

Comment on ‘Legal Futures: Where Should Artificial Intelligence Take Us?’

Dr Robert Mullins*

There are only two points that I wish to make in my comments tonight. Professor Bennett Moses has already made both points in her talk with considerable insight. Tonight, I hope to offer a few brief words of elaboration and illustration.

The first point is that, as miraculous as the outputs of generative AI can look, they are not the outputs of human cognition. Large Language Models (LLMs) function on an underlying logic that is (roughly) the same as the logic underlying predictive text. Given a sequence of tokens, the model predicts the next token that is most likely to be relevant or useful. In combination with reinforcement learning (using feedback from a human or other LLM), these models can be trained to generate outputs that are contextually appropriate, in the sense that they are both optimised for the task for which they are designed and unlikely to produce socially inappropriate content.

The second point, and the point which I want to spend the remainder of my time elaborating upon, is that the widespread use of these technologies will, I think inevitably, restructure or reorder our cognition. I think their use in law and as an aid to various processes of legal reasoning will quickly come to feel unremarkable and insignificant. Because the changes brought about by use of these technologies are not as dramatic as some are predicting, we are at risk, as Professor Bennett Moses suggests, of losing sight of what we stand to lose by adopting them. One need not be a thoroughgoing Luddite to reflect on these losses. As the late political philosopher Gerry Cohen observed, the ‘proposition...that humankind is a net beneficiary of modernization, is not a reason for not lamenting what has been lost’.¹

The Luddites named their movement after the folk hero King (Ned) Ludd, destroyer of stocking frames. Tonight, I wanted to start with the warnings of a famous proto-Luddite. In her discussion, Professor Bennett Moses has already observed that, to date, the principal technology that has reshaped legal practice is written and printed script. At various points in his dialogues, Plato gives voice to his suspicion of writing and written script. In the *Phaedrus*, Plato’s Socrates warns of the effects of writing on memory, and of its tendency to destroy the possibility of true dialogue with the author.² In the *Statesman* Plato likewise warns against the error of treating written laws as a substitute for the wisdom of the lawgiver. His ‘Stranger’ compares the lawgiver who issues laws in writing to a doctor who leaves a written note for their patient—like the doctor’s note, the written law is preferable to

* Senior Lecturer, School of Law, University of Queensland. Email: r.mullins@law.uq.edu.au.

¹ GA Cohen, ‘Rescuing Conservatism: A Defense of Existing Value’ in R Jay Wallace, Rahul Kumar and Samuel Freeman (eds), *Reasons and Recognition: Essays on the Philosophy of T.M. Scanlon* (Oxford University Press 2011) 213.

² *Phaedrus*, 274a-278a; see Plato, John M Cooper and DS Hutchinson, *Plato: Complete Works* (Hackett Publishing Company, Incorporated 1997) 683.

having no instruction at all, but it is a poor replacement for the wisdom of the writer.³ As Walter Ong observed in his landmark study, *Orality and Literacy*, many of Plato's objections to the use of writing and written codes of law bear more than passing resemblance to arguments mounted against computers in their early adoption.⁴ These worries, I should reiterate, are not wrong just because they are recurrent. Writing is a technology, and it is a technology that has restructured our thought. As a result, we have lost things that are valuable—much of the creativity and invention associated with traditions of oral poetry and oral law of great beauty and complexity is lost to us, except as we can piece it together from their written remnants.⁵

I do not mean to imply that the birth of generative AI is comparable in significance to the birth of writing. But if writing restructures thought, there is no reason to think that the widespread use of the various forms of AI that Professor Bennett Moses mentions will not do the same.

The effect of the widespread adoption of technologies of writing on our legal system can be seen in the evolution of the common law. The doctrine of precedent as we know it today was made possible by the practice of reporting written judgements.⁶ Early commentators on the English common law, like Coke, Hale, and Blackstone, wrote of precedents as merely providing 'evidence' of the law—they were not a source of law in their own right.⁷ Written reports were unreliable and incomplete. Precedents were often simple inconsistent. The formalization of the practice of official case reporting provided one of the conditions necessary for the emergence of the doctrine of *stare decisis*.

The digitization of case reporting and introduction of case citators has made the use of case authority even more exhaustive. On any question of law, an undergraduate student with a modicum of training can now have access to a near-exhaustive list of authorities from any jurisdiction, reported and unreported, together with the history of their reception by subsequent courts. The birth of AI-assisted legal research, which uses 'Retrieval-Augmented Generation' to improve the quality of LLM's output by allowing them to call on up-to-date primary and secondary legal sources, will make the process of accessing and synthesizing these authorities even more straightforward.⁸

One further effect of the more widespread use of LLMS may be to reify models of legal reasoning that are tractable and precise, and which allow for the processes by which we reason from legal sources through to conclusions to be represented as a computational process. Because the basic mechanism underlying LLMs is token-level inference, their ability to engage in complex problem-solving tasks is limited. Recent work suggests that the use of 'chain-of-thought' or 'tree-of-thought' prompting, which involves representing various problem-solving processes as a chain or tree of intermediate conclusions, can improve their problem-solving abilities.⁹ If LLMs are used to resolve legal

³ *Statesman*, 295b; *ibid* 466–467.

⁴ Walter J Ong, *Orality and Literacy* (3rd edn, Routledge 2013) 78.

⁵ Ong references the work of scholars like Milman Parry and Albert Law on the formulation of poetry in the oral tradition.

⁶ See e.g. T Ellis Lewis, 'The History of Judicial Precedent' (1932) 48 *Law Quarterly Review* 230.

⁷ WS Holdsworth, 'Case Law' (1934) 50 *Law Quarterly Review* 180, 182–183.

⁸ As with products like *Lexis+AI*, which has already been released in the US and is soon to be released in Australia: <<https://www.lexisnexis.com/en-us/products/lexis-plus-ai.page>> (accessed 30 April 2024).

⁹ See e.g. Jason Wei and others, 'Chain-of-Thought Prompting Elicits Reasoning in Large Language Models', *Proceedings of the 36th international conference on neural information processing systems* (Curran Associates Inc 2024); Shunyu Yao and others, 'Tree of Thoughts: Deliberate Problem Solving with Large Language Models' (2023) abs/2305.10601 ArXiv <<https://api.semanticscholar.org/CorpusID:258762525>>.

problems, their performance will be enhanced by modelling the various cognitive processes involved in legal reasoning. There is also an emerging recognition of the need for models used in automation to be explainable, in the broad sense that both the model and its outputs can be interpreted and explained to those affected by them. The Australian Human Rights Commission's Final Report on Human Rights and Technology included the recommendation that the Australian government should not make administrative decisions 'if the decision maker cannot generate reasons or a technical explanation for an affected person'.¹⁰ If the explanations given for automated decisions are to be socially appropriate and ethically sound, they will often need to be recognisable to their audience as involving sound legal reasoning or argumentation. Computational models of legal argumentation provide a promising basis for these explanations.¹¹

Lawyers do not need any high level of technical expertise to contribute to or understand computational models. In my own work I have explored the potential of so-called 'factor-based' models for the representation of legal reasoning and argumentation.¹² Factor-based models represent legal problems as collections of stereotypical patterns of fact ('factors' or 'dimensions') that strengthen or weaken the case for certain conclusions. A particular legal problem can be represented as a hierarchy of factors, with paths from the base-level to the top-level representing a particular chain of argument for or against a given conclusion (Figure 1).

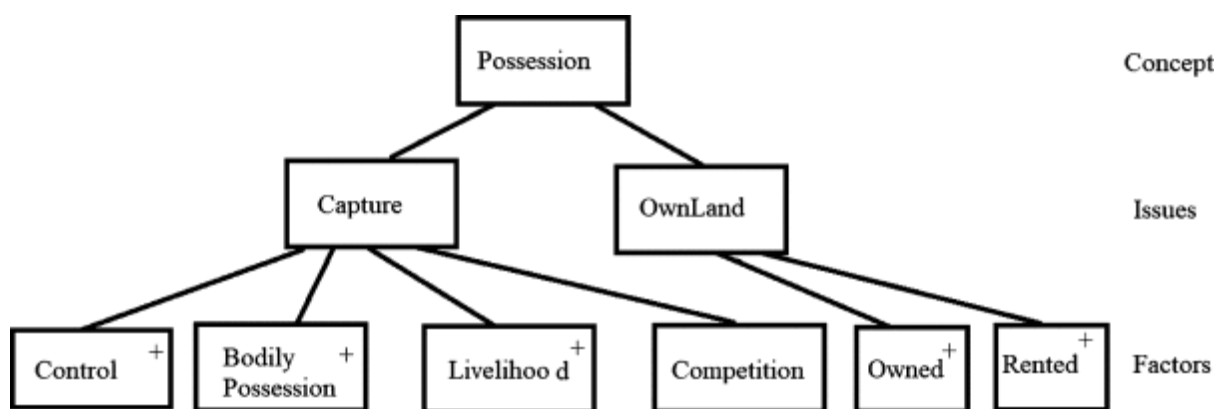


Figure 1: A simple factor hierarchy for a case involving a dispute over ownership of wild animals.¹³

These models have been used to teach legal reasoning to students with no technical expertise.¹⁴ And although most of the prominent computational models of legal reasoning emerged from earlier stages of AI and Law research, they have important precursors in legal scholarship. I think a similar concern for mapping the structure of legal argument can also be found in, for example, Hohfeld's canonical

¹⁰ Australian Human Rights Commission, *Human Rights and Technology*, Final Report (2021).

¹¹ See Katie Atkinson, Trevor Bench-Capon and Danushka Bollegala, 'Explanation in AI and Law: Past, Present and Future' (2020) 289 *Artificial Intelligence* 103387.

¹² See e.g. Robert Mullins, 'Two Factor-Based Models of Precedential Constraint: A Comparison and Proposal' (2023) 31 *Artificial Intelligence and Law* 703. For a non-technical introduction to factor-based and dimensional approaches see Kevin D Ashley, *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age* (Cambridge University Press 2017) ch 3.

¹³ Source: Atkinson, Bench-Capon and Bollegala (n 10).

¹⁴ Vincent Aleven, 'Teaching Case-Based Argumentation through a Model and Examples' (University of Pittsburgh 1997).

study of the informal logic legal rights, Wigmore's charts for the study of legal evidence (Figure 2), Montrose's obscure and difficult notation for the doctrine of precedent (Figure 3), or in Alf Ross's study of the structure of the legal concept of ownership.¹⁵ It is no coincidence, in fact, that work in AI has drawn on each of these authors in various ways, often quite consciously.¹⁶

The pattern which judgments may take is more easily discerned by the adoption of a notation which permits of a generalization of the various reasoning processes and may be applied to cases both of single judgment and of multiple judgments. The suggested notation is simple, viz.—

a+ signifies a finding that fact *a* exists.
 a= „ no finding whether fact *a* exists or not.
 a— „ a finding that fact *a* does not exist.
 x+ „ a determination that rule of law *x* exists.
 x= „ no determination whether rule of law *x* exists
 or not.
 x— „ a determination that rule of law *x* does not exist.
 R+ „ a decision that relief be granted to the plaintiff.
 R— „ a decision that relief be refused to the plaintiff.

Similarly, significations can be given for facts *b*, *c*, etc., by substituting *b* or *c* for *a*, and for rules of law *y*, *z*, etc. by substituting *y* or *z* for *x*.

Figure 1: Montrose's Precedent Notation.¹⁷

¹⁵ Wesley Newcomb Hohfeld, *Some Fundamental Legal Conceptions as Applied in Judicial Reasoning and Other Legal Essays* (Yale University Press 1919); Alf Ross, 'Tû-Tû' [1957] *Harvard Law Review* 812; JH Wigmore, *The Science of Judicial Proof: As given by Logic, Psychology, and General Experience, and Illustrated in Judicial Trials* (Little, Brown 1937); JL Montrose, 'The Language of, and a Notation for, the Doctrine of Precedent (Part II)' (1953) 2 *UW Austl. Ann. L. Rev.* 504.

¹⁶ Alf Ross's work has influenced Lindahl and Odelstad's study of intermediaries in normative systems; see e.g. Lars Lindahl and Jan Odelstad, 'Intermediaries and Intervenients in Normative Systems' (2008) 6 *Selected papers from the 8th International Workshop on Deontic Logic in Computer Science* 229. Although Montrose's work has not been particularly influential, it was an early precursor of the factor-based or dimensional approaches pioneered in AI and Law by Edwina Rissland and Kevin Ashley; see e.g. Kevin Ashley, *Modelling Legal Argument: Reasoning with Cases and Hypotheticals* (MIT Press 1990). Wigmore's evidence charts were one of the earliest attempts to represent arguments as tree diagrams, which has now become standard practice in the theory of argumentation, including in computational argumentation theory. See e.g. John L Pollock, *Cognitive Carpentry: A Blueprint for How to Build a Person* (MIT Press 1995). Hohfeld's informal logic of rights has influenced the study of 'normative positions' in AI and Law; Marek Sergot, 'Normative Positions' in Dov Gabbay and others (eds), *Handbook of Deontic Logic and Normative Systems* (College Publications 2013).

¹⁷ Source: Montrose (n 14).

Issue: Did U. Kill J.?
Prosecution's Case

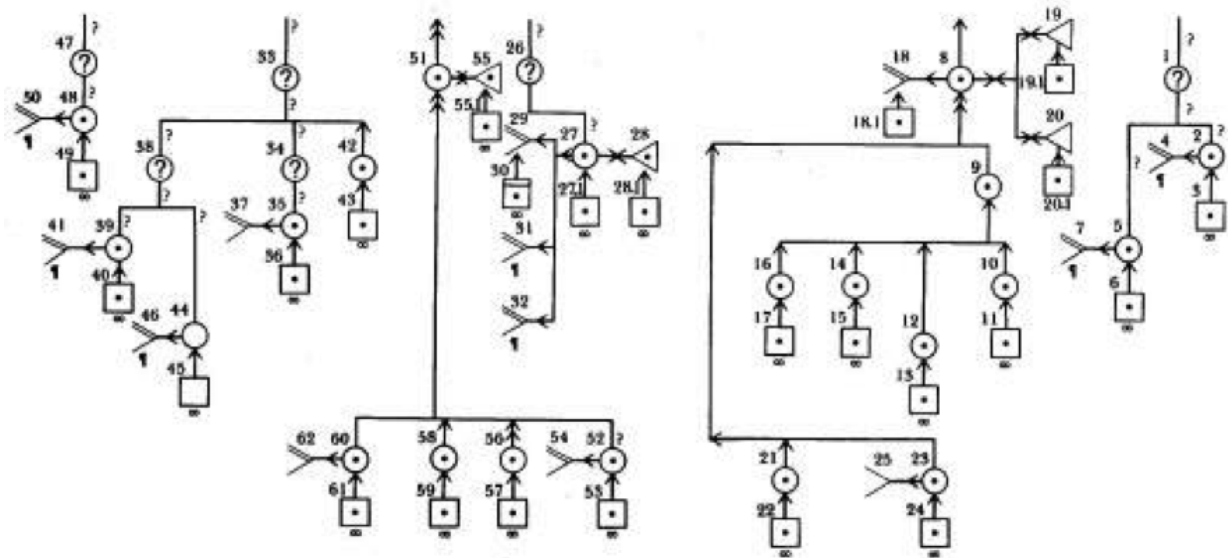


Figure 3: A Wigmore Chart.

Having reflected on how the use of AI might restructure legal thought, I want to conclude on a short note of Platonic caution, if not full blown Ludditism. I worry that we are at risk of losing track of the sense in which legal reasoning must be subordinated to social value. Consider the legal philosopher HLA Hart’s well-rehearsed example of a rule prohibiting vehicles from a public park. It is now relatively easy, if not trivial, for us to program the automatic classification of different objects as truck, cars, or bicycles, and for the classifier to group these together under the category of ‘vehicle’. But it is the social function of the law, and those who interpret the law, to decide whether a bicycle truly belongs in that category. As Hart noted, in a legal setting the act of classification is ‘consciously controlled by some identified social aim’.¹⁸ I think it is important that we retain our sense of conscious control over these acts of classification, and that we not mistake the simulation of human judgement in these matters for judgement itself.

¹⁸ HLA Hart, ‘Positivism and the Separation of Law and Morals’, *Essays on Jurisprudence and Philosophy* (Clarendon Press 1983) 67.