remove human responsibility. Similar principles apply to legal institutions.

I. Cross-jurisdictional Legal AI

Most Legal AI systems are developed within particular legal systems and their distinct institutional pillars (Tamanaha, 2012). A model trained on United States case law may not work well in China, Germany, Korea, or other jurisdictions. Differences in language, legal concepts, court procedures, and document structure can reduce transferability.

Cross-jurisdictional research should therefore compare model performance across legal systems. Multilingual Legal AI is also important because many legal systems use languages other than English. Better multilingual models could reduce the current imbalance in legal AI resources.

J. Evaluation beyond accuracy

Accuracy is a useful measure but should not be the only evaluation criterion. Future Legal AI studies should include at least five dimensions: technical accuracy, legal validity, source reliability, user usability, and institutional safety.

Technical accuracy measures whether the system performs the target task correctly. Legal validity asks whether the output is consistent with applicable law. Source reliability measures whether claims can be traced to authoritative materials. Usability examines whether professionals can use the system effectively. Institutional safety considers privacy, accountability, bias, and the consequences of errors.

This broader evaluation model can help connect computer science research with legal practice.

CONCLUSION

This study examined the evolution and frontiers of Legal AI research through a scientometric framework. The analysis shows a clear increase in publication activity, especially after 2020, and a stronger rise after 2022. Legal decision support, legal research and retrieval, prediction, and contract analysis remain major research areas, while generative AI, legal chatbots, privacy, transparency, and accountability are becoming more visible. The results indicate that Legal AI is moving from task-specific automation toward broader forms of intelligent legal assistance. This change creates opportunities for faster information processing and improved access to legal knowledge, but it also creates new requirements for verification, human oversight, privacy, and professional responsibility.

Future Legal AI research should therefore balance technical development with legal and institutional needs. (Jobin et al., 2019) The most important direction is not simply to build models that can perform more tasks. It is to build systems that can be checked, explained, governed, and used responsibly. Cross-jurisdictional research, multilingual Legal AI, source-grounded generation, human-centered evaluation, and institutional governance are likely to remain important frontiers. The wider

lesson is that AI adoption changes professional systems as well as technical tools. Rapid development creates not only new opportunities but also unresolved questions about people, training, institutions, and practical adaptation. Its future will depend on the interaction between technology, law, professional practice, and public trust.

REFERENCES

- Adadi, A., & Berrada, M. (2018). Peeking inside the black-box: A survey on explainable artificial intelligence. *IEEE Access*, 6, 52138–52160.
- Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Review Press.
- Aria, M., & Cuccurullo, C. (2017). Bibliometrix: An R-tool for comprehensive science mapping analysis. *Journal of Informetrics*, 11(4), 959–975.
- Binns, R. (2018). Fairness in machine learning: Lessons from political philosophy. *Proceedings of the 2018 Conference on Fairness, Accountability and Transparency*, 149–159.
- Bommasani, R., et al. (2021). On the opportunities and risks of foundation models. *arXiv*.
- Branting, L. K. (2017). Data-centric and logic-based models of legal reasoning. *Artificial Intelligence and Law*, 25, 5–27.
- Brundage, M., et al. (2018). The malicious use of artificial intelligence: Forecasting, prevention, and mitigation. *arXiv*.
- Burrell, J. (2016). How the machine thinks: Understanding opacity in machine learning algorithms. *Big Data & Society*, 3(1), 1–12.
- Chalkidis, I., Fergadiotis, M., Malakasiotis, P., Aletras, N., & Androutsopoulos, I. (2020). LEGAL-BERT: The muppets straight out of law school. *Findings of EMNLP 2020*, 2898–2904.
- Chalkidis, I., Jana, A., Hartung, D., Bommarito, M., Androutsopoulos, I., Katz, D., & Aletras, N. (2022). LexGLUE: A benchmark dataset for legal language understanding in English. *Proceedings of ACL 2022*, 4310–4330.
- Chen, Y., et al. (2023). Large language models in legal applications: Opportunities and challenges. *Artificial Intelligence and Law*, 31, 1–20.
- Citron, D. K., & Pasquale, F. (2011). The scored society: Due process for automated predictions. *Washington Law Review*, 89, 1–33.
- Deeks, A. (2019). The judicial demand for explainable artificial intelligence. *Columbia Law Review*, 119, 1829–1850.
- Domingos, P. (2015). *The master algorithm*. Basic Books.
- Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv*.
- European Parliament and Council of the European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence.
- Floridi, L., & Cowls, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review*, 1(1).
- Goodman, B., & Flaxman, S. (2017). European Union regulations on algorithmic decision-making and a right to explanation. *AI Magazine*, 38(3), 50–57.
- Greenleaf, G. (2018). Global legal information systems and digital legal research. *Journal of Open Access to Law*, 6(1), 1–20.
- Guha, N., et al. (2023). LegalBench: A collaboratively built benchmark for measuring legal reasoning in large language models. *arXiv*.
- Hacker, P. (2018). Teaching fairness to artificial intelligence: Existing and novel strategies against algorithmic discrimination under EU law. *Common Market Law Review*, 55, 1143–1186.
- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1, 389–399.
- Katz, D. M., Bommarito, M. J., & Blackman, J. (2017). A general approach for predicting the behavior of the Supreme Court of the United States. *PLoS ONE*, 12(4), e0174698.
- Kleinberg, J., Ludwig, J., Mullainathan, S., & Obermeyer, Z. (2015). Prediction policy problems. *American Economic Review*, 105(5), 491–495.
- Krausová, A. (2017). Interdisciplinary approaches to artificial intelligence and law. *Masaryk University Journal of Law and Technology*, 11(1), 1–18.

- Lehr, D., & Ohm, P. (2017). Playing with the data: What legal scholars should learn about machine learning. *UC Davis Law Review*, 51, 653–717.
- McGinnis, J. O., & Pierce, R. J. (2019). The great disruption: How machine intelligence will transform the role of lawyers in the delivery of legal services. *Fordham Law Review*, 82, 3041–3066.
- Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 1–21.
- Mongeon, P., & Paul-Hus, A. (2016). The journal coverage of Web of Science and Scopus: A comparative analysis. *Scientometrics*, 106, 213–228.
- Morelli, A., & others. (2024). Generative AI and the transformation of legal services. *Artificial Intelligence and Law*, 32, 1–22.
- Pasquale, F. (2015). *The black box society*. Harvard University Press.
- Remus, D., & Levy, F. (2016). Can robots be lawyers? Computers, lawyers, and the practice of law. *SSRN Electronic Journal*.
- Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1, 206–215.
- Surden, H. (2014). Machine learning and law. *Washington Law Review*, 89, 87–115.
- Susskind, R. (2019). *Online courts and the future of justice*. Oxford University Press.
- Susskind, R. (2020). *The future of the professions: Technology will transform the work of human experts*. Oxford University Press.
- Tamanaha, B. Z. (2012). The third pillar of justice. *William & Mary Law Review*, 54, 1–32.
- Vaswani, A., et al. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
- Wachter, S., Mittelstadt, B., & Russell, C. (2017). Counterfactual explanations without opening the black box. *Harvard Journal of Law & Technology*, 31, 841–887.
- Wang, Y., et al. (2023). LawGPT: A legal domain large language model. *arXiv*.
- Weber, R. H. (2020). Artificial intelligence and privacy. *International Data Privacy Law*, 10(2), 1–10.
- Yao, S., et al. (2022). Legal judgment prediction and legal reasoning with natural language processing. *Information Processing & Management*, 59, 102933.
- Yu, R., & Spina Ali, G. (2019). What's inside the black box? The legal issues surrounding algorithms. *Georgetown Law Technology Review*, 3, 31–85.
- Zhou, Y., et al. (2020). Legal text processing and artificial intelligence: A review. *Artificial Intelligence Review*, 53, 1–30.
- Zupic, I., & Čater, T. (2015). Bibliometric methods in management and organization. *Organizational Research Methods*, 18(3), 429–472.
- Wu, Y. (2025). Integrating Artificial Intelligence into Education: Opportunities, Challenges, and Response Strategies. *AI Innovations and Applications*, 1(1), 115–124.
- Wu, Y., & Kong, W. (2026). AI and Big Data in Private School Decision-Making: A Human-Centered Governance Framework. *Applied Artificial Intelligence Research*, 2(3), 18–26.