LSJ ONLINE

A world without lawyers: is automated decision-making the future?

AAA

BY LAINA CHAN - MAY 09, 2025 8:40 AM AEST

Snapshot

Legal practice has long relied on trust, relationships and reputation, yet AI is transforming the landscape.

efficiency.

In legal practice, the selection of a lawyer is rarely arbitrary. It is often grounded in personal recommendations, existing professional relationships, or the lawyer's reputation for competence and integrity. Trust plays a pivotal role. A client does not just seek technical expertise but reassurance – a feeling that their lawyer will advocate zealously and guide them through complex, often emotionally charged processes.

For example, a corporate client negotiating a high-stakes merger might choose a firm with which they have established trust through years of successful transactions. Similarly, a distressed individual facing a custody dispute will likely rely on a recommendation from a friend or family member who found comfort and resolution with a particular lawyer. In both scenarios, the human element – the ability to empathise, strategise and adapt – is indispensable.

Artificial intelligence ('**AI**'), however, operates differently. It promises speed, consistency and data-driven insights but it lacks the emotional intelligence and personal rapport central to human legal counsel. Can AI, a fundamentally impersonal tool, address the complexities of legal practice as well as replicate the trust and relational dynamics foundational to legal practice?

AI's expanding footprint in legal services

The legal sector has already embraced various AI-driven tools to enhance efficiency and reduce costs. Examples include:

1. **document review and e-discovery**: AI systems can sift through millions of documents, identifying relevant materials far more accurately and quicker than human associates;
2. **contract analytics**: platforms like Kira and Luminance promise to detect anomalies, missing clauses or potential risks in contracts with unprecedented speed;
3. **predictive analytics**: tools such as Lex Machina analyse judicial patterns to forecast case outcomes, enabling lawyers to strategise more effectively; and
4. **self-service legal assistance**: platforms like DoNotPay provide users with automated solutions for minor legal matters, such as disputing parking tickets or generating basic contracts.

These innovations highlight the growing potential of AI to handle repetitive and high-volume tasks. However, as AI systems evolve, so do the complexities associated with their adoption. Questions about transparency, ethical considerations and the role of human oversight remain unresolved. In a recent consultation paper, the Victorian Law Reform Commission identified several potential risks associated with using AI in the judicial system. These include accountability, transparency and explainability, bias and fairness, privacy and data security, impact on public trust, dependence on AI and loss of judicial discretion, and errors and reliability.

client. However, AI does not always identify the nuances in legal documents or summarise all the key points in those documents. Merely relying upon the inconsistences flagged by AI might lead to key issues being overlooked.

While there is the potential for the application of AI to extend to legal operations teams in multinational corporations, that future is not currently a reality. Existing large language models ('**LLMs**'), like ChatGPT, struggle with legal interpretation and this author is unaware of an existing due diligence system that is able to account for jurisdiction-specific nuances. While AI systems will eventually streamline due diligence, providing comprehensive risk assessments within hours rather than weeks, that will require the development of a complex system which incorporates finetuned LLMs that are specifically trained to read and understand law as well as jurisdiction-specific vector databases (a type of database optimised for storing, indexing and querying complex data, discussed below). This underscores the current necessity for human review and legal representation across various jurisdictions. It is difficult to envisage how this need for human review can ever be completely obviated.

Understanding legal AI research: the role of RAG systems

Legal AI research often involves the use of Retrieval-Augmented Generation ('**RAG**') systems which combine retrieval-based techniques with generative AI models. These systems are designed to search vast legal databases for relevant information and generate context-sensitive responses based on the retrieved material. RAG systems will increasingly be deployed to streamline legal research and improve efficiency.

How RAG systems work

The system uses a retriever component to search a legal corpus – such as case law, statutes or academic commentary – and identifies documents or snippets relevant to the input query. This process often relies on vector databases that transform legal text into numerical embeddings, enabling sophisticated semantic searches. After retrieving the most relevant information, a generative AI model processes the data to generate coherent, contextually accurate answers or summaries.

For example, a lawyer researching recent interpretations of section 63 of the *Civil and Administrative Tribunal Act 2013* (NSW) which deals with the power of the Tribunal to correct errors in its decisions, might input their query into a RAG-powered tool. The system retrieves relevant judgments and statutory interpretations, then generates a summary outlining the key principles discussed in those cases.

While RAG systems can significantly improve the speed and depth of legal research, they also introduce a range of challenges and risks, particularly when dealing with the complexities of legal information.

Transparency in decision-making is a cornerstone of justice. Litigants have the right to understand how and why decisions

Challenges arising in legal AI research and RAG systems

Hallucinations in RAG systems

One of the most significant issues with RAG systems is hallucination: the phenomenon where the generative component produces content that is factually incorrect, unsupported by the retrieved documents or entirely fabricated. This can lead to severe consequences in legal research where accuracy and credibility are paramount.

For example, a RAG system generating a summary of case law might include a fabricated precedent or misinterpret a judicial ruling. If a lawyer relies on this information without verification, it could lead to flawed legal arguments or even adverse outcomes for their client.

The adoption of these mitigation factors can reduce the risk of hallucinations to a negligible risk:

validation pipelines*:* implement processes to cross-check generated outputs against the retrieved documents. This ensures the content is grounded in actual data;

source anchoring*:* design the system to explicitly link generated content to specific excerpts from the retrieved materials, making verification easier;

human oversight*:* ensure lawyers review and refine AI-generated outputs to confirm their accuracy and relevance; and

training on legal contexts*:* improve the AI's training data by including diverse and representative legal materials, minimising the risk of hallucination due to gaps in knowledge.

Over-reliance on algorithms

Another risk of AI in legal research is the potential for the legal algorithm to only reveal some but not all the relevant cases. AI systems are designed to prioritise results based on complex ranking algorithms, often tailored by proprietary databases. While these systems excel at surfacing relevant cases, they may inadvertently exclude critical but less obvious precedents or arguments. For example, an AI system might prioritise recent case law over older but still authoritative decisions, leading to incomplete research.

There is the risk that a lawyer relying exclusively on AI tools might miss an obscure precedent from decades ago that fundamentally alters the legal argument. This oversight could jeopardise the outcome for the client. It is therefore important for the legal AI research tool to traverse the entire vector database to ensure the full picture is revealed. Human oversight is also critical as the experienced human lawyer may recall the obscure precedent.

Some potential solutions to this challenge could be:

auditability: legal AI tools should incorporate features that allow lawyers to trace the logic behind search results, ensuring accountability and trust;

cross-validation: using multiple AI tools or combining AI-driven research with traditional methods can help mitigate the risks of incomplete or biased results; and

The black box problem in legal AI

The 'black box' problem in AI refers to the opacity of decision-making processes in complex machine learning models, particularly those relying on neural networks. These systems process vast amounts of data through layers of algorithms, producing outputs that even their developers cannot fully explain. This lack of transparency is particularly concerning in the legal context where reasoned justifications are paramount.

For instance, a bail recommendation tool might classify a defendant as 'high risk' without revealing which factors were weighted most heavily. Was it prior convictions? Employment history? Neighbourhood demographics? Without transparency, it is impossible to evaluate whether the recommendation was fair or accurate.

Why transparency matters in legal practice

Transparency in decision-making is a cornerstone of justice. Litigants have the right to understand how and why decisions affecting them were made, whether by judge, jury or algorithm. This ensures:

accountability: a clear rationale allows for the identification and correction of errors;

appeals and oversight: judicial decisions are subject to appeal, but if an AI system's logic cannot be scrutinised, meaningful appeals become impossible; and

public confidence: trust in the legal system erodes when decisions appear arbitrary or opaque.

In the United States, the COMPAS algorithm – a recidivism risk-scoring model used throughout the criminal justice system – has been widely criticised for its opacity in predicting recidivism risks. Defendants and their attorneys have struggled to challenge its recommendations due to the lack of access to its underlying logic, raising concerns about due process.

Mitigating the black box problem

Several strategies can address the black box issue:

explainable AI ('**XAI**'): tools, like SHAP (SHapley Additive exPlanations), and methods, like LIME (local interpretable model-agnostic explanations), aim to clarify which data points influenced an AI's output most. For example, an AI recommending a precedent could highlight the factual similarities between the case at hand and the cited ruling;

audits and validation: regular independent audits can test AI systems for accuracy, consistency and fairness. For instance, a sentencing algorithm could be reviewed periodically to ensure it does not disproportionately penalise certain demographics;

open-source development: transparency increases when AI systems are built collaboratively and their code is open to scrutiny by legal and technical experts; and

human oversight: maintaining a 'human in the loop' ensures lawyers or judges can override AI recommendations when they conflict with ethical standards or common sense.

Imagine an AI system used for zoning decisions in urban planning. If the tool recommends rejecting a proposal, citing traffic concerns, explainability tools should identify specific datasets – traffic flow models, prior zoning rulings or demographic patterns – to clarify the basis of the decision.

for oversight.

Let's be honest: a world entirely without lawyers is unlikely. Not because machines cannot perform mechanical legal tasks, but because justice demands more than mechanical efficiency.

To mitigate this, legal practitioners can:

- negotiate service-level agreements requiring transparency from AI vendors. For example, the proprietary legal research tool might reveal how it prioritises cases within a legal system;
- cross-reference results with a secondary research platform or manual review; or
- manually review the results of the legal AI tool.

Bias in legal AI systems

AI systems are only as impartial as the data they are trained on. Bias arises from:

- **historical inequities**: if training data reflects systemic discrimination – for example, racially biased sentencing patterns – the AI will replicate and potentially amplify these biases;
- **feature selection**: algorithms might inadvertently rely on proxies for protected characteristics, such as using zip codes as a stand-in for race or socioeconomic status; and
- **underrepresentation**: minority groups or niche legal scenarios may be underrepresented in training datasets, leading to poorer performance for these cases.

Biased AI systems risk perpetuating injustice at scale. For instance, a predictive policing algorithm might disproportionately flag minority neighbourhoods for heightened surveillance. Alternatively, a bail recommendation tool could unfairly deny release to defendants based on historical data rather than individual circumstances.

In 2013, a study of COMPAS by ProPublica revealed it was twice as likely (45 per cent versus 23 per cent) to falsely label African-American defendants who did not recidivate over a two-year period as high risk compared to Caucasian defendants, underscoring the systemic issues embedded in historical data.

Measures to address such bias issues could include:

- **diverse and representative datasets**: training data must reflect the full spectrum of human experiences and minimise historical biases;

human overrides: lawyers and judges must retain the ability to challenge or disregard AI outputs that appear biased or unjust.

Some jurisdictions have proposed legislation requiring any AI used in judicial decision-making be auditable for fairness. For example, the European Union's *AI Act 2024* emphasises fairness and transparency in high-risk AI systems, including those used in justice. In the United States, the *Algorithmic Accountability Act 2023* called for regular audits and transparency in AI tools, providing safeguards against entrenched biases. However, on 20 January 2025, President Donald Trump signed an executive order rescinding the 2023 directive issued by former President Joe Biden on artificial intelligence.

So, should we replace lawyers with AI?

Arguments in favour of AI include:

cost reduction: AI can handle routine tasks more cheaply than human lawyers, making legal services accessible to a broader audience;

efficiency and speed: machines operate faster and with greater consistency, reducing delays in research and decision-making; and

data-driven insights: AI can analyse patterns across vast datasets, offering insights that no individual lawyer could discern.

In a multinational merger, AI tools can potentially significantly reduce the cost of cross-border due diligence. In litigation, properly designed AI tools can substantially reduce the cost of reviewing document disclosures.

Arguments against AI include:

lack of emotional intelligence: AI cannot empathise or build rapport with clients, which is essential, particularly in emotionally charged cases like divorce or criminal defence;

ethical concerns: machines lack moral judgment and may inadvertently produce recommendations that conflict with broader societal values;

transparency issues: the black box problem undermines accountability and due process; and

bias: AI risks embedding and scaling historical inequities.

Lawyers typically tailor their strategy based on a client's personal circumstances, a nuance that AI systems struggle to replicate.

Hybrid solutions: humans and AI in tandem

Rather than replacing lawyers, AI is likely to complement them. A hybrid model could leverage AI's strengths – speed, scalability and pattern recognition – while preserving the human elements of legal practice. Some examples of humans working well in tandem with AI could include:

AI-assisted research: an AI might identify relevant precedents, which a lawyer can analyse and adapt to the client's specific needs;

For example, in a data breach case, an AI system could generate a comprehensive chronology of events, while a lawyer uses the findings to craft arguments tailored to the client's business and reputational priorities. The combination of AI and human expertise will enable more informed and proactive legal strategies, such as risk prevention and predictive litigation management.

Conclusion: balancing transparency, trust and fairness

The legal profession stands at a crossroads. AI offers transformative potential, promising to democratise access to justice and enhance efficiency. Yet significant challenges remain, particularly regarding transparency, bias and preservation of the human touch.

Let's be honest: a world entirely without lawyers is unlikely. Not because machines cannot perform mechanical legal tasks, but because justice demands more than mechanical efficiency. It requires empathy, ethical judgment and the ability to adapt to the complexities of human life. By embracing a hybrid model that integrates AI's capabilities with human expertise, the legal profession can evolve while preserving its core values. Transparency, fairness and trust must guide this evolution, ensuring that technology serves justice without compromising its foundations.

This article has been prepared with the assistance of AI but has been thoroughly verified, edited and enhanced by the author.

Laina Chan is a barrister at 2 Selborne Chambers and CEO of Miai Law Pty Ltd.

TAGGED IN | **Legal updates** AI artificial intelligence future of law generative ai

Related Articles

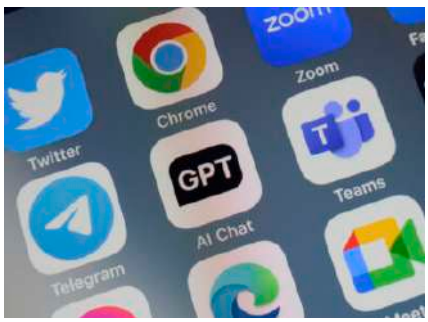

LSJ ONLINE