
THE CODIFICATION OF CODE: HOW AI IS RESHAPING LEGAL RESEARCH

Prama, Research Scholar, BJR Institute of Law, Bundelkhand University, Jhansi, India

Awadhesh Kumar Shukla, Research Scholar, BJR Institute of Law, Bundelkhand
University, Jhansi, India

Sudhanshu Bajpai, Junior Library Superintendent, Indian Institute of Technology (BHU),
Varanasi

ABSTRACT

This bibliometric study examines the intersection of artificial intelligence (AI) and law, mapping key research trends and influential contributions. A dataset of over 11,000 publications was analysed using citation analysis, co-authorship network mapping, and keyword co-occurrence techniques. The analysis reveals exponential growth in AI-law research output in the last decade and identifies a core set of highly cited works that have shaped the field's trajectory. Three major thematic clusters emerge: (1) AI for legal decision-making and knowledge systems – encompassing expert systems, legal reasoning, and text analytics; (2) AI in law enforcement and criminal justice – including crime prediction, risk assessment, and policing applications; and (3) Ethical, legal, and policy implications of AI – covering algorithmic fairness, data privacy, and AI regulation. Co-authorship and country collaboration networks show that research is dominated by a few countries (notably China and the USA) with moderate international collaboration. These findings demonstrate how bibliometric methods can quantitatively characterize the development of AI in law and highlight emerging trends (such as AI ethics and governance) that are likely to define future research. The study provides a structured overview of the field's evolution, influential authors and sources, collaborative patterns, and thematic shifts, offering valuable insights for researchers and policymakers interested in the nexus of AI and legal studies.

Introduction

Artificial Intelligence is increasingly transforming the legal domain, from automating routine tasks to informing high-stakes judicial and policy decisions. AI in law (or legal AI) refers to the application of AI technologies – such as machine learning, natural language processing, and knowledge-based systems – to legal problems (e.g. assisting lawyers, predicting case outcomes, or analysing legal documents) as well as the study of legal implications of AI (such as governance, ethics, and accountability of AI systems). Understanding the trajectory of research in this interdisciplinary field is important because it reflects how legal scholarship and practice are adapting to technological advances. Bibliometric analysis, which applies quantitative methods to scientific literature, is a powerful approach to uncover patterns and trends in a research domain. By systematically examining publication metadata (papers, authors, citations, keywords, etc.), bibliometrics can identify the evolution of a discipline, key contributors, and emerging hotspots of research. In the context of AI and law – a field straddling computer science, law, and ethics – a bibliometric study allows us to map its development from early conceptual works to the current landscape of big data analytics and AI regulation. This paper employs a comprehensive bibliometric methodology to analyse the global research output on AI in law. We use performance indicators (e.g. publication counts, citation counts, h-index) to gauge productivity and impact, and science mapping tools (citation networks, co-authorship networks, co-word analysis) to visualize the intellectual structure of the field. The goal is to shed light on how AI and law research has grown and transformed over time, who the leading authors and institutions are, what themes and technologies have been most prominent, and how collaboration and knowledge networks are configured. In the following sections, we describe our data collection and analysis methods, present findings from citation analysis, keyword trends, and collaboration networks, and discuss the thematic evolution of AI-law research. This provides a data-driven overview of the field's past and present, laying the groundwork for understanding future directions in AI and legal studies.

Literature Review

Artificial intelligence (AI) presents transformative potential in the legal field, impacting diverse aspects such as legal research, law enforcement, adjudication, ethics, and governance. The literature converges on several key themes: enhancement of legal processes, integration in law enforcement, ethical and regulatory challenges, and the need for interdisciplinary

collaboration.

Enhancement of Legal Processes

AI technologies are increasingly deployed to streamline and augment traditional legal workflows. AI-powered tools significantly reduce the time lawyers spend on legal research by swiftly analysing vast volumes of case law and identifying relevant precedents (Zakir et al¹, 2024; Rajendra & Thuraisingam, 2022²). Contract analysis is another vital area where AI assists in detecting compliance issues, uncovering risk factors, and improving accuracy in document review (Xudaybergenov, 2023³; Ramachandran & Rana, 2024⁴). Studies also affirm AI's growing role in automating data-intensive legal tasks traditionally handled by humans, thereby enhancing operational efficiency and accuracy in legal decision-making (Lunhol & Torhalo, 2024⁵; Beegum et al., 2023⁶). This aligns with Surden's (2019)⁷ realist approach that demystifies AI's actual capabilities, distinguishing between functional AI applications and speculative narratives. Nikolskaia and Naumov (2020)⁸ similarly highlight that while AI cannot imitate legal reasoning fully, it serves as a valuable support tool in managing routine legal work.

AI in Law Enforcement

AI's role extends into the domain of law enforcement, particularly in predictive policing, surveillance, and anomaly detection. Raaijmakers (2019)⁹ outlines AI's applications in suspect profiling, automated traffic monitoring, and child protection. Lunhol and Torhalo (2024)⁵

¹ Zakir, M. H., Bashir, S., Ali, R. N., & Khan, S. H. (2024). Artificial intelligence and machine learning in legal research: a comprehensive analysis. *Qlantic Journal of Social Sciences*, 5(1), 307-317. <https://doi.org/10.55737/qjss.203679344>

² Rajendra, J. B. and Thuraisingam, A. S. (2022). Artificial intelligence and its impact on the legal fraternity. *UUM Journal of Legal Studies*, 13. <https://doi.org/10.32890/uumjls2022.13.2.6>

³ Xudaybergenov, A. (2023). Toward legal recognition of artificial intelligence proposals for limited subject of law status. *International Journal of Law and Policy*, 1(4). <https://doi.org/10.59022/ijlp.55>

⁴ Ramachandran, D. and Rana, R. S. (2024). Ai-driven jurisprudence: navigating legal landscapes in the digital age. *International Journal of Law, Justice and Jurisprudence*, 4(1), 128-139. <https://doi.org/10.22271/2790-0673.2024.v4.i1b.103>

⁵ Lunhol, O. and Torhalo, P. (2024). Artificial intelligence in law enforcement: current state and development prospects. *Socratic Lectures 10 - Part II*. <https://doi.org/10.55295/psl.2024.ii12>

⁶ Beegum, S., Antony, B., & Kumar, S. (2023). The Impact of Artificial Intelligence on Law

⁷ Surden, H. (2019). Artificial Intelligence and Law: An Overview. *Georgia State University law review*, 35, 15109.

⁸ Nikolskaia, K., & Naumov, V. (2020). Artificial Intelligence in Law. *2020 International Multi-Conference on Industrial Engineering and Modern Technologies (FarEastCon)*, 1-4. <https://doi.org/10.1109/FarEastCon50210.2020.9271095>.

⁹ Raaijmakers, S. (2019). Artificial Intelligence for Law Enforcement: Challenges and Opportunities. *IEEE Security & Privacy*, 17, 74-77. <https://doi.org/10.1109/MSEC.2019.2925649>

report that AI systems, including facial recognition and crime prediction algorithms, offer advantages in resource allocation and proactive policing. However, these technologies raise serious ethical concerns about civil liberties, potential misuse, and algorithmic bias. Walters and Novak (2021)¹⁰ stress the absence of international standards governing such applications, especially in protecting vulnerable populations like children and minorities from unjust surveillance practices.

Ethical and Legal Challenges

As AI systems influence legal outcomes, ethical and legal scrutiny intensifies. Scholars have raised concerns regarding accountability, data privacy, algorithmic transparency, and the inadequacy of current legal frameworks (Kumar & Suthar, 2024¹¹; Kudeikina & Kaija, 2024¹²). The challenge of assigning liability for AI-driven decisions—especially when errors occur—remains unresolved (Xudaybergenov, 2023)³. Consequently, researchers recommend the creation of regulatory bodies to develop legal and ethical standards to guide AI's implementation in judicial systems (Ramachandran & Rana, 2024⁴). Chucha (2023)¹³ adds a psychological lens, suggesting that public trust in AI-assisted justice requires mechanisms for verifying decisions in legally and psychologically acceptable ways.

Theoretical, Regulatory, and Conceptual Developments

The absence of a unified legal definition of AI complicates efforts to regulate its use across jurisdictions. Walters and Novak (2021)¹⁰ reveal inconsistencies in defining what constitutes AI, especially in sensitive areas like military applications or consumer surveillance. Ponkin et al. (2018)¹⁴ go further by discussing the idea of AI as a potential “electronic person” and stress the need for legal systems to address AI's capacity for self-learning and autonomous decision-

¹⁰ Walters, R., & Novak, M. (2021). Artificial Intelligence and Law. *Cyber Security, Artificial Intelligence, Data Protection & the Law*. https://doi.org/10.1007/978-981-16-1665-5_3.

¹¹ Kumar, D. and Suthar, N. (2024). Ethical and legal challenges of ai in marketing: an exploration of solutions. *Journal of Information, Communication and Ethics in Society*, 22(1), 124-144. <https://doi.org/10.1108/jices-05-2023-0068>

¹² Kudeikina, I. and Kaija, S. (2024). Limits of the use of artificial intelligence in law – ethical and legal aspects. *ENVIRONMENT. TECHNOLOGIES. RESOURCES. Proceedings of the International Scientific and Practical Conference*, 2, 188-191. <https://doi.org/10.17770/etr2024vol2.8016>

¹³ Chucha, S. (2023). Artificial intelligence in justice: legal and psychological aspects of law enforcement. *Law Enforcement Review*. [https://doi.org/10.52468/2542-1514.2023.7\(2\).116-124](https://doi.org/10.52468/2542-1514.2023.7(2).116-124).

¹⁴ Ponkin, I., Владиславович, П., Redkina, A., & Игоревна, П. (2018). Artificial Intelligence from the Point of View of Law. , 22, 91-109. <https://doi.org/10.22363/2313-2337-2018-22-1-91-109>.

making. Such theoretical inquiries are foundational to creating laws that are technologically adaptive yet ethically grounded.

Interdisciplinary Collaboration and Future Directions

Addressing the multifaceted implications of AI in legal contexts demands an interdisciplinary approach. Collaboration among legal professionals, ethicists, technologists, and policymakers is essential to develop frameworks that protect fundamental rights while enabling innovation (Ramachandran & Rana, 2024⁴). Ashley (2017)¹⁵ envisions legal analytics as the future of jurisprudence, where computational models generate, evaluate, and explain legal arguments. Reiling and McCarthy (2020)¹⁶ emphasise that court systems must reconcile AI's operational efficiency with the procedural safeguards enshrined in human rights law.

In sum, the literature underscores that AI is no longer a speculative concept in law but a practical, evolving force that requires robust legal, ethical, and societal responses. As AI systems become embedded in legal institutions, ongoing research and policy dialogue are imperative to ensure justice systems remain transparent, accountable, and equitable.

Need for Bibliometric Analysis

The need for bibliometric analysis of artificial intelligence (AI) in law is significant as it offers several advantages in understanding the evolving landscape of this interdisciplinary field. Bibliometric analysis is a quantitative tool for analysing scientific publications, providing insights into trends, impact, and the development of knowledge within AI in law.

1. Mapping Research Trends: Bibliometric analysis systematically evaluates the literature related to AI in law, enabling researchers to identify trends over time, including prolific authors, influential publications, and prevailing themes. This allows for a historical perspective on how AI has been adopted within legal contexts and sheds light on emerging research areas that require further investigation (Batra et al., 2022)¹⁷. This capability is crucial

¹⁵ Ashley, K. (2017). Artificial Intelligence and Legal Analytics. . <https://doi.org/10.1017/9781316761380>

¹⁶ Reiling, A., & McCarthy, J. (2020). Courts and Artificial Intelligence. *International Journal for Court Administration*, 11. <https://doi.org/10.36745/ijca.343>.

¹⁷ Batra, S., Saini, M., Yadav, M., & Aggarwal, V. (2022). Mapping the intellectual structure and demystifying the research trend of cross listing: a bibliometric analysis. *Managerial Finance*, 49(6), 992-1016. <https://doi.org/10.1108/mf-07-2022-0330>

for scholars and practitioners seeking to align their work with current developments and gaps in the literature.

2. **Identifying Key Contributors and Institutions:** Employing bibliometric methods helps map the intellectual structure of AI research in law, pinpointing key contributors and institutions that dominate the field. This information can help emerging researchers identify leading scholars and institutions to collaborate with or seek mentorship from (Shim et al., 2017)¹⁸. Additionally, understanding which journals frequently publish relevant work can guide where to submit research findings within this area (Batra et al., 2022)¹⁷.

3. **Enhancing Interdisciplinary Collaboration:** AI's application in law intersects with various disciplines, and bibliometric analysis illustrates these interdisciplinary connections. By visualizing networks of citations and collaborations, researchers can identify external influences, such as insights from computer science or ethics, thus promoting broader collaborative efforts that can lead to innovative solutions to complex legal challenges

4. **Evaluating Research Impact and Quality:** Bibliometric studies often quantify citations, enabling an assessment of which works have significantly influenced the legal and AI research communities (Ellegaard & Wallin, 2015)¹⁹. This data can inform researchers, funding bodies, and policymakers about influential areas within AI in law, guiding resource allocation and future research initiatives. Furthermore, understanding the methodologies employed in influential studies equips scholars with foundational knowledge that can inform their work.

5. **Highlighting Ethical and Legal Considerations:** As AI technologies increasingly integrate into legal frameworks, it is essential to assess the ethical and legal implications of this integration. Bibliometric analysis can help surface discussions on ethics within AI law, balancing technological advancement with fundamental human rights (Hu et al., 2024)²⁰. By doing so, it fosters a comprehensive approach to researching and implementing AI within the

¹⁸ Shim, J., Bliemel, M., & Choi, M. (2017). Modeling complex entrepreneurial processes. *International Journal of Entrepreneurial Behavior & Research*, 23(6), 1052-1070. <https://doi.org/10.1108/ijebr-11-2016-0374>

¹⁹ Ellegaard, O. and Wallin, J. A. (2015). The bibliometric analysis of scholarly production: how great is the impact?. *Scientometrics*, 105(3), 1809-1831. <https://doi.org/10.1007/s11192-015-1645-z>

²⁰ Hu, Y., Yu, L., Du, W., Hu, X., & Shen, Y. (2024). Global hotspots and research trends of radiation-induced skin injury: a bibliometric analysis from 2004 to 2023. *Frontiers in Oncology*, 14. <https://doi.org/10.3389/fonc.2024.1430802>

legal domain, ultimately ensuring responsible and equitable deployment.

6. Future Research Directions and Hotspots: Through bibliometric analysis, researchers can identify both established trends and potential future directions for AI research in the legal sphere. This foresight is crucial for guiding new inquiries and directing attention to underexplored areas that may become increasingly significant as technology evolves. By visualising this data, researchers can formulate strategic plans that align with broader scientific discourse.

In conclusion, bibliometric analysis serves as a cornerstone for a comprehensive understanding and advancement in the field of AI in law. It illuminates research trends and contributions while fostering interdisciplinary dialogue, ultimately enhancing the quality and impact of legal scholarship involving AI technologies

Research Methodology

A. Data Sources and Collection: The study began by defining the scope of “AI and law” research and retrieving relevant publications. We conducted an extensive search in bibliographic databases (primarily Scopus) to gather literature at the intersection of artificial intelligence and law. The search strategy used a combination of AI-related keywords (“artificial intelligence”, “machine learning”, “expert system”, “natural language processing”, “deep learning”, etc.) and law-related keywords (“law”, “legal”, “legislation”, “justice”, “court”, “policing”, etc.), applied to titles, abstracts, and author keywords. We included journal articles, conference papers, and reviews, and we imposed no strict start date in order to capture the earliest works. The initial search results were then refined with a PRISMA-like approach: we removed clearly irrelevant records (e.g. papers where “Law” was an author’s surname or referring to scientific laws rather than legal context) and duplicates. We also excluded non-scholarly items and retracted papers to ensure data quality. The final dataset comprised 11,281 documents, spanning 1911 to early 2025. This broad timeframe covers over a century of publications, though the significant research activity begins in the late 20th century (as discussed later). Each record in the dataset contained metadata including title, authors, affiliations, source (journal or conference), year, keywords, and citation counts. These metadata were exported and processed using the bibliometrix R package and its Biblioshiny application, which facilitated cleaning and analysis. Steps were taken to standardise author names (to

address variations or initials), unify institutional names, and consolidate keywords (e.g. merging singular and plural forms, handling synonyms) for accurate analysis.

B. Bibliometric Analysis Tools: We employed both performance analysis and science mapping techniques. Performance analysis assesses productivity and impact through metrics like publication counts, citation counts, and the h-index. The h-index was calculated for authors in the dataset to identify influential researchers – an author has index h if they have at least h papers each with $\geq h$ citations. We also examined average citations per paper and the total citations of articles, which together highlight high-impact work. Key bibliometric indicators used include: total publications by author, institution, country; total citations and average citations; and each entity's h-index within the dataset. We identified top contributing authors, institutions, countries, and sources (journals/proceedings) by these measures. Science mapping techniques were used to explore the relationships within the field. A citation analysis was performed to see how articles cite each other and to pinpoint seminal works. We built citation networks and identified the most cited references. A co-authorship network was constructed to analyse collaboration patterns: in this network, nodes represent countries, and an edge between two nodes indicates they co-authored at least one publication. By visualizing co-authorship clusters, we can see the community structure of the field and key collaboration hubs. To map the conceptual structure, we conducted a co-word analysis using author keywords. Frequently occurring keywords indicate major topics, and by mapping co-occurrence of keywords in documents, we identified thematic clusters of research. We used network modularity algorithms to detect clusters of interrelated terms, which correspond to subtopics within AI and law. A strategic diagram (thematic map) was also reviewed to understand the centrality and density of these themes (classifying them as motor themes, niche themes, emerging or declining topics). Finally, a trend analysis was done to observe how keywords rise or fall in usage over time – highlighting emerging topics. We broke the timespan into sub-periods to perform a thematic evolution analysis, tracking how clusters in earlier periods evolve (merge, split, or vanish) in later periods. All analyses were conducted with bibliometrix's functions and results (tables, networks, and graphs) were exported for interpretation. The combination of these tools provides both quantitative metrics and visual maps, ensuring a comprehensive understanding of the field's performance and structure.

C. Performance Indicators: The dataset's main information indicates an annual growth rate of 4.66% in publications, but this average masks a recent surge (detailed in the thematic

evolution section). The overall average citations per document is 13.6, though highly cited papers push the upper tail of the distribution. We calculated that the total citations received by all AI-law papers in the dataset exceed 75,000 (with a median of 2 citations per paper, indicating many papers are still new or niche). The collective h-index of the field (for all papers considered as one set) is 109, meaning at least 109 papers have 109+ citations. At the author level, the highest individual h-index in this field (within our dataset) was 34. We also note that there are 23,754 authors represented, of whom about 2,847 (12%) authored single-authored documents (the rest collaborated). The co-authors per document is 2.92 on average, reflecting a moderate collaboration intensity. About 17% of the publications involve international co-authorship, signalling a moderate level of global collaboration in AI and law research (similar to many interdisciplinary fields). These indicators set the stage for deeper analysis: the next sections delve into the results of citation analysis, keyword analysis, collaboration networks, and thematic trends.

Citation Analysis This section identifies the most influential publications, authors, and journals in AI and law research based on citation metrics. Citations serve as a proxy for impact, and a small set of highly cited works often forms the knowledge base of a field.

Most Cited Publications: Table 1 lists the top five most cited documents in the AI-law dataset, along with their citation counts. Notably, the highest cited works are foundational AI research that have been widely referenced in the legal-AI literature, underscoring the interdisciplinary nature of the field.


Table 1. Top five most cited publications on AI in Law.

Rank	Publication (citation details)	Total Citations
1.	Zadeh, L. A. (1975). Information Sciences. – Introduced fuzzy logic for approximate reasoning.	11,953
2.	Schmidt, M., & Lipson, H. (2009). Science. – Discovered equations governing natural systems (“free-form laws”).	2,017
3.	Feldman, M., Friedler, S. A., Moeller, J., et al. (2015). KDD Proc. – Proposed algorithm to detect and reduce bias (“disparate impact”).	1,192

Rank	Publication (citation details)	Total Citations
4.	Mitchell, M., Wu, S., et al. (2019). ACM FAT Proc.* – Introduced “model cards” for AI transparency and fairness.	1,129
5.	Kusner, M. J., Loftus, J., et al. (2017). NeurIPS Proc. – Formulated the concept of “counterfactual fairness” in algorithmic decisions.	927

These highly cited works highlight key foundations and concerns of AI in the context of law. Zadeh (1975), with nearly 12k citations, is a seminal work on fuzzy logic – a mathematical framework for reasoning under uncertainty. Its influence on legal AI comes from providing a way to handle the vagueness inherent in legal reasoning (e.g. interpreting linguistic terms in regulations). The presence of this paper atop the list indicates that core AI methodologies (here, fuzzy sets) have profoundly informed later research on legal decision support systems. Schmidt & Lipson (2009), published in *Science*, demonstrated how AI can derive scientific laws from data. Although it addresses laws of nature, its inclusion (with 2,017 citations) reflects interdisciplinary crossover – the idea of automated discovery of patterns (“laws”) has inspired work on discovering legal rules or predictors from case data. Moving to more recent work, Feldman et al. (2015) (1,192 citations) is a pioneering study in algorithmic fairness. It introduced a method to detect and mitigate “disparate impact” bias in machine learning models, directly relevant to legal concerns about discrimination in AI systems. Likewise, Mitchell et al. (2019) (1,129 citations) and Kusner et al. (2017) (927 citations) are influential contributions from the burgeoning AI ethics and fairness arena. Mitchell et al. proposed Model Cards as a standard for AI system transparency (important for accountability and regulatory compliance), and Kusner et al. developed the notion of counterfactual fairness, linking causal reasoning with legal definitions of fairness. The high citation counts (nearly 1k each, in just a few years) signal how rapidly this subfield has grown. In summary, the most cited publications underscore two driving forces in AI-law research: foundational AI techniques (fuzzy logic, machine learning) and the emergent focus on ensuring AI systems comply with legal and ethical standards (fairness, transparency). These works have become touchstones for subsequent studies and are frequently cited as the conceptual or methodological basis in the literature.

Leading Authors:



The field of AI and law is supported by a diverse group of scholars from computer science, law, and related disciplines. In our analysis, the most prolific authors (by number of publications) include several with common surnames such as “Wang”, “Zhang”, “Liu”, and “Li”, each appearing 50+ times. For example, Wang Y. (which may correspond to multiple individuals) has 88 publications, and Zhang Y. has 69. These counts are partially an artifact of name aggregation (many different authors share these surnames and first initials), but they indicate a very strong contribution from Chinese researchers, as such surnames are common in China. When considering individual scholars, a few renowned researchers stand out. Trevor Bench-Capon (UK) and Giovanni Sartor (Italy) each have around 40 publications in the dataset, reflecting decades-long involvement in AI & Law research (both have been active since the 1980s-90s and contributed foundational work on legal reasoning and argumentation). Kevin D. Ashley (USA) is another prominent figure (with 31 publications), known for case-based reasoning in law. Other notable contributors include Edwina Rissland (USA) who was a pioneer in integrating AI and legal reasoning, and Bart Verheij (Netherlands) in argumentation, among many others. The presence of these established names alongside a long tail of authors who have a few papers each shows that the field has a core group of experts as well as broad participation. In terms of impact, some of these authors also have high citation counts and h-


index values. For instance, several of Bench-Capon's and Sartor's papers are highly cited in the AI & law community, giving them h-index values in the 20s or 30s within this domain. It is also notable that many prolific authors are affiliated with leading institutions discussed below, illustrating how individual and institutional outputs are connected.

Leading Journals and Conferences:




Research on AI in law is disseminated across a mix of specialized journals and conferences. The most active publication venues (sources) in our dataset are dominated by computer science conference proceedings, reflecting the field's technical side. The single largest source is the Lecture Notes in Computer Science (LNCS) series, with 639 papers. LNCS covers proceedings of numerous AI-related conferences and workshops, indicating that a substantial portion of AI-law research appears in conference form (which is common in computer science). For example, the biennial International Conference on Artificial Intelligence and Law (ICAIL) – the flagship conference started in 1987 – contributes a large number of papers (its proceedings account for 227 papers in our data). Another important series is the JURIX conference (International

Conference on Legal Knowledge and Information Systems), often published in IOS Press's Frontiers in Artificial Intelligence and Applications (87 papers). Besides these, general AI conference series such as AAAI, IJCAI, and NeurIPS also appear, as they occasionally include law-related studies (e.g. on fairness or legal text analysis). Among journals, the top contributor is the dedicated Artificial Intelligence and Law journal (Springer), with 166 papers. This journal, established in 1992, has been a central forum for the community, publishing research on all aspects of AI applied to legal problems. Its prominence in both quantity and influence is clear – it has an h-index of 33 within our dataset and over 3,600 citations to its articles, making it a core outlet. Another notable journal is Computer Law & Security Review (81 papers), which approaches AI from the legal scholarship side (covering topics like data protection, cyber law, and AI policy). The journal AI & Society (50 papers) provides an interdisciplinary venue focusing on societal and ethical implications of AI, including legal aspects. Additionally, law-specific journals and general science/engineering journals have occasionally featured impactful AI-law articles – for example, IEEE Transactions on Systems, Man, and Cybernetics and Expert Systems with Applications each have a handful of highly cited papers on legal applications of AI. In terms of citation impact, we observe that some of the highest citation averages come from general science or AI venues: e.g., papers in Science or IEEE Transactions (as seen in Table 1) garnered hundreds or thousands of citations, even if those venues contributed fewer papers in count. Meanwhile, within law-focused outlets, Artificial Intelligence and Law journal's articles have a healthy citation rate (averaging 22 citations per paper) and Computer Law & Security Review averages 21, indicating that work published in these specialized journals does get referenced by the community. Overall, the dispersion of sources illustrates the interdisciplinary nature of AI and law research – computer science conferences drive technical innovations, while law and society journals drive doctrinal and ethical discussions. This duality ensures the field progresses on both technological and legal fronts.

Citation Trends:

The citation trend over time reflects both the growth of the field and the lasting influence of key publications. Early foundational works from the 1980s and 1990s (for example, early expert system papers or legal reasoning models) accumulated citations steadily over the subsequent decades. Many of these have become “classic” references in the AI & law canon, continually cited as the field expands. From the mid-2000s onward, as publication volume increased, the total citations per year rose significantly. We find that around 2015 there was a steep uptick in both the number of publications and citations. This corresponds to new landmark papers (such as those on fairness and transparency in AI) that quickly attracted attention. For instance, a 2016 paper on predicting court decisions or a 2017 paper on AI fairness might accumulate dozens of citations per year soon after publication. By 2020, the field’s citation counts were being boosted not only by the aging of older influential work but by rapid citation of new work. The average citations per document in the dataset (13.6) is moderate, but skewed by a few extremely cited papers – indeed, we observe a “long-tail” distribution where a small core of publications (like the top 1–2% listed in Table 1) account for a large share of total citations. Most papers (especially recent ones) have under 10 citations so far. This indicates that while interest in AI-law is broad, the guiding literature consists of


relatively few but very influential pieces, often comprehensive surveys or methods that many others build upon. Another pattern is the role of survey and review articles as “citation magnets.” Much like other fields, broad overview articles (for example, a 2020 survey on bias in AI or a 2019 law review on AI governance) tend to garner a lot of citations because they serve as entry points for researchers. We see evidence of this in our data – such papers have high citation counts relative to original research papers of similar years. In summary, the citation analysis shows a field maturing: early seminal ideas continuing to be acknowledged, and new critical concerns (ethics, fairness, etc.) rapidly crystallizing into highly cited references. It underlines the importance of interdisciplinary contributions – AI breakthroughs and legal analyses – in pushing the domain forward, and suggests that future influential work may arise from bridging technical advances with legal insight.



Keyword Trend Analysis: The vocabulary used in AI and law publications provides insight into the prevailing research themes and how they have evolved. We analysed the frequency of keywords and their co-occurrence patterns to identify common topics and emerging trends in the field.

Common Keywords: As expected, the term “artificial intelligence” is ubiquitous – it appears over 6,000 times in the corpus, underscoring that AI is the central focus of every document. In contrast, the generic term “law” appears 332 times, often in contexts like “law enforcement” or “laws and legislation”. Interestingly, the composite term “laws and legislation” (an index term referring to legal statutes/regulation) appears 1,481 times, indicating that a significant subset of the literature deals with legal rules, regulatory frameworks, or the impact of AI on legislation. This points to strong representation of scholarship on how AI is governed by or used to navigate legal mandates. Beyond those fundamental terms, several clusters of keywords emerge:

Tree



- AI Techniques and Paradigms:** Keywords related to machine learning are very prominent. “Machine learning” itself appears 525 times (with an additional 378 occurrences of the hyphenated “machine-learning”), reflecting the surge of data-driven approaches in legal tech. Specific methodologies like “deep learning” (450 occurrences) and “neural networks” (416) rank high, showing that modern AI methods

(especially since the mid-2010s) have permeated the field. Older AI paradigms are also present: “expert systems” (161 occurrences) and “knowledge-based systems” (229) appear frequently, echoing the early era of rule-based legal AI. We also see “natural language processing” (239 occurrences, sometimes phrased as “NLP systems”), indicating the importance of text analysis in law (unsurprising given legal sources are text-heavy). Other technical terms include “algorithms” (376) and “data mining” (278), highlighting a focus on data analysis techniques, and “optimization” (205) and “classification” (180) hinting at common tasks in AI applications (e.g., classifying legal documents or optimizing resource allocation in legal processes). The term “robotics” (255) also appears, likely in context of robotics law or autonomous systems, and “automation” (201) pointing to automated legal processes (like contract review automation).


- **Legal Application Areas:** A major theme is AI in the justice system and public safety. The keyword “law enforcement” is one of the top non-AI terms (390 occurrences), indicating many studies on using AI for policing, crime prevention, and related law enforcement tasks. Similarly, “crime” appears 300 times, and more specific terms like “police”, “criminal justice”, or specific crime types are present (e.g., “cybercrime” appears often in context of AI for cybersecurity law). We also see “decision support systems” (410 occurrences) which often refers to tools for judges or lawyers to make decisions, and this aligns with terms like “risk assessment” (202) – likely referring to risk assessment instruments (for bail, parole, etc.) that use AI. Another important area is judicial and legal analytics: terms such as “court” (appearing in combinations like “court decisions” 26 times, “Supreme Court” 28 times) show interest in AI predicting or analysing court judgments. “Legal reasoning” (110 occurrences) and “legal decision-making” are recurring phrases, reflecting the enduring pursuit of modelling how judges or lawyers think. In the domain of specific laws and domains, “intellectual property” appears (57 times, plus 46 for “intellectual property rights”), pointing to AI’s use in patent law or copyright (and also legal questions about AI-generated IP). “Contracts” (31) is another key term – relevant to contract analytics and automation. We also observe “health” and “medical” in some keywords, suggesting overlap with healthcare law (e.g., AI in medical legal issues) and “fintech” or “finance” in others, indicating AI in legal aspects of finance (like fraud detection).

- **Ethical and Societal Issues:** Terms related to ethics and policy are highly frequent, evidencing a strong thread of research on AI governance. Besides the aforementioned “laws and legislation”, the term “ethics” appears 128 times on its own, and in combinations like “ethical technology” (241) which refers to technology ethics, “ethical issues”, “ethical implications”, etc. Notably, “data privacy” (176 occurrences) and “privacy” (137) are among the top keywords. This is expected, as data protection (e.g., GDPR) and privacy-preserving AI are major legal concerns. We also see specific phrases like “privacy by design” and “privacy preserving”, indicating technical approaches to privacy in AI systems. The concept of fairness in AI, while crucial, appears under multiple terms: “fairness” (31), “bias” (16), “algorithmic bias” (6), etc., which individually rank lower but collectively signify a robust focus on algorithmic justice. Indeed, if we combine related terms (fairness, bias, discrimination), they form a substantial thematic group. Another term, “accountability”, appears in the context of AI accountability (though not in top 20, it is noted in thematic analysis). Human-centric terms like “human” (666) and “humans” (435) show up due to indexing (denoting studies involving human subjects or implications for humans). The presence of “female” and “male” in keywords (often alongside “human”) is usually an artifact of demographic indexing in experiments or surveys, rather than a thematic focus, but it does indicate when studies involve human participants (e.g. user studies of AI legal tools).


From these common keywords, we glean that the field’s literature heavily features AI methodologies (machine learning, NLP, etc.) and ties them to legal contexts (law enforcement, judicial decision, IP, privacy, ethics). The dominant themes revolve around using AI to assist or automate legal tasks (decision support, crime analysis, document analysis) and analysing the impacts of AI on law and society (through regulatory and ethical lenses). There is a clear indication that technical research (e.g. on algorithms, systems) goes hand-in-hand with normative research (on law, ethics) – a dual nature reflected in the keywords.

Emerging and Evolving Keywords: Tracking keyword occurrences over time reveals shifts in interest. Early literature (1980s–1990s) was replete with terms like “expert systems”, “knowledge representation”, “legal reasoning”, reflecting the focus on rule-based AI. As we moved into the 2000s, we see increased frequency of terms such as “data mining”, “text mining”, “ontology” (ontologies were a hot topic in the 2000s for knowledge representation of

legal domains). In the 2010s, “machine learning” and “big data” became buzzwords as computational power and data availability grew – their rise is evident in the publication record (e.g., “big data” appears 245 times, mostly in post-2010 papers). Specifically, the mid-2010s onward show an explosion of terms relating to deep learning (which was virtually absent pre-2010 and then surged) and fairness/ethics (many terms in this realm start appearing after 2015). For example, the term “algorithmic fairness” first appears around 2015 and grows thereafter, mirroring real-world concerns about AI bias coming to prominence. Another example is “GDPR” (Europe’s data protection law) which appears as a keyword in multiple recent papers (post-2018) as scholars study its impact on AI development. “Explainability” or “explainable AI” also emerges in the late 2010s as a significant theme (though in keyword lists it may appear under various terms like “interpretability”, “explanation”, etc.). On the legal side, new subtopics have gained traction: for instance, “legal analytics” and “computational law” appear as phrases denoting data-driven analysis of legal information, and “smart contracts” (with the advent of blockchain technology) enters the lexicon in recent years. The term “AI governance” or “AI regulation” becomes more common in the 2020s, reflecting scholarly attention to how societies manage AI. In contrast, some older terms plateau or decline: “expert system” usage has dropped off significantly after the 1990s, and generic terms like “knowledge-based system” are less used in favour of more specific AI technique names. We also see that contextual terms have shifted – e.g., early works often talked about “legal expert systems” in broad terms, whereas newer works use specific context like “AI for contract review” or “AI for policing,” indicating a move from general exploration to targeted applications. Another noteworthy trend is the appearance of terms related to specific regulations or frameworks: aside from GDPR, terms like “algorithmic accountability”, “transparency”, “ethics guidelines” (possibly referencing documents like the EU Ethics Guidelines for Trustworthy AI) are showing up, which wasn’t the case a decade ago. This signals that the research is increasingly engaged with real-world legal frameworks and policy discussions.



Overall, the keyword evolution paints a picture of a field that started with an emphasis on applying AI to legal logic and decision-making, and has expanded to include data-centric techniques and societal implications. The co-occurrence network of keywords further helps to delineate how these terms group into thematic clusters, which we discuss next.




Thematic Clusters (Co-Word Analysis): By analysing which keywords appear together in the same publications, we identified major thematic clusters in the AI-law research network. Each cluster represents a set of topics that are closely related, indicating a subfield or research theme. We found three primary clusters (visualized via a co-word network graph, where node size reflects keyword frequency and edges reflect co-occurrence frequency):

This cluster groups terms related to legal reasoning and decision aids. Keywords like “legal reasoning”, “expert systems”, “knowledge-based system”, “case-based reasoning”, “legal ontology”, and “decision support” are central here. These often co-occur with specific legal

domains (e.g., “contract”, “evidence”) and with AI terms indicating methodology (e.g., “rule-based”, “reasoning”, “inference”). This cluster essentially represents the traditional AI & Law research track that aims to model and support legal decision-making using AI. Papers in this cluster might describe systems that help judges or lawyers by encoding legal knowledge (rules, cases) and providing advice or predictions. The presence of terms like “ontology” and “knowledge representation” point to efforts to formally model legal knowledge, while “case-based reasoning” and “argumentation” relate to how AI can mimic legal argument. This cluster is strongly associated with the earlier decades of AI & law research, but it persists today, evolving with new techniques (for instance, integrating machine learning with symbolic reasoning).

The second cluster centres on the ethical, social, and legal implications of AI itself. Key terms here include “human”, “law”, “fairness”, “privacy”, “regulation”, “laws and legislation”, and “policy”. These tend to co-occur in publications that discuss how to govern AI technologies or examine the impact of AI on society and legal rights. For instance, “privacy” and “data protection” appear in discussions of big data analytics under privacy laws. This cluster has grown significantly in recent years as AI’s societal impact has become a pressing issue. It represents the confluence of legal scholarship and AI – many papers in this cluster come from legal scholars analysing AI, or collaborations between lawyers and technologists.




Currently, all three clusters are actively pursued, often in a complementary manner – e.g., to build a fair legal decision support system, one must merge technical, practical, and ethical knowledge.

Institutional and Country Contributions

AI and law research is a global endeavour, but contributions are not evenly distributed. Here we highlight the leading countries and institutions in terms of research output, and examine collaboration patterns.

Leading Countries:

Country Scientific Production



The production of AI-law research is dominated by a few countries. China and the United States are the two biggest contributors by a wide margin. China accounts for about 3,715 documents (roughly one-third of all publications), and the USA about 3,198 (around 28%). Together, they contribute over 60% of the literature. This imbalance is partly reflective of larger trends in AI research – China and the US have heavily invested in AI R&D in recent years – and these two nations also have large communities of computer scientists and legal scholars engaging with AI. India is the third-largest contributor with 1,421 documents (12.6%), which is notable and

likely driven by the country's growing interest in AI for governance and its strong IT sector. Close behind, several European countries make substantial contributions: the UK (1,353 docs, 12%), Italy (1,072, 9.5%), Germany (912), Spain (593), Netherlands (549), and France (533) are all in the top 10 by output. This indicates that Western Europe, collectively, forms a significant portion of the research, aligning with the presence of long-standing AI & Law research groups (in the UK, Italy, Netherlands, etc.) and the EU's active role in AI regulation (which spurs academic work). Other countries in the top tier include Australia (538) and Canada (450, just outside top 10), showing engagement in the Commonwealth and North America beyond the US.

In terms of influence, the USA leads in total citations (25,735) and has a high average citations per paper (32.4), suggesting that US-affiliated research (often published in high-impact venues or dealing with widely discussed topics like fairness) is frequently cited. China, while prolific, has a lower average citation rate (12.7 per paper), possibly because many Chinese papers are newer or published in local venues, or because of language and accessibility factors. The UK, Canada, and Singapore have relatively high citation averages (22–37) indicating that their contributions often appear in impactful collaborative works or key journals. An interesting data point: Singapore has fewer papers (150 by primary count) but a very high average citation (36.7), likely due to a few highly cited policy papers or collaborations (Singapore has been active in AI governance discourse). Similarly, countries like Finland and Iran show high average citations from small outputs, often reflecting one or two seminal papers (for instance, an Iranian-affiliated author co-wrote a top fairness paper, contributing to Iran's total citations disproportionately).

The international collaboration level is moderate. Overall, about 17% of publications involve authors from multiple countries. We observed that China and the US – the leaders – have mostly domestic collaborations: 89% of Chinese publications are co-authored by Chinese institutions only (with 11% involving international partners). The US has around 14% international collaborations. This suggests that both countries have large enough internal networks to conduct research without needing as much cross-border collaboration, or they publish in forums where domestic co-authorship is common. In contrast, European countries have higher international collaboration rates: for example, the UK and Italy each have 27–28% papers with international co-authors. This reflects the collaborative nature of EU research projects and the smaller size of individual countries necessitating partnerships. It's common to see UK, Italian,

Dutch, German researchers co-authoring papers (often through networks like the European AI & Law community, e.g., the JURIX conference). India has about 13% international papers, indicating most work is domestic but a portion involves collaborations (often with US, UK, or other countries, given many Indian researchers have ties abroad). One consequence of these patterns is that certain highly cited papers are international efforts – for example, a landmark paper on AI ethics might involve authors from the US, UK, and Germany, hence boosting all those countries' citation counts.

Overall, the geography of AI and law research shows a concentration in Asia (China, India), North America (USA, with Canada contributing), and Europe (UK, Italy, etc.), with growing contributions from other regions (e.g., Australia and Singapore in Asia-Pacific). Africa and South America are underrepresented, though there are emerging efforts (South Africa and Brazil have a few contributions, often in AI ethics context, but not yet large in number).

Leading Institutions: The institutional landscape is similarly dominated by a mix of established Western academic centres and a few large Chinese universities. Table 2 shows the top five institutions by the number of publications in AI-law.

Table 2. Top five institutions contributing to AI and Law research.

Rank	Institution	Country	Publications
1.	University of Bologna – A pioneering center for legal informatics (hosts CIRSFD, etc.).	Italy	83
2.	University of Oxford – Noted for its AI research and law-tech initiatives.	UK	80
3.	University of Amsterdam – Home to research on law & AI (e.g., Leibniz Center).	Netherlands	77
4.	University of Liverpool – Known for logic and AI in law (key researchers in argumentation).	UK	76
5.	University of California (combined campuses) – UC Berkeley, UCLA, etc., collectively active in AI policy and tech.	USA	71

These institutions have been at the forefront of AI-law research. The University of Bologna is ranked first with 83 papers; it boasts one of the earliest academic groups in this domain (going back to the 1980s) and has produced extensive research on legal ontologies, normative systems, and more (professors like G. Sartor have been key). Oxford (80 papers) has strengths in AI ethics and law (e.g., its law faculty and computer science department collaborate on AI governance, and the Oxford Internet Institute contributes on AI policy). Amsterdam (77) has the Leibniz Centre for Law, which historically specialised in knowledge-based systems for law, and more recently in legal information retrieval and AI. Liverpool (76) is notable for work on computational models of argument and multi-agent systems applied to law (with T. Bench-Capon's long-term contributions). The University of California entry (71) likely aggregates multiple campuses: UC Berkeley (with its law school and AI labs working on fairness and policy), UCLA (with work on AI and the law, such as autonomous vehicle law), etc., which combined make the UC system a major contributor.

Just outside the top five, we have Tsinghua University in China with 70 papers. Tsinghua, known for its computer science excellence, has in recent years turned attention to AI ethics and law (including an AI Institute for governance). Its presence signals China's academic input – another Chinese institution in the top 10 is Beihang University (52 papers) known for AI and policy research. The list also includes Stanford University (60 papers) and Carnegie Mellon University (59) at ranks 7–8, reflecting strong American research interest: Stanford's CodeX center focuses on legal tech and AI, and CMU has contributions in areas like AI for social good (which includes legal applications). University College London (UCL) (54) and University of Pittsburgh (54) are tied around 9th place. UCL has groups on legal analytics and an AI centre, and Pittsburgh (with Kevin Ashley and colleagues) has been a hub for case-based legal reasoning research. It's also notable to see National University of Singapore (NUS) (52) and University of Cambridge (52) among the upper ranks, indicating significant work on AI and law in Singapore (likely AI governance and fintech law) and Cambridge (which has centres for law, medicine and AI, etc.). An interesting entry is Uttaranchal University (52) from India, which suggests a strong niche effort in India on AI/law (possibly in areas like cybersecurity law; this might also reflect an indexing quirk or a particular prolific group there).

In terms of institutional impact, many of these top institutions are also the ones producing highly cited work. For example, Oxford and Cambridge researchers co-authored influential policy papers, Stanford and CMU produced core AI algorithms cited in legal AI contexts, etc.

The collaboration networks often show these institutions connecting: e.g., there are joint projects between European institutions (Bologna, Amsterdam, Liverpool often link via EU projects), and between US and European (Oxford-Stanford collaborations on AI ethics). The data also show that leading institutions tend to specialise: some are more technical (CMU, Tsinghua, focusing on algorithms), some more on law/policy (Oxford, NUS), and some bridging both (Stanford, Bologna).

The co-authorship network at the institution level reveals clusters often along geographic lines. For instance, European universities form a cluster (with frequent EU collaborations), while many U.S. universities collaborate among themselves or with Canadian/UK partners. Chinese institutions collaborate heavily within China, though we see some links abroad (e.g., Tsinghua with US universities on AI governance studies). International organisations or labs (like IBM Research or Microsoft Research) also contribute but were not listed in the top academic institutions; however, their presence as co-authors (with academia) is part of the network.

In summary, a relatively small number of institutions (primarily in Europe and North America, with a couple in Asia) serve as centres of excellence in AI and Law, producing a large share of the research. These hubs drive innovation and also train many researchers (who then spread globally, further contributing to collaboration networks). Over time, we might expect more institutions – especially in Asia (China, India) – to climb the ranks as interest grows, but currently the historical centres in Europe (Bologna, Amsterdam, etc.) and elite universities in the US/UK remain extremely influential.

Conclusion

This bibliometric analysis has mapped the development of research at the intersection of artificial intelligence and law, revealing a dynamic and rapidly growing field. The analysis shows that AI & law research has grown from a handful of theoretical explorations in the late 20th century into a substantial interdisciplinary domain in the 21st century. The citation landscape is characterised by a core set of influential works – notably, foundational AI papers (like Zadeh's on fuzzy logic) and recent works on fairness and transparency – that have guided the research agenda. These works are widely regarded as the building blocks of current knowledge, evidenced by their high citation counts and frequent appearance in reference lists. At the same time, the long tail of many modestly cited papers indicates a broad base of exploratory studies and niche applications, suggesting a healthy diversity in the field's research

questions.

The keyword and thematic analysis highlight how the focus of AI-law research has broadened over time. Early focus on rule-based expert systems and logic gave way to data-driven analytics and predictive models, and more recently to considerations of ethics, bias, and governance. Terms like “machine learning”, “neural networks”, and “NLP” now coexist with “legal reasoning” and “case-based argumentation”, reflecting an integration of modern AI techniques with traditional legal AI approaches. Furthermore, emerging keywords related to policy (e.g., “AI regulation”, “accountability”, “transparency”) underscore the field’s engagement with real-world legal challenges posed by AI. The thematic clusters identified – legal decision support, law enforcement analytics, and AI ethics/policy – encapsulate the field’s major branches. Importantly, these branches are interrelated and increasingly inform each other: for instance, technical advances in legal analytics raise ethical questions that legal scholars address, and legal requirements (like explainability) drive the development of new technical methods.

Analysing authorship and collaboration patterns revealed that the AI-law research community is globally distributed but with concentrations of expertise. A relatively small number of authors and institutions (many in the US, Europe, and China) have disproportionately driven the publication output and innovation in this field. The network analysis shows several well-connected clusters of researchers – often aligned with geographic or institutional proximity – who collaborate frequently. International collaboration, while present, could be further strengthened (currently 17% of works are inter-country). Greater cross-pollination between regions (for example, more Asia-Europe or Asia-America collaborations) could enhance the field, combining diverse perspectives (e.g., different legal systems’ approaches to AI). The data also highlight the role of interdisciplinary collaboration: many impactful papers are co-authored by computer scientists and legal scholars, a trend that should be encouraged as it bridges the gap between technical feasibility and legal practicality.

In terms of venues, the presence of both computer science conferences and law journals as major publication outlets suggests that researchers are successfully reaching both audiences. However, it also implies that knowledge can be siloed – legal insights might not always trickle into AI conferences and vice versa. Efforts like workshops, joint panels, and special issues can continue to foster dialogue between communities. The bibliometric findings can guide such

efforts by identifying which forums have been most influential and where interdisciplinary exchange is happening (e.g., the ICAIL conference, the AI and Law journal, etc., as key meeting points).

Looking ahead, the bibliometric trends point to several future directions. The surge in AI ethics and policy research will likely continue, especially as governments enact new laws (such as the EU AI Act) – researchers will study the implementation and effects of these regulations and possibly propose frameworks to audit and comply with them. Transparency and explainability of AI systems will remain a central concern in legal contexts, so we can expect more work on technical solutions that satisfy legal standards of explanation. Another growing area is AI applications in government and public administration (sometimes termed “RegTech” or regulatory technology), where AI helps interpret or enforce regulations – this intersects with law and could become a larger theme. The advent of very large AI models (like GPT-4 and beyond) and their ability to perform some legal reasoning tasks (e.g., drafting documents or answering legal questions) will surely spawn research on their accuracy, reliability, and legal ramifications. As suggested by our analysis, the field tends to respond to AI advances; thus, managing and harnessing generative AI in the legal domain could become a prominent research theme in the immediate future.

In conclusion, AI and law have evolved into a vibrant field tackling some of the most pressing questions of how society will integrate advanced AI systems in a lawful and just manner. The bibliometric approach used in this study provided an evidence-based overview of this evolution – charting the growth in publications, spotlighting influential works and contributors, and revealing the shifting thematic priorities. Such an analysis is valuable not only as a retrospective but also as a planning tool: by understanding where the field has been, researchers and policymakers can better see where it should go. The insights suggest a field that is maturing, in that it has established foundational knowledge and is now increasingly impact-driven, seeking to inform real legal practice and policy. The role of AI in legal studies will likely expand, with the boundaries between “AI and law” and mainstream “law” continuing to blur as AI becomes part of standard legal workflows and legal thinking becomes part of AI system design. We anticipate that future bibliometric studies, say a decade from now, will reflect further convergence of these areas – perhaps treating “AI and Law” not as a separate niche, but as an integral component of both the AI research universe and the legal scholarly landscape.