FOREWORD

THE HON CHIEF JUSTICE ANDREW BELL*

In chapter 1 of the recently published *Cambridge Handbook of Generative AI and the Law* ('*Handbook*'),[1] my review of which appears in this Issue of the *University of New South Wales Law Journal* ('*Journal*'), Tom Melham observes that:

> [t]echnical progress in generative artificial intelligence – at least in certain directions – has been both rapid and voluminous, and any account of the frontier would quickly be out of date.[2]

The use of the term 'frontier' is so apposite. When I was appointed Chief Justice of New South Wales in March 2022, I did not anticipate for one second that so much of my time less than four years later would be taken up with reading, thinking, and worrying about the advent of generative artificial intelligence ('GenAI'), which is something I now do almost every single day. Indeed, it was not until November 2022, just three years ago, that ChatGPT first emerged publicly and the world changed.

I do not think that this is a hyperbolic statement. Virtually no part of our society remains untouched by GenAI, and it is already having a profound impact throughout the world, attracting massive investment and requiring vast computing power and energy. And yet, so much about GenAI remains unclear and opaque and is potentially problematic, at least in the context of the administration of justice – and I am not just referring to the phenomenon of 'hallucinations' or, as Bret Walker AO SC recently put it in a public lecture, the propensity of GenAI to 'lie', albeit with a high degree of articulate self-assurance.[3]

So-called 'hallucinations' are, however, a major theme of this Issue, and indeed an important feature of any exposition of the relationship between GenAI and the law.[4] In chapter 24 of the *Handbook*, Martin Ebers cites a 2024 Stanford study showing the result of tests on GPT-3.5, Llama 2, and PaLM 2 with more than 200,000 queries to investigate the phenomenon of so-called hallucinations

---

* Chief Justice, New South Wales. I am grateful to my tipstaff, Sebastian Braham, for his assistance in the preparation of this foreword.

1 Mimi Zou et al (eds), *The Cambridge Handbook of Generative AI and the Law* (Cambridge University Press, 2025) <https://doi.org/10.1017/9781009492553>.

2 Tom Melham, 'Generative AI: An Introduction' in Mimi Zou et al (eds), *The Cambridge Handbook of Generative AI and the Law* (Cambridge University Press, 2025) 5 <https://doi.org/10.1017/9781009492553.004>.

3 Bret Walker, 'Lawyers and Money: 20 Years On' (Lecture, Banco Court, 14 October 2025).

4 See also A S Bell, 'Change at the Bar and the Great Challenge of Gen AI' (Speech, Address to the Australian Bar Association, 29 August 2025) [66]–[82], where I raised at least eight 'categories' of 'hallucinations' and discussed the cases in which they arose.

in typical legal tasks. According to the study, Ebers reports, 'hallucination rates in response to specific legal queries ranged from 69 to 88 percent for state-of-the-art language models'.[5] Similarly, in Vivi Tan, Jeannie Marie Paterson and Julian Webb's article in this Issue, on 'Generative Artificial Intelligence in Small Value Consumer Claims: Hallucination Risk, System Design, and Governance in Online Dispute Resolution',[6] a 2024 study by Matthew Dahl et al is cited which found that 'when asked a direct, verifiable question about a randomly selected Federal Court case, LLMs hallucinate between 58% (ChatGPT 4) and 88% (Llama 2) of the time'.[7] On any view, this should ring alarm bells.

One commonly ventilated view is that GenAI is particularly well suited to use by parties in 'small value consumer disputes',[8] as a means of improving access to justice by 'providing non-legal experts with accessible and timely information about legal claims'.[9] That is an attractive claim but it all depends on the quality and reliability of the output, which is in part at least a function of the skill and pre-existing knowledge of the user. Tan, Paterson and Webb's article points out that the use of GenAI by non-legally qualified parties comes with additional risks,[10] which, I would add, is wholly consistent with my experience dealing with self-represented litigants in the New South Wales Court of Appeal.[11] A similar point is made in Michael Legg, Vicki McNamara and Armin Alimardani's article 'The Promise and the Peril of the Use of Generative Artificial Intelligence in Litigation', in which self-represented litigants are said to be particularly vulnerable to the misuse of GenAI, owing to a lack of 'alternative sources of legal information' which makes the technology both (a) attractive and (b) difficult to verify: a concerning combination.[12] Local courts and administrative tribunals in which self-represented litigants most frequently appear are notoriously 'time poor' with huge volumes of cases. A significant additional burden is involved when such judicial officers and tribunal members need to be on the lookout for false cases and references.

Just as it deals with its weaknesses, this Issue also discusses the perceived *strengths* of GenAI, and indeed its first contribution, Lachlan Robb et al's '"It's the End of the World as We Know It": Generative Artificial Intelligence and the

5 Martin Ebers, 'LawGPT: LLMs under Legal Services Regulation' in Mimi Zou et al (eds), *The Cambridge Handbook of Generative AI and the Law* (Cambridge University Press, 2025) 431 <https://doi.org/10.1017/9781009492553.030>.

6 Vivi Tan, Jeannie Marie Paterson and Julian Webb, 'Generative Artificial Intelligence in Small Value Consumer Claims: Hallucination Risk, System Design, and Governance in Online Dispute Resolution' (2025) 48(4) *University of New South Wales Law Journal* 1257.

7 Ibid 1280, quoting Matthew Dahl et al, 'Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models' (2024) 16(1) *Journal of Legal Analysis* 64 <https://doi.org/10.1093/jla/laae003>. Curiously, Dahl et al found that LLMs performed most poorly on tasks related to material from lower court, or less recent, decisions: Dahl et al (n 7) 66, 80.

8 Tan, Paterson and Webb (n 6) 1283. This development can be observed in the Guangzhou Arbitration Commission in China, and also British Columbia's Civil Resolution Tribunal, an online alternative dispute resolution system which provides 'end-to-end' virtual solutions to small claims disputes.

9 Ibid 1266.

10 Ibid.

11 See, eg, *May v Costaras* [2025] NSWCA 178.

12 Michael Legg, Vicki McNamara and Armin Alimardani, 'The Promise and the Peril of the Use of Generative Artificial Intelligence in Litigation' (2025) 48(4) *University of New South Wales Law Journal* 1196, 1198.

Changing Landscape of Legal Practice and Education',[13] conducts an empirical analysis of some 51 practitioners across 27 organisations (including lawyers in large law firms, in-house, and in community legal centres) to convey the myriad ways in which GenAI is being used to enhance the efficiency of legal work, especially by completing some of the administrative grunt work typically reserved for graduate lawyers.[14] This may be the area where GenAI has the most to offer in the least problematic way.

Alongside the commonly ventilated 'draw cards' of efficiency and access to justice, Amy Wojciechowski and Daniela Simone's contribution to this Issue[15] raises the concept of our shared national interest in developing a 'sovereign Australian AI capability'.[16] Such a capability is said to be cultivated by offering a 'favourable location for all the stages of the AI supply chain'.[17] Other jurisdictions have taken that challenge on, not least of all Singapore, epitomised through its adoption of 'Project SEA-LION' (South-East Asian Languages in One Network), a family of LLMs tailored to Southeast Asian languages and with the agenda of 'nurtur[ing] Singapore as a global leader in AI solutions'.[18]

The rapid pace of private sector supply and demand has unsurprisingly and by and large outstripped regulatory responses and such regulatory responses as there have been vary significantly throughout the globe.

It is critical for those who seek to use and embrace GenAI to be properly informed and aware of its limitations as well as its claimed advantages, including the ramifications for its use and consideration of critical questions in relation to responsibility and accountability for errors in the outworkings of GenAI.

It is similarly critical for those who would seek to regulate the use of GenAI in the disparate areas in which it may be sought to be deployed to have as full as possible an understanding of these possibilities and limitations so as to ensure that regulation strikes an appropriate balance. It is difficult to regulate technology which is moving so quickly and constantly being refined and modified. But regulation there must be, and three articles in this timely Issue of the *Journal* address important questions of regulation.

A striking regulatory blind spot, exposed in José-Miguel Bello y Villarino et al's article, 'Are We Regulating the Right Digital Systems? Testing Emerging Artificial Intelligence Frameworks against Real-World Public Sector Systems',[19] is

13 Lachlan Robb et al, '"It's the End of the World as We Know It": Generative Artificial Intelligence and the Changing Landscape of Legal Practice and Education' (2025) 48(4) *University of New South Wales Law Journal* 1129.

14 Ibid 1149.

15 Amy Wojciechowski and Daniela Simone, 'Do You Copy?: Attributing Copyright Infringement to Actors Involved in Text-to-Image Generative Artificial Intelligence' (2025) 48(4) *University of New South Wales Law Journal* 1319.

16 Ibid 1327.

17 Ibid.

18 Jason Grant Allen and Jane Loo, 'Singapore's Evolving AI Governance Framework' in Mimi Zou et al (eds), *The Cambridge Handbook of Generative AI and the Law* (Cambridge University Press, 2025) 165 <https://doi.org/10.1017/9781009492553.014>.

19 José-Miguel Bello y Villarino et al, 'Are We Regulating the Right Digital Systems? Testing Emerging Artificial Intelligence Frameworks against Real-World Public Sector Systems' (2025) 48(4) *University of New South Wales Law Journal* 1165.

that emerging legal frameworks (such as the *EU AI Act*[20] or the New South Wales 'AI Assessment Framework'[21]) focus exclusively on artificial intelligence ('AI') systems, and overlook the risks posed by automated decision-making ('ADM') systems which do not utilise AI technology. With reference to Australia's Robodebt scandal and the Horizon Post Office scandal in the UK, the authors mount a compelling case that 'AI is a poor proxy for risks from technology adoption',[22] and that a general risk-based regulatory approach should be adopted without reference to specific technical criteria. That conclusion resonates as inherently intuitive, given that the speed of the development of GenAI (and reports that 'superintelligence' may be on the horizon)[23] will soon render redundant any regulatory regime which ties itself to rigid technological definitions of the systems it purports to regulate.

The theme of regulating not only AI, but all ADM systems, is carried over into Dancy's article, entitled 'Guiding Action in the Age of Artificial Intelligence'.[24] This absorbing piece addresses the regulation, or lack thereof, of predictive algorithms which are used as risk assessment tools, including, for example, the 'Level of Service Inventory – Revised' tool used in New South Wales, Victoria, and the Australian Capital Territory to 'infor[m] decisions about prison placement' and 'a range of therapeutic and other corrective interventions, and parole'.[25]

One fundamental vulnerability of predictive algorithms, and indeed all GenAI systems, is the so-called 'black box' dilemma, whereby their internal logic or reasoning is opaque, even to their designers. In light of that dilemma, which I have discussed in more detail elsewhere,[26] I embrace Dancy's conclusion that 'any judicial use of risk assessment should be such that the variables used to make predictions should be ascertained and communicated to those affected'.[27] Indeed in my view, at least at the current stage of the technology, the incipient use of GenAI more generally by some judges in some jurisdictions in the drafting of their decisions[28] raises large issues going to legitimacy, transparency, and trust in the judiciary.

---

20 *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence and Amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act)* [2024] OJ L 2024/1689.

21 NSW Government, 'The NSW Artificial Intelligence Assessment Framework' (Framework, 2024) <https://www.digital.nsw.gov.au/policy/artificial-intelligence/nsw-artificial-intelligence-assessment-framework>.

22 Bello y Villarino et al (n 19) 1167.

23 Jan Philipp Burgard, 'When Sam Altman Predicts a "Superintelligence" Might Arrive', *Politico* (Web Page, 25 September 2025) <https://www.politico.com/news/magazine/2025/09/25/sam-altman-ai-interview-axel-springer-00580997>.

24 Tatiana Dancy, 'Guiding Action in the Age of Artificial Intelligence' (2025) 48(4) *University of New South Wales Law Journal* 1236.

25 Ibid 1238.

26 A S Bell, 'Fabrication and Delegation: AI in International Arbitration' (Speech, International Arbitration Conference 2025, 13 October 2025) [71]–[72].

27 Dancy (n 24) 1256.

28 See, eg, in the United States, *Snell v United Specialty Ins Co*, 102 F 4th 1208 (11th Cir, 2024). In the United Kingdom, see Hibaq Farah, 'Court of Appeal Judge Praises "Jolly Useful" ChatGPT after Asking It for Legal Summary', *The Guardian* (online, 15 September 2023) <https://www.theguardian.com/technology/2023/sep/15/court-of-appeal-judge-praises-jolly-useful-chatgpt-after-asking-it-for-legal-summary>. In the

Another incident of the so-called 'black box' problem is its propensity, or perceived propensity, to mask the bias of GenAI outputs.[29] The theme of bias is thoroughly addressed in Sergio Sulmicelli's article on 'Queer-Responsive Regulation for Artificial Intelligence in Healthcare: A Comparative Study'.[30]

Sulmicelli uses the term 'queer-responsive' to refer to 'laws and policies that are more attentive and respond more effectively to the challenges of sex, gender and sexual minorities'.[31] That ideal is said to be threatened by AI's 'predictive and autonomous decision-making capabilities' which 'risk exacerbating the vulnerabilities of already marginalised groups, including women, transgender, intersex and non-binary individuals'.[32] One type of bias touched on is 'pre-existing bias',[33] which refers to the historical under-representation of certain groups in the datasets on which AI systems are trained, leading to outputs that 'appea[r] neutral but replicat[e] the initial inequality'.[34]

Part of the problem, as Sulmicelli explains, is that the 'rigid classifications' embedded in certain datasets, such as healthcare data and practices, make it difficult for GenAI models to adapt to new social standards or understandings, including in relation to gender or sex. That has obvious implications for the law, which relies on incremental change to befit changing social standards. Indeed, Sulmicelli's article reminded me of the observations of the Hon Michael Kirby, in a recent public seminar touching on the possibility of AI judges, where he remarked that the decision in *Mabo v Queensland [No 2]* (1992) 175 CLR 1 would not have eventuated through reliance on an artificial network trained on historical data.

In its final three articles, this Issue pivots to an increasingly salient and legitimate source of work for lawyers: the attribution of liability to a human actor lying behind a GenAI system, or the so-called pursuit for 'the natural person behind the AI'.[35]

The attribution of liability to human actors lying behind GenAI networks has both a technical and a normative aspect, each of which is addressed in these final three articles. The technical aspect is that it requires the disentanglement of AI operators along the design pipeline.[36] Wojciechowski and Simone's 'Do You

---

Netherlands, see Gelderland Court, 10664071 CV23-2321, 7 June 2024, ECLI:NL:RBGEL:2024:3636, reported in *V-N Vandaag* 2024/1668.

29 Bell (n 26) [64]–[78]. See also Dan M Kotliar, 'Data Orientalism: On the Algorithmic Construction of the Non-Western Other' (2020) 49(5) *Theory and Society* 919 <https://doi.org/10.1007/s11186-020-09404-2>.

30 Sergio Sulmicelli, 'Queer-Responsive Regulation for Artificial Intelligence in Healthcare: A Comparative Study' (2025) 48(4) *University of New South Wales Law Journal* 1288.

31 Ibid 1301.

32 Ibid 1289.

33 Ibid 1291.

34 Ibid.

35 Beatrice Panattoni, 'Generative AI and Criminal Guilt: When No One Meant to Harm' in Mimi Zou et al (eds), *The Cambridge Handbook of Generative AI and the Law* (Cambridge University Press, 2025) 400 <https://doi.org/10.1017/9781009492553.027>.

36 See Peter Henderson, 'Challenges for Foundation Model Liability and Regulatory Regimes: An Analysis of US Law' in Mimi Zou et al (eds), *The Cambridge Handbook of Generative AI and the Law* (Cambridge University Press, 2025) 123–4 <https://doi.org/10.1017/9781009492553.012>. See also Panattoni (n 35).

Copy?: Attributing Copyright Infringement to Actors Involved in Text-to-Image Generative Artificial Intelligence'[37] and Foong's 'Generative Artificial Intelligence Models and Copyright Infringement: Doctrinal Challenges and Regulatory Gap-Filling Using Unfair Competition Principles'[38] each go some way to disentangling the 'supply chain'[39] of GenAI models in the context of copyright claims.

The normative dimension relates to the extent to which established legal categories, including laws relating to data protection,[40] copyright,[41] privacy,[42] antidiscrimination,[43] competition,[44] and even equitable doctrines,[45] might be expanded to provide a cause of action for loss caused by GenAI systems, and, relatedly, the extent to which judges should facilitate and encourage that development, as a part of what Gageler CJ has described as 'the Dixonian ideal of a common law system [as being] one of incremental judicial development of legal principle to meet changing societal needs'.[46]

Wojciechowski and Simone's article,[47] together with Foong's,[48] grapple with that very question against the backdrop of the proliferation of copyright suits which have been brought against AI organisations in the United States, often on behalf of individual artists the work of whom has been incorporated, without their consent, by GenAI systems both at the training (input) and production (output) stages.

Both articles issue a powerful warning against 'knee-jerk reactions'[49] which seek convenient outcomes at the expense of legal coherence, by 'stretch[ing] copyright concepts' beyond the doctrine's normative goal of 'fostering a particular system of (human) public communication'.[50] Further, the authors point towards other, countervailing policy considerations weighing against the expansion of

37 Wojciechowski and Simone (n 15).
38 Cheryl Foong, 'Generative Artificial Intelligence Models and Copyright Infringement: Doctrinal Challenges and Regulatory Gap-Filling Using Unfair Competition Principles' (2025) 48(4) *University of New South Wales Law Journal* 1361.
39 Wojciechowski and Simone (n 15) 1321.
40 Hannah Ruschemeier, 'Generative AI and Data Protection' in Mimi Zou et al, *The Cambridge Handbook of Generative AI and the Law* (Cambridge University Press, 2025) ch 15 <https://doi.org/10.1017/9781009492553.020>.
41 Foong (n 38) 1364; Wojciechowski and Simone (n 15) 1326.
42 Elana Zeide, 'Generative AI and the Fundamental Limitations of US Privacy Law' in Mimi Zou et al (eds), *The Cambridge Handbook of Generative AI and the Law* (Cambridge University Press, 2025) ch 16 <https://doi.org/10.1017/9781009492553.021>.
43 Riccardo de Caria, 'Generative AI and Non-discrimination Law in the EU' in Mimi Zou et al (eds), *The Cambridge Handbook of Generative AI and the Law* (Cambridge University Press, 2025) ch 14 <https://doi.org/10.1017/9781009492553.019>.
44 Sylvia Papadopoulos, 'Redefining Rivalry: Generative AI and the Evolving Landscape of Competition Law' in Mimi Zou et al (eds), *The Cambridge Handbook of Generative AI and the Law* (Cambridge University Press, 2025) ch 19 <https://doi.org/10.1017/9781009492553.024>.
45 May Fong Cheong and Mimi Zou, 'Resolving Wrongs in Algorithmic Contracting: Applications of the Doctrines of Unconscionability' (2025) 48(4) *University of New South Wales Law Journal* 1399.
46 Stephen Gageler, 'The Coming of Age of Australian Law' in Barbara McDonald, Ben Chen and Jeffrey Gordon (eds), *Dynamic and Principled: The Influence of Sir Anthony Mason* (The Federation Press, 2022) 8.
47 Wojciechowski and Simone (n 15).
48 Foong (n 38).
49 Ibid 1397; Wojciechowski and Simone (n 15) 1358.
50 Wojciechowski and Simone (n 15) 1357.

copyright law, including, to pick up a concept discussed in some of the earlier articles, that a lack of training data 'contribute[s] to embedding harmful biases in models that risks broader harms in Australian society in some cases'.[51]

Finally, May Fong Cheong and Mimi Zou's article 'Resolving Wrongs in Algorithmic Contracting: Applications of the Doctrines of Unconscionability', argues that the statutory and equitable unconscionability doctrines should be reassessed and expanded to regulate disputes arising out of algorithmic-based contracting. That discussion arises in response to the Singapore Court of Appeal's decision in *Quoine Pte Ltd v B2C2 Ltd* [2020] 2 SLR 20 ('*Quoine*'), where the Court upheld the validity of algorithmically executed contracts in the form of cryptocurrency trades which were executed at approximately 250 times the prevailing market rate due to a system error. The authors argue that the need for contract law doctrines to 'remain relevant and responsive to modern contracting systems that increasingly employ machine learning and other advanced technologies' justifies the widening of the doctrines.[52]

Ultimately, as more cases such as *Quoine* arise, regulators, lawmakers, and, where appropriate, judges, will need to determine: (i) how established legal causes of action and principles accommodate the myriad issues that will be presented by the pervasive and ubiquitous adoption of GenAI technology in all fields of human endeavour; (ii) which issues require legislative intervention; and (iii) which ought not be regulated at all. That is a difficult and multi-faceted exercise, behind which lies many of the different policy considerations canvassed in this Issue, not least of all the push and pull of international regulatory competition, as Australia establishes its sovereign AI capability.

The Editor and all the contributors are to be commended for this constructive contribution to what is now a substantial academic field dealing with the relationship between AI and the law.

51 Ibid 1325.
52 Cheong and Zou (n 45) 1426.

THEMATIC
ARTIFICIAL INTELLIGENCE AND THE LAW

Artwork by Bronwyn Verwey