



Memory, Judgment and the Future of Law

By Laina Chan

Snapshot

- AI can recall everything the law has said, but only humans can decide what it means.
- When AI is trained to reason — to connect issues, rules and outcomes — its explanations become more plausible and its results more trustworthy.
- Judgment cannot be automated: the legitimacy of law depends on reasoning that can be seen, tested and justified.

Artificial intelligence has given the profession something close to a perfect memory. Every judgment, every argument, every precedent can now be summoned instantly. For the first time in history, we face a machine that seems to know everything. But perfect recall is not the same as understanding. Knowing every case is not the same as knowing which one matters.

That difference between information and judgment has always defined the craft of advocacy. The best lawyers do not win by citing the most authorities. They win by choosing the right ones, by discerning which principle is controlling and which case must be distinguished. Advocacy has always depended on judgment, not accumulation. AI can now retrieve everything that has ever been decided, but it still cannot tell us what will persuade.

Recent academic research has confirmed what every lawyer already knows instinctively: that reasoning, not recall, is the heart of law. In 2024, Irene Benedetto and her colleagues published a study in *Artificial Intelligence and Law* showing that when AI systems were trained to recognise **legal entities that carry meaning**, such as the parties, statutes, and relationships between issues, the models will produce more plausible explanations that are

more aligned with human reasoning.¹ In my experience, the more a system resembles a lawyer carefully building an argument, the better it will perform. It is not the volume of data that matters but the structure of reasoning. The research demonstrated that explanation and accuracy rise together; an AI that can show its working is more reliable than one that merely guesses well.

A companion study by Luyao Ma and co-authors in 2021 drew an even sharper distinction between prediction and reasoning.² Their team tested judgment-prediction systems on 70,000 real Chinese court cases in the private lending category. Most prior models had been trained on **judge-edited summaries** — the polished, factual narratives written after trial. On those datasets, the machines achieved impressive accuracy.

However, when the same models were tested on the messy, contradictory transcripts of live trials, with incomplete evidence and competing versions of events, accuracy dropped dramatically. The finding was both technical and philosophical: machines could imitate judicial language, but not judicial doubt. They performed well when the facts were fixed, but law is lived in uncertainty. That gap between the cleaned hindsight of data and the contested reality of decision-making remains the difference between machine logic and human judgment.

Of course, not everyone agrees that predicting judicial decisions is appropriate, regardless of how advanced the technology may become. As Lance Eliot observes, the question is not just whether AI can predict, but whether it *should*. In 2019, France drew a clear line when it amended its *Justice Reform Act* to criminalise the publication of “judge analytics.” The new Article 33 prohibits the reuse of personally identifiable data about judges or court clerks “for the purpose or effect of evaluating, analysing or predicting their actual or supposed professional practices,” and violations can attract penalties of up to five years’ imprisonment and a €300,000 fine.³ The law applies to researchers, companies and individuals alike. Its message is simple but profound: transparency cannot come at the expense of judicial

¹ Irene Benedetto et al., “Boosting Court Judgment Prediction and Explanation Using Legal Entities,” *Artificial Intelligence and Law* (2024) at page 609.

² Luyao Ma et al., “Legal Judgment Prediction and Explanation Using Legal Entities,” *Proceedings of SIGIR* (2021) at Table 3 (MSJudge-MTL micro-F1 86.5%); pp 3–4 (prior work uses judge summaries); p5 (70k debates; noise vs facts).

³ *Loi n° 2019-222 du 23 mars 2019 de programmation 2018-2022 et de réforme pour la justice* (France) art 33, published *Journal officiel de la République française* 24 March 2019; see also Lance B. Eliot, “Legal Judgment Prediction amid the Advent of Autonomous AI Legal Reasoning,” *arXiv preprint* arXiv:2009.10350 (2020) at 8–9, citing Rebecca Loescher’s English translation.

independence. The legitimacy of the courts rests not only on what is decided, but on the assurance that no one is ranking the judges who decide it.

Judgment prediction is sometimes described as if it were a single, solvable problem—a matter of matching facts to outcomes. In reality, it is an entire constellation of questions. As Lance Eliot notes, prediction depends on how a case is framed, which facts are deemed material, which precedents are selected, how the relevant principles are interpreted, and how competing policy considerations are weighed. Each of these steps involves value choices, not mechanical ones. Compressing that complexity into an algorithm risks mistaking the surface of law for its substance. A program can model patterns in language or outcome, but it cannot capture the layers of reasoning, discretion and ethical judgment that turn information into a decision. The process of judgment is not a formula to be computed; it is an argument to be justified.⁴

Another way of viewing this is that Legal Judgment Prediction is impossible given the exigencies of litigation. The outcome of a serious dispute is rarely certain. In every substantial case there are arguments that cut both ways, and reasonable minds will differ about who has the better of it. A trial judge's view may be reversed on appeal, and that reversal may itself be overturned by a higher court. This is not weakness; it is how the law grows. Uncertainty is the condition that makes reasoning—not prediction—the discipline of our profession.⁵

MiAI Law was therefore built to reason. In law, a wrong answer that can be explained is better than a right answer that cannot. Proof, not probability, is the measure of integrity. Every MiAI Law report is footnoted, audit-ready and transparent. Each conclusion can be traced to primary authority, and when the evidence is insufficient the system says so. It reasons from first principles but leaves judgment where it belongs — with the lawyer.

Courts decide human disputes, not mathematical problems. An algorithm may map correlations between words and outcomes, but it cannot weigh hesitation, emotion or credibility. The law's moral dimension — the empathy that recognises suffering, the courage to depart from precedent when justice requires it — remains beyond computation. The High Court's decision in *Mabo v Queensland (No 2)* illustrates the point. No model trained solely on historical precedent could have produced that outcome. It took human reasoning to see that law must evolve when fairness demands it.

⁴ Eliot, L. B. (2020), "Legal Judgment Prediction amid the Advent of Autonomous AI Legal Reasoning," arXiv preprint arXiv:2009.10350 at p.9.

⁵ See also Eliot, L. B. (2020), "Legal Judgment Prediction amid the Advent of Autonomous AI Legal Reasoning," arXiv preprint arXiv:2009.10350 at p. 8 quoting Atkinson et al.

The provision of reasons is a core requirement and the fundamental entitlement of parties to any litigation. Each step toward automated decision-making requires greater transparency to preserve confidence. Experience suggests that systems that assist and explain will be accepted; those that conceal their processes will be rejected. That principle guided MiAI Law's design. Its architecture is intentionally limited to **assistive reasoning** — automating retrieval and structure while keeping interpretation in human hands. Transparency is not an afterthought. It is the boundary that sustains trust.

These insights reaffirm something fundamental. The law's legitimacy rests not on omniscience but on method — on the discipline of reasoning that can be inspected, challenged and improved. AI can assist that process, but it must not replace it. The temptation to automate judgment entirely misunderstands what law is. A judgment is not merely the product of rules applied to facts; it is the expression of responsibility for a decision that affects human lives.

MiAI Law was built to make that responsibility clearer, not lighter. It preserves the discipline of verification. If a user relies on an answer without checking the underlying authorities, the system gently reminds them to review the sources. It encourages diligence rather than complacency — a design choice that turns professional ethics into engineering.

Even perfect recall can become its own danger. When every authority is instantly available, discernment becomes scarce. The best lawyers will be those who use the machine's speed to sharpen their judgment, not to drown in material. In advocacy, as in design, simplicity still persuades. A well-chosen precedent will always carry more weight than a flood of minor authorities.

The future of the profession therefore lies not in competing with machines but in mastering them. Technology can extend our memory, but only we can give that memory meaning. It can lighten the labour of research, but not the duty of care. It can make the law more accessible, but only if we preserve its humanity.

MiAI Law's purpose is to keep those values intact. It encodes transparency, enforces accountability and respects the boundary between reasoning and judgment. It is a bridge between the tradition of proof and the promise of progress — a reminder that the law's greatest innovation has always been its insistence on explanation. AI can learn the shape of our arguments, but only we can ensure they remain just.