



ARTIFICIAL INTELLIGENCE AND THE FIELD OF LAW

Adina Maria PAICU*

*“Constantin Brâncuși” University of Târgu Jiu

Abstract: *Artificial intelligence is increasingly present in the legal field, from everyday research and drafting tools to systems that inform decisions inside courtrooms. This paper examines that shift along several lines. First, it looks at how AI tools are already changing everyday legal work, and whether their use strengthens or weakens legal reasoning in students and young practitioners. Second, it discusses a well-documented controversy over an algorithmic risk-assessment tool used in American criminal courts, which revealed a racial disparity in its errors. Third, it compares two regulatory models: the European Union’s Artificial Intelligence Act, which classifies AI used in the administration of justice as high-risk, and the more decentralised, case-by-case guidance that has emerged across United States courts. Finally, it offers practical recommendations for legal education. The paper concludes that artificial intelligence is unlikely to replace legal professionals, but that its responsible use depends on keeping a human professional genuinely able to check, understand, and take responsibility for its output.*

Keywords: *artificial intelligence, legal profession, algorithmic bias, AI regulation, access to justice*

**Contact details
of the
author(s):** *paicu.adina@gmail.com*

INTRODUCTION

For most of its history, the legal profession has changed slowly. Law libraries grew, procedures were digitised, and communication moved from letters to email, but the basic shape of legal work — reading, reasoning, writing, arguing — stayed remarkably constant. Artificial intelligence (AI) is the first technology in a generation that looks likely to change that shape itself, not just the tools used to do it. Over the past few years, AI systems have moved from the edges of legal practice — spam filters, simple document search — toward its centre: drafting contracts, summarising case law, and even, in limited and closely supervised ways, informing decisions inside courtrooms (Biresaw, 2023).

It helps to separate two very different ideas that often get mixed together in public conversation. One is AI as a helper: a tool that drafts a first version of a document, searches through



case law, or answers a simple question, while a trained professional still checks the result. The other is AI as a decision-maker: a system whose output directly shapes what happens to a real person, such as whether they are granted bail or how long their legal case takes to resolve. Much of the current debate about AI and law becomes much clearer once this distinction is kept in mind, because the risks, and the right safeguards, are not the same for both.

This paper looks at that question from several angles. It first examines how AI tools are already changing everyday legal work, and whether their use strengthens or weakens legal reasoning in students and young practitioners (Biresaw, 2023; Bednar et al., 2026). It then turns to a more contested use — AI inside the machinery of justice itself, through the lens of a well-documented controversy over algorithmic risk assessment in American courts (Angwin et al., 2016), before comparing how the European Union and the United States have each tried to regulate this technology, and what practical steps legal education might take in response (European Union, 2024).

MAIN TEXT

AI as a working tool. The least controversial, and by far the most widespread, use of AI in law is as a research and drafting aid. Legal research has traditionally meant searching through enormous volumes of statutes, case law, and commentary by hand or with basic keyword tools; AI-powered legal research platforms promise to do this faster, and in some cases to surface connections a human researcher might miss (Biresaw, 2023). Generative AI has extended this further. Large language models can now draft routine contracts, summarise depositions, and produce first drafts of legal memoranda in a fraction of the time a junior associate would need (Villasenor, 2024).

That shift brings a specific risk into focus: the reliability of the output itself. A 2025 study that tested both legal-specific and general-purpose language models on structured legal reasoning tasks, using the IRAC framework (Issue, Rule, Application, Conclusion) that law students are trained in, found that the models could perform a basic version of this analysis, but were limited by thin, underdeveloped answers, an unwillingness to commit to a clear conclusion, and a tendency to sound confident even when wrong (Peoples, 2025). For a profession where a wrong but confident-sounding answer can mislead a client or a court, that combination of traits is a genuine concern, not a minor inconvenience.

In practice, this already shows up in ordinary, unglamorous ways: a paralegal asking an AI assistant to summarise a hundred pages of emails before a trial, a small firm using AI to translate a foreign contract quickly before deciding whether a certified human translation is worth the cost, or a junior associate asking a chatbot to explain an unfamiliar area of law before turning to the official sources. In each case, a person still reads, judges, and signs off on the result — a pattern that, for now, is far more common than any dramatic vision of a fully automated law firm.

Does AI erode legal reasoning, or sharpen it? A natural worry follows from the research above: if AI can produce a passable legal analysis on demand, will lawyers — especially those still in training — lose the habit of reasoning through a problem themselves? A 2026 empirical study set out to test this directly, examining whether law students and junior lawyers who used AI tools showed any measurable decline in their understanding of the underlying legal principles afterward (Bednar et al., 2026). The results were more mixed than either AI’s critics or its enthusiasts tend to assume: the study found no overall decline in downstream comprehension, but concluded that AI’s effect on legal reasoning depends heavily on when and how it is used — as a shortcut that replaces thinking, or as a tool that structures and accelerates it. That conditional finding matters for legal education as much as for practice, since it suggests the relevant question is not simply whether to allow AI, but how to teach its use so that it supports rather than substitutes for independent judgment.



AI inside the courtroom: the COMPAS controversy. The most closely scrutinised example of AI actually influencing legal outcomes, rather than just supporting legal work, comes from the use of algorithmic risk-assessment tools in American criminal courts. COMPAS (Correctional Offender Management Profiling for Alternative Sanctions) is a proprietary tool used by judges and parole officers in several U.S. states to estimate the likelihood that a defendant will reoffend, informing decisions on bail, sentencing, and parole (Angwin et al., 2016).

In 2016, the investigative newsroom ProPublica studied COMPAS scores given to over seven thousand defendants in Florida, and compared them with what actually happened to those people two years later. The results showed a troubling pattern: Black defendants who never reoffended were mistakenly labelled “high risk” far more often than white defendants who also never reoffended (Angwin et al., 2016). The company that makes COMPAS disagreed with this conclusion, arguing that its scores were equally accurate for both groups overall. Later academic studies have confirmed that both sides had a point — the disagreement came from two different, both reasonable, ways of measuring “fairness,” which do not always point to the same answer.

The COMPAS case is useful precisely because neither side was simply wrong. It is not evidence that AI should never be used anywhere near a courtroom, nor is it evidence that human judges, on their own, are free of similar biases — earlier research had already found troubling patterns in purely human sentencing decisions too. What it does show is that handing a decision to an algorithm does not automatically make it more objective; it can simply make an existing problem harder to see, because a number on a screen can look more neutral than it actually is. That is precisely why transparency and the possibility of outside scrutiny — the kind that made the ProPublica investigation possible in the first place — matter so much whenever these tools are used in a legal setting.

Regulating legal AI: the EU’s approach, briefly. Partly in response to controversies like COMPAS, the European Union’s Artificial Intelligence Act, in force since 2024, classifies AI systems used by or on behalf of a judicial authority as “high-risk,” requiring providers to manage risks carefully and keep a human genuinely involved in the decision (European Union, 2024). As discussed below, that last condition — genuine human involvement — is also the hardest one to guarantee in practice (Banks, 2026).

A different approach: the United States. The European Union’s binding, risk-based regulation is not the only model. In the United States, there is no single federal law governing AI in the courts; instead, a patchwork of guidance has emerged court by court, and sometimes judge by judge. In July 2025, the Administrative Office of the United States Courts distributed interim guidance to all federal judges, advising them not to delegate core judicial functions to AI systems and instructing them to independently verify any material an AI tool produces (Administrative Office of the United States Courts, 2025).

That caution followed real incidents: federal judges themselves had issued orders containing AI-generated errors, and lawyers in multiple cases submitted filings with fabricated case citations that an AI tool had simply invented, sometimes resulting in sanctions (Administrative Office of the United States Courts, 2025). At the state level, approaches differ further still. The New York State Unified Court System adopted its first formal AI policy in October 2025, requiring mandatory training for judges and staff and stressing that judges “bear the ultimate responsibility” for anything AI helps produce, even while permitting its use as a supporting tool (New York State Unified Court System, 2025). Other states, such as Illinois, have taken a noticeably lighter touch, encouraging responsible AI use without requiring lawyers to disclose when they have relied on it.

The contrast is not only procedural but philosophical. The EU’s approach assumes that the categories of risk can be anticipated in advance, in enough detail to write binding rules around them



before most courts have even adopted the technology. The American approach assumes, in effect, the opposite: that the real risks will only become visible once lawyers and judges start using these tools in practice, and that courts are better placed to respond to specific, observed problems — a fabricated citation here, an over-confident judge there — than to predict every failure mode from a distance.

Each approach has a cost that the other avoids. Because the EU’s rules apply uniformly from the moment they take effect, a court in a small Romanian town and a court in a major Western European capital face the same baseline obligations, regardless of how much experience either has with the technology in practice. The American patchwork, by contrast, means a defendant’s lawyer might operate under a strict disclosure requirement in one courtroom and no requirement at all one county over, which is difficult to defend as a matter of basic fairness, even if it allows faster local experimentation.

Compared with the EU’s single, binding framework, the American approach is more decentralised and more reactive — responding to specific incidents as they surface, court by court, rather than setting uniform rules in advance. Neither model has fully solved the underlying problem. The EU’s high-risk classification sets clear conditions but, as discussed above, depends on oversight that is not always as meaningful as it appears on paper; the American patchwork reacts quickly to specific abuses but leaves considerable unevenness in how AI is treated from one courtroom to the next. Comparing the two suggests that neither binding regulation nor case-by-case guidance is, by itself, a complete answer — what both share, when they work, is an insistence that a human being remains answerable for the final decision.

Beyond courts and firms: AI and access to justice. Most discussion of legal AI focuses on large firms and courts, but a quieter application may end up mattering just as much: using AI to help ordinary people who cannot afford a lawyer understand basic legal information. Chatbots that explain tenants’ rights, generate simple wills, or help someone fill out a court form in plain language are already in use in several countries, and research on AI’s role in legal work notes that one of its likely effects is broadening who can access legal information in the first place (Biresaw, 2023). This use case sits in an interesting middle ground: a chatbot that explains general legal concepts to the public is not, in itself, an AI system used by a judicial authority to decide a case, so it does not automatically fall into the Act’s high-risk category for the administration of justice (European Union, 2024). Yet the stakes for the person using it can be just as high as in a courtroom if the information is wrong and they act on it without ever consulting a professional — a risk made more concrete by the same 2025 finding that language models can sound confident even when their legal analysis is thin or wrong (Peoples, 2025).

Practical recommendations for legal education and training. If AI’s effect on legal reasoning genuinely depends on how it is used, as the evidence above suggests (Bednar et al., 2026), then legal education has a direct role to play in shaping that use, rather than simply reacting to it after students graduate. A first, modest step is explicit instruction: teaching students not only how to prompt an AI tool effectively, but also how to recognise the specific failure modes documented in the research discussed above, such as a confident-sounding answer that turns out to be incomplete or simply wrong (Peoples, 2025).

A second step is building verification into the workflow itself, rather than treating it as an optional extra step students may skip under deadline pressure. The disclosure and verification requirements that have appeared in several U.S. federal court orders — requiring lawyers to certify that AI-generated citations were checked against the real case law — offer a useful model that law schools could adopt earlier, in the classroom, well before students encounter the consequences of skipping that step in a real courtroom (Administrative Office of the United States Courts, 2025).



A third step is mandatory, periodically refreshed training, rather than a single orientation session. The New York State Unified Court System’s decision to require ongoing AI training for all judges and non-judicial staff, not just a one-time briefing, reflects an implicit acknowledgment that both the technology and the risks around it keep changing (New York State Unified Court System, 2025). Law faculties preparing students for a career that will span decades have good reason to adopt the same logic: a single lecture on “AI in law,” delivered once, is unlikely to remain relevant for the whole of a student’s career, given how quickly both the tools and the surrounding rules continue to change.

Taken together, these recommendations point toward a modest but consistent principle: legal education should treat AI literacy the way it already treats legal research or courtroom procedure — as a core, examinable skill, taught explicitly, rather than something students are simply expected to absorb on their own. Institutions that take this seriously now are likely to produce graduates who are better prepared for a profession that, as this paper has argued throughout, is changing shape around them whether they engage with that change deliberately or not.

CONCLUSION

Open questions for future research. Several questions remain open, and are worth naming rather than glossing over. First, the EU AI Act’s obligations for high-risk systems used in the administration of justice only become fully applicable from December 2027 onward (AI Act Explorer, 2026), so most of the evidence about how well its human-oversight requirement works in practice is still, necessarily, anticipatory rather than observed. Second, almost all of the empirical research reviewed in this paper comes from a small number of jurisdictions, overwhelmingly the United States and the European Union; how these dynamics play out in legal systems with different court structures, different languages, and different levels of digital infrastructure remains comparatively understudied.

Third, the COMPAS controversy has become something close to a standard reference point in this literature, but it describes one specific tool, in one country, used for one narrow purpose — recidivism risk scoring. Whether its particular pattern of error generalises to other kinds of legal AI, such as document review software or AI-assisted legal-aid chatbots, is an empirical question rather than something that can simply be assumed. Future research that tracks these newer, more ordinary applications with the same rigour that journalists and academics brought to COMPAS would do a great deal to clarify which of this paper’s conclusions hold more generally, and which are specific to high-stakes criminal justice settings.

Artificial intelligence is not, on the evidence gathered here, poised to replace lawyers or judges outright. What it is doing is changing the shape of legal work at both ends of the profession: making routine research and drafting faster and cheaper, while raising hard new questions about fairness, oversight, and reliability wherever AI output starts to influence decisions that affect someone’s liberty or rights.

The throughline across the research examined in this paper is not that AI is simply good or bad for law, but that its effects depend heavily on how it is embedded into existing legal institutions and habits of thought, and on which regulatory model — binding and centralised, like the EU’s, or decentralised and case-by-case, like the United States’ — happens to surround it. Used as a tool whose output is checked, understood, and ultimately owned by a human professional, AI appears to speed up legal work without eroding the reasoning underneath it. Used as an unquestioned source of scores, drafts, or predictions, it risks importing exactly the kind of opaque, hard-to-challenge authority that legal systems have spent centuries building procedures to prevent (Banks, 2026).



It is worth pausing on a concrete illustration of how quickly that caution can be vindicated. Within a single year of generative AI tools becoming widely available to lawyers, courts on more than one continent had already encountered filings containing case citations that did not exist — invented by an AI system with no indication, to the lawyer who submitted them, that anything was wrong (Administrative Office of the United States Courts, 2025). The lesson was not that generative AI is unusable in legal practice; firms and courts that had built a verification step into their process caught the problem before it caused harm. The lesson was that the technology’s fluency can outpace, by a wide margin, a user’s instinct to doubt it — and that instinct, not the technology itself, is what legal training now has to reinforce deliberately.

That, in the end, is the single standard this paper has tried to defend throughout: not whether artificial intelligence belongs in the field of law, but whether the humans working alongside it remain genuinely in charge of the judgment it is used to support.

In the meantime, the most reasonable working position for students, practitioners, and institutions alike is a cautiously engaged one: neither dismissing these tools as a passing novelty, nor adopting them uncritically because they are fast and fluent, but treating every AI-assisted output as a draft that earns trust only once a qualified human has actually checked it.

A final, practical note follows from all of this for a reader based in Romania. As a European Union member state, Romania is directly bound by the AI Act, and Romanian courts, law schools, and law firms will face the same high-risk classification, the same oversight obligations, and the same open question about whether that oversight will be meaningful in daily practice (European Union, 2024). The American experience reviewed above is not a blueprint Romania is obliged to follow, but it does offer an early warning: fabricated citations and over-trusted AI output are not hypothetical risks confined to one legal system, and a Romanian legal education that takes AI literacy seriously now, before a comparable incident forces the issue, is likely to serve its graduates considerably better than one that waits to react.

This also helps explain why the paper has treated the European and American regulatory responses as complementary rather than competing. A binding rule, like the EU’s high-risk classification, sets a floor: it guarantees that certain safeguards exist even where no individual judge or firm thinks to ask for them. Case-by-case guidance, like the orders issued by individual American judges after specific incidents, adds a ceiling of practical detail that a general rule, written years in advance, cannot always anticipate. Read together, the two approaches sketch out what a mature legal response to AI probably needs: enforceable minimum standards, combined with the institutional flexibility to respond quickly when those standards turn out to have gaps.

That same caution is worth extending to this paper’s own conclusions. The research reviewed here spans only a few years, and both the technology and the rules around it are still moving targets: a law student reading this paper five years from now may find that some of its specific examples, like the exact wording of a court’s interim AI guidance, have already been superseded by more settled rules. What is less likely to go out of date is the underlying principle running through every section of this paper — that the value of an AI tool in a legal setting depends less on how sophisticated it is, and more on whether a human professional remains genuinely able, and genuinely willing, to question what it produces.

None of this is a reason for alarm, nor for uncritical enthusiasm. Technologies that genuinely change a profession rarely do so overnight, and the legal profession, more than most, has institutional habits — precedent, procedure, professional ethics rules — that tend to absorb new tools slowly and deliberately. That slowness is not simply inertia; it reflects the fact that errors in legal work can cost someone their liberty, their property, or their family, and a profession built around that responsibility



has good reason to be cautious about any tool, however impressive, that it does not yet fully understand.

For law students and young legal professionals in particular, this points toward a practical takeaway rather than an abstract one. Learning to use these tools well is likely to matter just as much, in the coming years, as learning to use a law library mattered a generation ago — not as a replacement for legal judgment, but as a new and increasingly unavoidable part of how that judgment gets exercised in practice. The lawyers best placed to benefit from AI will probably not be the ones who trust it blindly, nor the ones who refuse to engage with it at all, but the ones who learn early to treat it the way a careful professional treats any assistant: useful, fast, occasionally wrong, and never a substitute for their own final judgment.



REFERENCES

- Administrative Office of the United States Courts. (2025). *Letter to the Senate Judiciary Committee regarding Judiciary use of artificial intelligence*.
https://www.grassley.senate.gov/imo/media/doc/ao_to_grassley_re_judiciary_use_of_ai.pdf
- Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016). *Machine bias*. ProPublica.
<https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- AI Act Explorer. (2026). *High-level summary of the AI Act*. <https://artificialintelligenceact.eu/high-level-summary/>
- Banks, I. (2026). Judges-in-the-loop? Judicial involvement in human oversight of high-risk decision support systems under the EU AI Act. *International Journal of Law and Information Technology*, 34, eaag001.
<https://doi.org/10.1093/ijlit/eaag001>
- Bednar, N., Cleveland, D. R., Erbsen, A., & Schwarcz, D. (2026). Artificial intelligence and human legal reasoning. *Minnesota Legal Studies Research Paper 2026-21*. <https://ssrn.com/abstract=6525800>
- Biresaw, S. M. (2023). *The impacts of artificial intelligence on research in the legal profession*.
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4727017
- European Union. (2024). *Regulation (EU) 2024/1689 (Artificial Intelligence Act), Recital 61 – High-risk AI systems in the administration of justice*. <https://ai-act-law.eu/recital/61/>
- New York State Unified Court System. (2025). *Interim policy on the use of artificial intelligence*.
<https://www.nycourts.gov/LegacyPDFS/a.i.-policy.pdf>
- Peoples, L. F. (2025). Artificial intelligence and legal analysis: Implications for legal education and the profession. *SSRN*.
<https://ssrn.com/abstract=5123122>
- Villasenor, J. (2024). Generative artificial intelligence and the practice of law: Impact, opportunities, and risks. *Minnesota Journal of Law, Science & Technology*, 25(2), 25. <https://scholarship.law.umn.edu/mjlst/vol25/iss2/8>