IJHSR
ijhighschoolresearch.org
RESEARCH ARTICLE

AI On Trial: A Holistic Review Of The Promise And Perils Of Artificial Intelligence In The Practice Of Law

Ginevra V. Catenacci

St. Stephen's School, Via Aventina 3, Rome, 00153, Italy; ginevracatenacci08@gmail.com

ABSTRACT: While Artificial Intelligence (AI) holds the potential to dramatically enhance the efficiency and accessibility of legal systems, its integration into the legal field also poses deep structural and ethical challenges. This paper offers a comprehensive synthesis of the debate, bringing together three critical dimensions often examined in isolation: (1) AI's benefits in enhancing legal workflows, outcome prediction, and advanced research; (2) the threats posed by bias, opacity, and unequal access; and (3) the ethical frameworks required to govern its use. By weaving these perspectives together, the paper provides a uniquely holistic lens for understanding not only what AI can do for law but also what it risks taking away if left unchecked. Our analysis argues that any integration of AI in law must be guided by robust legal frameworks capable of embedding principles of justice, explainability, and human oversight into AI systems from the beginning. This review thus serves as both a caution and a roadmap: the goal is not to resist technological change but to ensure that it serves the values at the heart of democratic legal systems rather than undermining them, offering space for responsible innovation rooted in legal integrity and justice.

KEYWORDS: Robotics and Intelligent Machines, Machine Learning, Legal Ethics, Artificial Intelligence, Algorithmic Bias.

Introduction

Imagine a courtroom where a machine identifies the most relevant case precedents, predicts the likely verdict, and proposes the optimal litigation strategy, all within seconds. Though this may sound like a scene from a fiction movie, it reflects the accelerating reality of artificial intelligence (AI) in the legal field.

The intersection of AI and law is not new; it has a rich history that set the foundation for today's rapid advancements. Its roots extend back over half a century, from the 1970s research on legal reasoning and logic, through the 1980s development of expert systems and case-based reasoning, to even earlier inquiries by legal scholars and mathematicians like Gottfried Leibniz in the 1600s, who envisioned applying mathematical thinking to law.[1] With AI tools already automating contract reviews, aiding judicial risk assessment, and transforming litigation research, law has become a central domain of AI disruption.

Artificial Intelligence has emerged as a transformative force in the legal profession since the mid-2000s, particularly in the wake of the so-called Industry 4.0 era, which has ushered in data-driven automation across industries. As a consequence, the legal field, once characterized by time-intensive manual processes, now stands on the cusp of a major shift.[2] The most immediate and widely discussed impact of AI is its transformative potential in improving efficiency and accuracy in legal practice. Multiple studies detail how machine learning systems can streamline outcome prediction, contract analysis, and litigation strategy, opening the door to faster decision-making and cost-effective legal services.

AI-driven platforms like ROSS Intelligence, Blue J Legal, and eBrevia can process vast legal datasets at unprecedented speeds, dramatically reducing the time and cost of legal research and document analysis. For instance, JPMorgan's COIN program is able to process and extrapolate specific details from over 12,000 contracts in mere seconds, saving the firm more than 36,000 billable hours annually.[3] Predictive analytics allow lawyers to assess litigation risks more objectively, while Natural Language Processing (NLP) enables the real-time parsing of legal documents, transforming the way in which legal arguments are crafted.[3] These capabilities are not just acting as technical upgrades, but are reframing the legal landscape by enabling a more strategic and data-informed approach to law.

However, while AI technologies have already demonstrated their value in streamlining legal workflows and augmenting decision-making processes among judges and scholars alike, they also raise critical legal and ethical concerns about bias, transparency, data privacy, and due process. As pointed out in "Artificial Intelligence and Law: An Overview" by Harry Surden, current machine-learning models often operate by themselves, so that the rationale behind a prediction or decision is obscure even to its creators.[1] Thus, when courts begin to rely on algorithmic risk assessments, there is a danger that these systems may encode and reproduce the very biases present in their training data, potentially undermining the fairness of legal proceedings.

Further, the growing deference to AI-generated outputs threatens to erode the authority of human legal professionals. Legal scholars such as Rosengrün (2022) argue that AI is increasingly acting as the de facto regulator of legislation and of entire societies, shaping legal norms and decisions without democratic oversight.[4] This shift raises the question of whether AI's efficiency gains are coming at the expense of some of law's most fundamental values: accountability, interpretability, and human judgment. Additionally, legal decisions are gradually being abstracted into coded logic, making it more difficult for

© 2026 Terra Science and Education

DOI: 10.36838/IJHSR811.58

individuals to contest, or even fully understand, the processes that affect their rights.

Adding to these concerns is the issue of unequal accessibility concerning a digital divide that could exacerbate existing inequalities in access to legal representation and justice. Moreover, over-reliance on these tools may contribute to the demeaning of legal professionals, diminishing the emphasis on human reasoning and critical judgment that have long defined the legal profession.

Similarly, though these systems would require less expertise and work, without adequate oversight, AI may also embed systemic biases into legal processes, reproducing social inequalities under the guise of objectivity. In addition, issues like AI hallucinations, which are very common and not transparently noted by the system when they happen, increase the risks of untrue and unobjective facts being presented to the court.

Overall, these developments point to an urgent need for a regulatory and ethical framework that can accommodate innovation while safeguarding justice. As noted by Getman *et al.*, responsible integration of AI into the legal system requires more than technological refinement; it demands interdisciplinary engagement and normative reflection.[5]

This paper undertakes a critical analysis of the impact and ethical use of AI within the U.S. legal field, drawing examples from other jurisdictions where relevant, with peculiar attention to legal research, litigation, and decision-making. It explores the promises and perils of this shift, analyzing how AI is reshaping legal practices, where it risks undermining core legal values, and the ethical guardrails around its deployment. By drawing on both academic literature and real-world applications, the analysis aims to trace the dual-edged nature of AI's influence on legal practice, an intersection where the stakes are nothing less than the future of justice itself.

Examining AI's role in the legal profession is extremely important nowadays. Developments in AI technologies mean that algorithmic decision-making is no longer hypothetical, but it is actively shaping access to justice, the fairness of proceedings, and public trust. Without critical evaluation, the promise of AI could come at the expense of accountability and transparency; thus, understanding both its potential and perils is essential to ensuring that technological progress strengthens the core values of the legal system. Ultimately, the paper seeks to outline the existing relationship between artificial intelligence and legal practice and to present the path forward as one that embraces the opportunities of AI without losing sight of justice, fairness, and integrity.

To structure this inquiry, the exploration is divided into three thematic sections: the benefits of AI in legal practice, the threats it poses to legal institutions and professional judgment, and the principles guiding the ethical uses of AI within law. To unpack these questions, this paper adopts a critical review methodology, drawing exclusively on a wide body of scholarly work to provide a balanced and evidence-driven understanding of the intersection of AI and law.

Section I: The Benefits of AI in Legal Practice

Artificial Intelligence has already proven to be a tool that is reshaping how law is researched and argued by streamlining workflows, enhancing predictive accuracy, and expanding the scope of legal research. Consequently, these technologies are enabling law practitioners to operate with unparalleled efficiency and precision. This section examines how AI augments legal practice and argues that these benefits could wholly redefine the role of legal professionals in the 21st century. Specifically, we address the benefits of AI in: 1) enhancing efficiency in legal workflows; 2) improving the prediction of legal outcomes; and 3) enabling more advanced, data-rich legal research and problem solving. Together, these advancements point to an evolving professional ecosystem where human expertise and machine intelligence operate in tandem.

A. Enhancing efficiency in legal workflows:

The legal profession, once defined by hours of manual document review, citation checks, and precedent hunting, is undergoing a transformation powered by technological software that is automating routine tasks and dramatically reducing the time required for traditionally labor-intensive legal tasks. According to Assistant Professor Shanin Kabir, the software of Kira Systems, founded by Noah Waisberg, can speed up the process by up to 40% for first-time users and up to 90% for experienced users.[3] In addition, eBrevia, founded by Ned Gannon and Adam Ngueyn, claims to be able to review over 50 documents in under a minute with accuracy rates surpassing manual review by approximately 10%.[3] For example, Lawgeex claims to reduce contract review time by up to 80% and costs by 90%. These kinds of operational improvements free up legal professionals to focus on strategic planning rather than repetitive review tasks and save up time and resources for lawyers, clients, and law firms. According to Harry Surden, technology-assisted review and predictive coding now allow lawyers to efficiently filter through millions of documents, a process that would previously have required teams of paralegals working for weeks.[1] Consequently, these systems are not merely speeding up work but are redefining what is feasible within litigation timelines and transforming the due diligence process by raising efficiency, reducing errors, and automating various tasks.

Machine learning, which involves the development of computer systems that are able to learn and adapt without explicit instructions, is currently the most impactful approach to Artificial Intelligence. AI tools operate through machine learning systems, whose rise has been fueled by a massive increase in the amount of data on the Internet, to "learn" from user feedback and patterns and incorporate the findings into their results, and, as more data gets added over time, improve the rates of accuracy and precision. As a consequence, the results of AI research are often similar to those a human would've reached, over a much longer period of time, without ever even reaching human-level cognition. Algorithms like this can be used to assist lawyers by extrapolating the relevant information within hundred-page-long documents, identifying the core legal

principles at stake, and predicting case outcomes. In general, AI works extremely well for tasks that have clear, definite, and unambiguous patterns, such as Surden mentions, chess, fraud detection, and tumor identification. [1]

Moreover, AI's ability to assist in the drafting of legal documents is another major area of impact. Automating legal drafting through pattern recognition and data extraction enables rapid production of contracts and pleadings that comply with jurisdiction-specific requirements.[2] While human oversight remains essential, AI dramatically reduces the time needed to produce reliable, legally sound drafts.

B. Prediction of Legal Outcomes:

Beyond efficiency gains, AI's ability to analyze large datasets also enables powerful predictive tools that help lawyers forecast legal outcomes and make more data-driven strategic decisions, especially in the early stages of litigation. In their paper, Alarie *et al.* emphasize that while traditional human judgment often relies on intuition and experience, such "quick and dirty" predictions are frequently unreliable and prone to bias.[6] Drawing from decades of studies across various fields, including parole decisions, medicine, and insurance, the authors demonstrate that formal, algorithmic methods consistently match or surpass human judgment in predictive accuracy.[6] This foundational insight applies equally to the legal field, where data-driven analytics can provide more objective and reliable forecasts than lawyers working unaided by technology. One of the central challenges in applying predictive analytics to law has historically been the availability and usability of data. Nowadays, however, given that most documents are electronically published, research and discovery are way more accessible to people who don't have the resources normally required to initiate, continue, or respond to a lawsuit, including time and money. John Markoff, in fact, wrote in the New York Times that e-discovery software can analyze documents in a fraction of the time for a fraction of the cost. For instance, Blackstone Discovery of Palo Alto, California, analyzed 1.5 million documents for less than $100,000.[7] For example, Alarie *et al.* say that more than one million Canadian court and tribunal decisions are publicly accessible through platforms like CanLII.[6] Additionally, regulators such as the Canada Revenue Agency and securities commissions publish vast numbers of rulings and decisions, creating a rich dataset that far exceeds the capacity of any individual lawyer to fully absorb.[6] Thus, by processing this diverse data, AI algorithms can detect patterns and generate predictions about legal outcomes that are both faster and more accurate than human analysis alone.

Moreover, the advantages of AI-driven prediction extend beyond accuracy as they operate at a remarkable speed, delivering predictions within seconds or minutes, in stark contrast to the often lengthy deliberation human lawyers require. Furthermore, unlike humans, algorithms do not suffer from fatigue, emotional fluctuations, or inconsistent judgment, factors that could compromise decision quality. This consistency is particularly valuable in legal contexts where impartiality and fairness are paramount. In addition, Alarie and his colleague mention how machine learning algorithms, including neural networks, pioneered by experts like Professor Geoff Hinton, learned from data to identify hidden connections between variables and outcomes without relying on restrictive assumptions.[6] This is extremely useful as these algorithms automatically adapt to complex patterns, improving the precision of legal predictions.[6] The authors also caution about the importance of data quality in machine learning applications. The famous "garbage in, garbage out" principle applies: if the input data is flawed or poorly curated, the predictive model's outputs will also be unreliable. Nonetheless, with well-curated data and robust model selection, machine learning offers significant improvements over traditional statistical techniques.[6]

Looking toward the future, we can foresee an exponential increase in computing power and cognitive computing capabilities, such as natural language processing, which will enable AI systems to process data like court opinions and rulings more effectively. By analyzing millions of data points, these algorithms reduce inconsistencies and errors, helping to avoid unnecessary detentions and associated social and financial costs. Such uses of AI illustrate the broader impact of predictive tools in promoting fairness and efficiency in the justice system. Furthermore, by providing clearer predictions about likely outcomes, these technologies increase the probability of settlement and reduce litigation, focusing court resources on genuinely ambiguous cases where legal development is most valuable.[6] As machine learning tools become more sophisticated and widespread, lawyers may face legal obligations to incorporate these technologies into their due diligence processes, ensuring that advice and litigation strategies are informed by the best available data-driven predictions.

Blue J Legal:

Many aspects of the legal landscape, in its complexities, are often described as grey areas, where no single fact is dispositive, and outcomes hinge on how courts weigh an interrelated web of considerations. For lawyers, this means navigating an enormous volume of precedent cases, each with subtle factual differences that can influence the final determination. The challenge of thoroughly researching and interpreting hundreds of relevant decisions makes providing confident legal advice difficult and resource-intensive. In this context, a compelling example presented by Alarie *et al.* is Blue J Legal, a Toronto-based startup that harnesses machine learning to predict legal outcomes in Canadian courts, particularly in tax law, by analyzing thousands of past cases and identifying the patterns that underlie judicial decisions.[6] Rather than relying on simple checklists or formulas, their algorithms evaluate multiple variables simultaneously, learning from the complex interactions among factors that human researchers might overlook. This technology significantly improves the speed and accuracy of predicting case outcomes, also quantifying the likelihood of the predicted outcome, giving lawyers a powerful tool to cut through ambiguity and focus on the most relevant precedents. This capacity to model uncertainty and reveal the grey zones within the law is particularly valuable in disputes where facts may be contested or open to interpretation.

DOI: 10.36838/IJHSR811.58

Beyond enhancing individual case assessments, Blue J Legal's system also incorporates user feedback to refine its models continuously. By aggregating diverse professional judgments, the platform mitigates individual biases and provides a more balanced and reliable analysis over time. This decentralized learning process ensures that the tool evolves alongside the law itself, adapting to new cases and shifting legal standards. Ultimately, such innovation promises to reduce costly litigation by helping parties and their lawyers reach clearer, more informed agreements early in the process and empowering more confident and efficient legal decision-making.

C. Advanced Legal Research and Problem Solving:

AI's impact on legal research is similarly transformative. Historically, legal research was an arduous, manual endeavor, requiring hours of combing through case reporters, statutes, and secondary sources. This process, though thorough, was time-consuming, susceptible to incomplete or outdated information, and inherently limited in its capacity to explore the full breadth of available legal data. Today, AI engines not only save time but also often uncover relevant materials that traditional searches might miss, enhancing the depth and quality of legal arguments.[11] Modern AI systems bring together natural language processing (NLP), machine learning (ML), and predictive analytics to conduct complex legal research in ways that mimic, yet significantly amplify, the cognitive processes of experienced attorneys. NLP enables semantic search, allowing queries to be posed in natural language rather than through rigid keywords, while ML models learn from patterns in data to refine their results over time. Systems such as ROSS Intelligence, Lexis, and CaseMine exemplify this shift, enabling lawyers to ask sophisticated legal questions and receive highly targeted results.[11] In practice, this means lawyers can uncover precedents and legal arguments that traditional search approaches might overlook, increasing both the depth and quality of their reasoning. For instance, BakerHostetler's use of ROSS Intelligence in its bankruptcy department, navigating 27 terabytes of data and delivering rapid and precise answers, illustrates the transformative potential of such tools.[3]

The expanding range of commercial AI tools demonstrates the width of these applications and underscores how automation enhances strategic decision-making across legal practice. LEVERTON, an offshoot of the German Institute for Artificial Intelligence, uses AI-powered data extraction to process contracts in over 20 languages, extracting key financial and operational details such as rent, maintenance obligations, and expiration dates, and presenting them in a structured spreadsheet format to expedite real estate due diligence.[3] This showcases AI's ability to distill complexity into actionable understanding.

Furthermore, ThoughtRiver's Fathom Contextual Interpretation Engine, developed with Cambridge University expertise, similarly demonstrates AI's capacity to interpret nuance: by automatically flagging risky clauses and generating plain-language summaries, it enables faster and more consistent risk assessments.[3] Legal Robot and Casetext's CARA extend this capability further; the former converts legal text into numerical representations for machine analysis, automatically flagging potential issues for human review, while the latter anticipates opposing arguments and surfaces relevant precedents, aiding lawyers in proactive litigation strategy. Each of these tools exemplifies AI's ability to integrate vast, disparate legal data into coherent insights. In addition, while early systems were hampered by cost and computational limitations, modern AI systems incorporate heuristic methods that allow them to improve with use and to recognize complex, non-linear relationships in data.[8] Together, these examples reveal that AI's true advantage lies not merely in efficiency, but in augmenting the lawyer's cognitive toolkit, transforming legal research into a predictive, interpretive, and ultimately more strategic form of problem-solving.

AI's strengths also lie in enhanced information retrieval and the mechanization of juridical activities. It can quickly locate the relevant legal texts, precedents, and studies across multiple jurisdictions and levels of law while simultaneously integrating administrative regulations, judicial decisions, and academic commentary.[10] This capacity allows AI to detect contradictions between sources, reveal ambiguities in legal provisions, and even propose innovative legal solutions that advance the doctrine of law. By bridging different branches of law, AI supports a more holistic, cross-disciplinary approach to legal problem-solving, enabling insights that human researchers might overlook. Moreover, centralized legal information systems, powered by AI, could provide courts, legislatures, and administrative agencies with immediate, consistent access to the same high-quality legal data, reducing duplication of effort and fostering a more coherent application of law across jurisdictions.

In sum, the integration of AI into legal practice has already transformed the profession's capacity to work with unprecedented speed, depth, and accuracy. Taken together, these developments show that AI is not merely an auxiliary tool but a transformative force capable of redefining the very mechanics of legal work by streamlining processes and uncovering insights and connections that might otherwise remain hidden. This evolution marks a decisive shift from merely retrieving information to generating richer, more nuanced legal reasoning. Yet, as these tools become more sophisticated and deeply embedded in legal workflows, they inevitably surface new questions about their influence on the profession's methods and values that demand careful consideration in the broader conversation about AI's role in law.

■ Section II: Threats of AI on the Legal Practise

While the advantages of artificial intelligence in legal practice are compelling, its rapid adoption also introduces profound risks that threaten the integrity of the profession. These risks are not speculative but already visible in the ways AI is reshaping legal workflows, reasoning, and access to justice. Scholars warn that unchecked reliance on algorithmic systems could erode professional expertise, distort outcomes through biased or flawed data, exacerbate inequalities in access to legal resources, and undermine transparency and accountability at the core of the rule of law. These concerns reveal that AI's influ-

ence extends beyond efficiency gains into the very structure of legal reasoning and practice. In order to fully understand the stakes, it is necessary to examine the specific domains in which AI poses threats.

A. Erosion of Legal Expertise and Over-Reliance on AI:

One of the most pressing threats posed by AI in legal practices is the risk of eroding professional expertise through over-reliance on automated systems. The legal profession has long been structured around a system of apprenticeship, where junior associates traditionally develop analytical ability, familiarity with legal sources, and judgment through labor-intensive processes such as reviewing documents, cross-checking citations, and drafting pleadings under the supervision of senior colleagues. This "apprenticeship model," though repetitive and time-consuming, formed the bedrock of professional development, allowing young lawyers to cultivate the analytical habits, interpretive skills, and contextual knowledge that underpinned independent judgment. Yet, as Harry Surden observes, the rapid adoption of machine learning tools such as predictive coding, e-discovery platforms, and automated contract analysis is transforming the very work that once trained generations of lawyers.[12] While this enhances efficiency, these technologies, by processing millions of documents at speeds far beyond human capacity, are gradually supplanting the entry-level tasks that were once essential for developing professional competence.

This concern reflects a broader "deskilling" effect that legal scholars have begun to document in which lawyers become passive consumers of algorithmic outputs rather than active interpreters of legal texts. This threatens to demolish the rigor of the legal system, where the value of arguments depends not only on citing precedent but also on deploying creativity, contextual nuance, and rhetorical judgment. If AI systems are allowed to dominate core tasks, lawyers may gradually defer to machine-generated results without interrogating their accuracy or context, thus producing a generation of lawyers whose interpretive instincts are underdeveloped. AI tools such as Lawgeex or eBrevia, capable of reviewing contracts with greater accuracy and speed than human paralegals, might save clients money and streamline transactions, but they also eliminate countless hours of hands-on experience that previously helped junior associates learn how to spot subtle issues, weigh alternative interpretations, and anticipate opposite arguments. Under this lens, a legal culture shaped by over-reliance on AI risks elevating efficiency above interpretive quality, eroding the human capacities that give the law its adaptability.

Historical perspectives reinforce the gravity of this problem. Early computer scientists working on automated legal reasoning already anticipated that computers would eventually outperform humans in the speed and thoroughness of data retrieval, but they also implicitly assumed that lawyers would always remain central to the discipline.[6] If lawyers instead abdicate these interpretive responsibilities to machines, the result may be not just efficiency gains but also a narrowing of legal imagination. Additionally, humans often cannot fully articulate the reasoning rules they use when solving complex problems; thus, this inability to "teach" the computer everything makes the role of human oversight irreplaceable. Over-reliance on AI risks collapsing this delicate balance, creating a profession that prizes output speed but sacrifices intellectual depth.

This danger becomes most visible in contexts such as predictive coding and automated drafting. For instance, predictive coding software, widely used in e-discovery, can filter out millions of irrelevant documents with minimal human input, leaving lawyers with a smaller subset of potentially relevant materials.[12] While this undoubtedly increases efficiency and saves resources, it also strips lawyers of the experiential learning that comes from sifting through large volumes of material, spotting unexpected connections, and developing intuition about patterns of evidence. Similarly, AI-assisted drafting tools such as those offered by Kira Systems or ThoughtRiver use pattern recognition and clause libraries to generate contracts or risk assessments that once required hours of human labor. These outputs may be accurate within established parameters, but they deprive junior lawyers of the opportunity to practice drafting from scratch. In the short term, this saves time and reduces costs, but in the long term, it risks depriving lawyers of the opportunity to cultivate the judgment necessary to adapt forms creatively to new contexts.

Over-reliance also invites complacency at higher levels of the profession. Even experienced lawyers may be tempted to defer to machine-generated outputs, especially when these are presented with the confidence of quantitative scores or forecasts that project an aura of objectivity. Lawyers, accustomed to reasoning in qualitative, argumentative terms, may find themselves increasingly deferring to these apparently "scientific" measures. The danger is not that AI will replace judgment outright, but that it will subtly reorient decision-making around what the system identifies as most "probable" or "relevant," sidelining considerations that fall outside its parameters. This is particularly troubling in "grey area" cases, where outcomes hinge on how judges balance competing principles. If lawyers grow accustomed to leaning on these outputs without supplementing them with human interpretive skills, the profession risks narrowing its engagement with ambiguity and uncertainty. In addition, by reducing the need for armies of junior associates to conduct research and document review, AI may risk reducing job opportunities for many workers as large firms may shrink their base, relying on a smaller cadre of associates supported by powerful technological platforms. Thus, if the interpretive skills of human workers are allowed to erode under the weight of automation, the law may become faster and cheaper but also narrower, less adaptable, and less humane.

B. Bias, Error, and the Limits of Data-Driven Reasoning:

If efficiency and prediction are the headline advantages of AI in legal practice, the flip side is that these very capacities depend heavily on the quality, neutrality, and completeness of data. Unlike human lawyers, who can contextualize facts, weigh competing values, and recognize when precedent should not be followed, AI systems reason primarily by detecting patterns in historical datasets. This reliance introduces vulnerabilities that can produce errors, perpetuate biases, and distort outcomes in

DOI: 10.36838/IJHSR811.58

ways that threaten the legitimacy of legal decision-making. A central principle in machine learning is that outputs are only as good as their inputs (the classic "garbage in, garbage out" problem that was already mentioned).[9] If a dataset is mistaken and reflects historical inequalities, prejudiced judgments, or incomplete records, an algorithm trained on that data will reproduce and amplify those biases. Yet, when presented as "objective" or "data-driven", such outputs risk reinforcing discriminatory practices under the guise of neutrality. In law, where decisions carry immense social consequences, this is not a theoretical issue but a tangible danger.

Beyond bias, the sheer complexity of legal reasoning exposes the limits of AI's pattern-based logic. While machine learning excels in domains with clear, unambiguous rules, like chess, the law is replete with grey zones. Many disputes hinge not on straightforward fact-matching but on competing interpretations, normative judgments, and context-specific equities. For instance, two cases may share factual similarities but diverge in outcome because a court interprets a principle differently, or because broader social or constitutional values are at stake. An AI trained on surface-level textual patterns may miss these subtleties, producing results that are technically accurate in a statistical sense but shallow jurisprudentially. This limitation reveals the risk of "automation bias," in which lawyers or judges place undue confidence in outputs that seem precise but fail to capture the law's depth.[12]

Errors, too, manifest in ways unique to AI systems. Unlike human lawyers, who are constrained by rules of logic, evidence, and professional responsibility, AI models, particularly large language models, are prone to "hallucinations." This phenomenon occurs when an AI system generates information that is false but presented with the same stylistic confidence as genuine knowledge. In legal research, this can mean fabricating case law, misquoting statutes, or citing precedents that do not exist. In 2023, for example, a New York attorney was sanctioned after ChatGPT, an OpenAI algorithm, generated a legal brief citing six entirely fictitious cases.[12] Even when AI systems retrieve real cases, they may misstate holdings, omit key limiting facts, or misinterpret their doctrinal weight, producing outputs no competent human researcher would accept. The danger lies not merely in occasional errors but in the difficulty of detecting them, because, unless outputs are meticulously verified by human lawyers, hallucinations can pass unnoticed, corrupting arguments and misleading courts. For instance, the case of Mata v. Avianca brought national attention to the ethical pitfalls of using AI. In this instance, two attorneys submitted AI-generated materials in a brief, which' included six fabricated cases, resulting in a sanction under Federal Rule of Civil Procedure Rule 11.[13,14] This underscores how AI, at least at the still underdeveloped stage it is at currently, cannot replace the lawyer's fundamental duty of diligence and independent verification.

Compounding these risks is the opacity of error detections, as traditional legal research, however tedious, offers transparent lines of accountability: a lawyer cites a case, and opposing counsel or the court can verify it. With AI-driven research, by contrast, errors may be buried within opaque "black box" models that provide no clear explanation of how results were generated. Commentators stress that the lack of transparency and explainability undermines accountability, as lawyers cannot easily reconstruct the algorithm's reasoning or test its assumptions. This not only heightens the risk of unintentional reliance on flawed results but also raises difficult questions of professional responsibility, since any errors introduced by AI ultimately remain the lawyer's liability. In addition, the limits of data-driven reasoning are further highlighted by jurisdictional and linguistic complexities. AI models, trained primarily on U.S. or English-language legal data, may perform poorly in other jurisdictions, failing to capture civil law tradition or multilingual statutes.[10] Errors introduced by such blind spots could disadvantage lawyers and litigants in non-dominant legal systems, widening global inequities in access to reliable AI tools.

What makes these problems particularly troubling is the perceived authority of AI outputs because AI tools are marketed as faster and smarter than any individual. Ultimately, law is not simply the retrieval of precedent but the balancing of competing values, the weighing of fairness alongside consistency, and the capacity to innovate beyond past data. By contrast, AI's outputs are tethered to historical patterns and probabilities. This mismatch means that errors, hallucinations, and biases are not accidental flaws but structural byproducts of current AI design. Until systems can provide verifiable, context-sensitive reasoning with transparency and accountability, they will remain valuable aids, but ones whose outputs demand rigorous human oversight.

C. Bias, Inequality, and Access to Justice:

The integration of artificial intelligence into legal systems also has profound implications for fairness, equality, and accessibility. At its core, the rule of law is premised on legal certainty, equality before the law, and equal access to justice, yet AI often threatens these very principles by replicating existing biases, amplifying socioeconomic disparities, and transforming law from a system of human accountability into one governed by opaque code. Scholars warn that this shift risks moving society from a rule of law to a "rule of code," where algorithms and the corporations controlling them exercise regulatory authority without transparency or democratic oversight.[4] This transformation has practical consequences across legal practice, from automated decision-making in welfare systems to predictive policing and e-discovery, and these consequences disproportionately affect disadvantaged groups.

One of the most striking examples of algorithmic bias undermining justice comes from the Netherlands, where the SyRI fraud detection system flagged hundreds of families for welfare fraud. The algorithm disproportionately tagged lower-income neighborhoods and migrant communities, compelling many families to repay benefits they had rightfully received. Ultimately, a Dutch court struck down the system for violating privacy rights and enabling socioeconomic discrimination, underscoring the dangers of entrusting opaque algorithms with high-stakes legal determinations.[9] This example demonstrates how AI, far from correcting human error, often exacerbates ex-

isting inequalities by embedding them into technical systems presented as objective and neutral.

The reliance of AI tools on historical data also undermines the impartiality of legal decision-making. Machine learning models learn from precedent, but when that precedent reflects discriminatory practices, the models inevitably reproduce those injustices. This reliance on "the world as it was" risks turning law into a backward-looking mechanism that perpetuates entrenched inequities rather than enabling progress toward justice.[12] Far from democratizing access, automation can exacerbate exclusion by making high-quality legal services increasingly scarce and expensive while proliferating lower-quality automated alternatives for those least able to contest their shortcomings. Moreover, AI also undermines access to justice by destabilizing employment and reshaping the legal profession in ways that affect service provision, illustrating the tension between efficiency and equity. This could have serious consequences for accessibility, as fewer trained lawyers may limit the diversity and reach of legal services available to ordinary citizens.

Data privacy further complicates the picture as AI systems rely on massive datasets to function effectively, but the collection and analysis of sensitive legal information create significant risks of misuse, surveillance, and breaches of confidentiality. Legal data, including client files, case histories, and personal records, is particularly sensitive and must be handled with robust safeguards. In her paper, Dr. Madaoui Nadjia emphasizes the importance of encryption, access controls, and compliance with data protection laws such as the GDPR to prevent unauthorized access or misuse of legal information.[9] Yet the sheer volume of data required for machine learning increases the likelihood of vulnerabilities, while the opacity of many AI systems makes accountability for breaches difficult to enforce.[9] In this way, the integration of AI into legal practice can weaken client trust and undermine one of the cornerstones of the profession: confidentiality.

The broader societal consequences of AI in law highlight the risk of democratic erosion. When code begins to function as law, regulation shifts away from transparent, contestable processes toward technical architectures controlled by private corporations.[4] Search engines determine which firms and cases are visible, predictive systems shape outcomes in dispute resolution, and contract enforcement becomes automated. These developments place critical legal functions outside of public oversight, empowering corporate actors to influence legal processes in ways that may prioritize efficiency or profit over fairness and equality. Even regulatory initiatives such as the EU Artificial Intelligence Act often risk entrenching corporate control rather than limiting it, thereby legitimizing the very systems that threaten to undermine access to justice.[4] Taken together, these developments reveal that the risks of bias, inequality, and limited access to justice are not incidental side effects of AI adoption but structural consequences of its integration into law. As a result, ensuring that legal systems remain equitable and accessible will require sustained vigilance, stronger safeguards, and recognition that technological progress is not synonymous with justice.

Taken together, these concerns demonstrate that while AI offers unprecedented speed and efficiency, its unchecked integration into legal practice threatens to erode the very foundations of the professions. The keyword being "unchecked," over-reliance on automated systems risks hollowing out legal expertise, flawed or biased data undermines the legitimacy of outcomes, errors and hallucinations reveal the limits of probabilistic reasoning, and opaque, unevenly distributed technologies jeopardize transparency, fairness, and equal access to justice. What emerges is not simply a technical challenge but a structural one: the profession's ability to preserve its standards of judgment, accountability, and equality before the law is at stake. These risks underscore the urgency of ensuring that AI remains a tool subject to law and ethical oversight, rather than an obscure authority that reshapes legal practice on its own terms.

Section III: Ethics

As previously seen, artificial intelligence promises several transformative benefits for the legal profession, yet its integration raises profound ethical challenges regarding fairness, transparency, privacy, accountability, and the preservation of human judgment. As scholars have frequently emphasized, AI is not merely a neutral tool but one that can either reinforce existing inequities or help expose them. Our position on the matter is clear: AI is inevitable and potentially beneficial, but its deployment must be bounded by human oversight and guided by ethical principles.

A. Core ethical issues:

Going back to section 2, bias and fairness remain some of the most significant risks in applying AI to legal contexts. These systems risk embedding systemic injustice into legal decision-making; thus, researchers underscore that addressing such algorithmic biases requires rigorous auditing and the use of diverse, representative datasets. Indeed, as Surden notes, one of the greatest disadvantages of AI in the legal system lies in its data dependency and bias. Because AI learns patterns from historical data that may already reflect social or institutional inequalities, its outputs risk reproducing and amplifying these same biases in ways that are difficult to detect or correct, creating a dangerous illusion of objectivity.[15] Without proactive measures, algorithms may keep perpetuating socioeconomic, racial, or gender-based inequities. At the same time, AI can serve as a diagnostic tool to uncover discriminatory human patterns; thus, bias is both a threat and an opportunity in that AI can exacerbate injustice, but it can also reveal inequities that warrant reform.

Moreover, many AI systems operate as "black boxes" and lack explainability behind their decisions, which raises several concerns over the transparency of their actions and decisions. Scholars warn that explanations like " the statistical model determined this outcome" are insufficient to prove legitimacy, as legal systems demand meaningful, interpretable reasoning that respects the right to be heard. Accountability is similarly at stake. Nadjia stresses the need for comprehensive guidelines

DOI: 10.36838/IJHSR811.58

and codes of conduct for AI developers and legal professionals to ensure responsibility does not dissipate into technical systems.[9] When AI contributes to an unjust outcome, the question of whether responsibility lies with the developer, judge, or lawyer cannot remain unanswered; thus, clear frameworks are required to delineate responsibility and prevent the deflection of blame to the algorithm.[5]

In the legal field, particularly, privacy and data protection concerns are magnified because of the vast amount of sensitive documents and case records that are at stake. Getman and his coworkers highlight the importance of encryption, access controls, and role-based data permissions to preserve confidentiality.[5] As a result, compliance with data-protection laws such as GDPR is essential, requiring informed consent and transparency in processing.

B. Our Recommendations for Ethical AI Use:

Drawing on these concerns and the literature, we propose the following principles for responsible integration of AI into the legal system:

1. Human Oversight is Essential:

AI should augment, not replace, human judgment. Judges, lawyers, and policymakers must remain accountable for outcomes, with AI outputs treated as advisory rather than determinative. Oversight mechanisms could include mandatory human review of any AI-generated recommendation, periodic audits of AI-assisted decisions, and possibly explicit documentation of how AI influenced the decision-making process. This would ensure that professional responsibility remains central, mitigating the risk that overreliance on AI diminishes human accountability.

2. Mandated transparency and explainability:

AI systems must provide both technical transparency and user-facing explanations. Technical transparency would involve audit trails, documentation of model logic, and logging of inputs/outputs. User-facing transparency means that litigants, lawyers, and judges can understand how the AI reached its conclusions in plain language. This would address the "black box" problem, preserve trust, and ensure decisions are interpretable and contestable.

3. Rigorous bias and equity safeguards:

AI should undergo systematic bias testing before deployment and prioritize fairness to reduce the risk of perpetuating historical injustices. This involves auditing datasets for representativeness, employing fairness metrics, and running simulations to detect disparate impacts on different demographic groups. Moreover, AI can be used proactively to identify structural inequities in existing human practices. Implementation strategies include third-party bias audits, mandatory reporting of bias mitigation efforts, and continuous monitoring for emerging biases over time.

4. Strengthened Data Privacy and Security Protocols:

Legal AI systems must adhere to robust privacy standards, including encryption, access controls, anonymization, and compliance with regulations like GDPR. Additional measures could involve: role-based access to sensitive information (ensuring only necessary personnel can access specific data); secure deletion of data once the purpose of processing is fulfilled; regular penetration testing and monitoring for unauthorized access. These safeguards protect client confidentiality, uphold attorney-client privilege, and preserve trust in the legal system.

5. Clear Accountability frameworks:

This means that responsibility must be explicitly assigned at all stages of AI deployment. Judges remain accountable for rulings, lawyers for their strategic decisions, and developers for system design and implementation. Best practices include formal accountability protocols, ethical codes for AI developers, and legal requirements for documenting AI's role in decision-making. Such frameworks prevent diffuse responsibility and ensure ethical and legal recourse when AI-assisted decisions cause harm.

6. Equitable access to AI tools:

To prevent a two-tiered system, AI resources should be extended beyond elite law firms to public defenders, legal aid organizations, and courts serving marginalized populations. Mechanisms could include government-funded AI tools, shared legal AI platforms for under-resourced practitioners, and training programs to build AI literacy among public legal workers. This way, equitable access ensures that AI enhances justice rather than exacerbates disparities.

7. Training for Legal Professionals:

Effective AI integration requires that lawyers, judges, and policymakers understand both its capabilities and limitations. Mandatory training programs could cover algorithmic literacy, bias detection, privacy compliance, and ethical decision-making. Empowering legal professionals ensures that AI is interpreted correctly and applied reasonably, mitigating risks of misuse or overreliance.

8. Continuous Ethical Review:

AI evolves rapidly, so ethical review must be ongoing. Independent oversight bodies should monitor compliance with transparency, fairness, privacy, and accountability standards. Some implementation strategies could include periodic audits of AI performance and outcomes, updating ethical guidelines as AI technologies and legal practices evolve, and establishing grievance mechanisms for litigants affected by AI-assisted decisions. If applied uniformly and correctly, continuous oversight ensures that AI remains aligned with societal values and legal principles, rather than drifting toward unethical or opaque uses.

9. Ethical Use Monitoring and Feedback:

Beyond audits, AI systems should incorporate feedback mechanisms that allow stakeholders to flag errors, bias, or misuse. Real-time monitoring and corrective processes enable continuous improvement and ensure that AI outputs remain aligned with ethical and legal standards.

Discussion and Conclusion

This paper has examined the promise and perils of artificial intelligence (AI) in the legal profession, arguing that while AI is an inevitable and potentially transformative force, its integration must be bounded by human oversight and guided by ethical principles. Our central thesis has been that AI can enhance justice through efficiency, predictive accuracy, and expanded access, but with robust safeguards, it risks undermining fairness, transparency, privacy, and accountability.

Section I highlighted AI's benefits in reshaping legal practice. From accelerating document review to enabling predictive analytics and more advanced legal research, AI tools already help lawyers and judges work with unprecedented speed and accuracy. Section II, however, explored the threats AI poses to professional judgment, fairness, and access to justice. Bias embedded in training data threatens to replicate systemic inequalities, while "hallucinations" and opaque algorithms jeopardize accountability and due process. Section III turned to the ethical framework for responsible AI use, where we outlined guiding recommendations. We argued that human oversight must remain essential, and transparency and explainability should be mandatory, amongst other things. Together, these measures highlight that AI should support, not supplant, human judgment in law.

While this study offers an overview of the intersection between AI and law, it is limited by its reliance on secondary sources. Given the rapid evolution of generative and predictive AI technologies, much of the current literature may quickly become outdated, reflecting early-stage implementations rather than long-term effects. Moreover, the selected studies often emphasize Western legal frameworks, particularly those of the United States and the United Kingdom, which may not represent the diversity of global regulatory and ethical contexts. This focus introduces potential bias in understanding how AI interacts with differing legal traditions, levels of technological adoption, and societal values. Consequently, future research should aim to incorporate empirical, cross-jurisdictional, and longitudinal analyses to deepen understanding of AI's evolving role in legal systems. Though comprehensive in its nature, constraints of this study itself include limitations as to the geographical and geopolitical standings of the countries that the case studies examined refer to (mostly, again, the Westernized ones). Nonetheless, taken cumulatively, these findings suggest that the larger outcomes of AI's rise in law will depend on governance and ethics, not technological capacity alone. AI can either reinforce existing inequities or expose them for reform, depending on how it is developed, distributed, and overseen. This dual potential raises urgent questions for scholars and policymakers alike: How can we ensure equitable access to AI tools across the justice system? What methods best detect and mitigate algorithmic bias in legal contexts? Should there be international standards for transparency and accountability in legal AI systems? And how might legal education evolve to prepare future practitioners for a world where human judgment and machine intelligence must co-exist? Future research should continue to interrogate these questions, exploring interdisciplinary methods that combine computer science, ethics, and jurisprudence. In closing, this review underscores that AI's integration into the legal system is not simply a technical shift but a normative one, implicating the future of justice itself. By embracing AI with transparency, equity, and accountability, the legal profession can ensure that technology strengthens rather than subverts its most fundamental values.

Acknowledgments

I am deeply grateful to my mentor, Dr. Siddarth Krishnan, for his guidance, feedback, and encouragement throughout the development of this project. His insights have been invaluable in sharpening the arguments and ensuring the paper reflects both scholarly rigor and ethical reflection. Above all, I sincerely thank my family for their encouragement and support.

References

1. Surden, Harry, "Artificial Intelligence and Law: an Overview", *SSRN*. Published on 30 June, 2019; Last revised 6 May 2025; Last accessed on Aug. 26, 2025. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3411869
2. Agrawal, Aishwarya, "A comprehensive Review on Impact of Artificial Intelligence on Litigation", *Law Bhoomi*. Published on November 17, 2023; Last Accessed on Aug. 26, 2025. https://lawbhoomi.com/a-comprehensive-review-on-impact-of-artificial-intelligence-on-litigation/
3. Kabir, Shanin, "The role of AI technology for legal research and decision making", *International Research Journal of Engineering and Technology, Volume 10*. Published in July 2023; Last accessed on Aug. 26, 2025. https://www.irjet.net/archives/V10/i7/IRJET-V10I7148.pdf
4. Rosengrün, Sebastian, "Why is AI a threat to the Rule of Law", Volume 1, Article Number 10. Published on August 9th, 2022. Last accessed on Aug. 26, 2025. https://link.springer.com/article/10.1007/s44206-022-00011-5
5. *Getman, Yaroshenko, Shapoval, Prokopiev, Maryna Demura*, "The impact of Artificial Intelligence on legal decision-making", Mykolas Romeris University. Published Sept. 2023. Last accessed on Aug. 26, 2025. https://www.ceeol.com/search/article-detail?id=1210529
6. Alarie, *et al.*, "How Artificial Intelligence will affect the practice of law", *UTP*. Published on March 9th, 2018, in the University of Toronto Law Journal, Volume 68, Number supplement 1. Last accessed on Aug. 26, 2025. https://utppublishing.com/doi/abs/10.3138/utlj.2017-0052.
7. Markoff, John, "Armies of Expensive Lawyers, Replaced by Cheaper Software", *New York Times*. Published on March 4, 2011; Last accessed Aug. 26, 2025. https://www.nytimes.com/2011/03/05/science/05legal.html
8. Buchanan, B.G., & Headrick, T.E., "Some Speculation about Artificial Intelligence and Legal Reasoning", 23 Stanford Law Review, 40, 1970. Last accessed on Aug. 26, 2025. https://digitalcommons.law.buffalo.edu/cgi/viewcontent.cgi?article=1866&context=journal_articles

DOI: 10.36838/IJHSR811.58

9. Madaoui, Nadjia, "The Impact of Artificial Intelligence on Legal Systems". Published in 2024, Last accessed on Aug. 26, 2025.
10. Mehl, L. . " Automation in the Legal World". Mechanisation of Thought Processes conference, Teddington, England. Published in 1958
11. Patel, Shivani, "The impact of artificial intelligence on legal research and case analysis", *The Amikus Qriae*. Last accessed on Aug. 26, 2025, https://theamikusqriae.com/the-im
12. Dana, Remus & Levy, Frank, "Can Robots be Lawyers? Computers, Lawyers and the Practice of Law", SSRN. Published on Dec. 11, 2015; Last revised on Nov. 30, 2016. Last accessed on Aug. 26, 2025,https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2701092.
13. Callister, Paul D. "Generative AI and Finding the Law", *Law Library Journal*, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4608268
14. Weiser, Benjamin, "Here's What Happens When Your Lawyer Uses ChatGPT", *New York Times*. Published on May 27th, 2023, https://www.nytimes.com/2023/05/27/nyregion/avianca-airline-lawsuit-chatgpt.html
15. Surden, Harry, "The Ethics of Artificial Intelligence in Law: the Basic Questions", Chapter 38 https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3441303

Author

Ginevra is a rising senior at St. Stephen's School in Rome, Italy. She was born and raised in Rome and is applying to universities in both the US and UK with the intention of studying law, hoping to pursue a career in the legal field.