



Generative Artificial Intelligence in Legal Practice: Epistemic Expertise, Unethical AI, and the Role of Virtue

Julian Webb¹

Received: 28 May 2025 / Accepted: 18 February 2026
© The Author(s) 2026

Abstract

This paper discusses the impact of generative AI (genAI) on the legal profession and the importance of virtue ethics in addressing the epistemic and ethical challenges for legal practice posed by genAI technology. It traces the development of legal genAI, and argues that its functionality currently leaves it ethically lacking (relative to the dream of AI as a substitutive technology) in terms of its competence, integrity, and compliance with ethical standards. The paper further argues that genAI also fails to contribute to the more normative, non-principlist, ethical expectations that we associate with excellent practice. While it can perform well on certain well-defined tasks, it struggles with the nuances of real-world legal problem-solving, particularly complex relational problems where practical wisdom and contextual understanding play a significant role, over and above technical domain expertise. This is because the tools currently lack the capacity for moral discernment and the ability to autonomously reason either from first principles, or from virtue-based reasoning, both of which are significant for professional decision-making. It suggests that, without more, the deep deployment of genAI in legal practice may thus lead to a diminution of practical wisdom, rather than its augmentation or enhancement. The paper concludes with some recommendations on how the legal ecosystem might manage genAI adoption in ways that better enable AI tools and the lawyers using them to augment rather than diminish virtue.

Keywords Legal technology · Generative AI · Virtue · Practical wisdom · Legal ethics

✉ Julian Webb
julian.webb@unimelb.edu.au

¹ Melbourne Law School, University of Melbourne, Melbourne, VIC 3010, Australia

1 Introduction

Professionals, it has been said, constitute ‘the preeminent institutional agents of our time’ (Scott, 2008, p.219). They have the power to shape intellectual and practical fields, and thence society itself, protecting or undermining different stakeholder interests, and, oftentimes, advancing their own as well. The influence of AI on the agentic work of professionals is thus a matter of substantial public interest, with writers variously proclaiming the liberation of professionals from the drudgery of routine, or, amongst somewhat contradictory grander narratives, a future of digital access to justice, liberated from professional monopoly; the death of conventional law practice, or even of traditional status professionalism as we know it (eg, McGinnis & Pearce, 2014; Webley et al., 2019; Susskind, 2008; Susskind & Susskind, 2022)..

This paper tries to take a more modest approach, blending descriptive and normative insights in considering the ethical consequences of ‘legal AI’ and specifically generative AI (genAI) on the practice of law.¹ It argues that moves to AI-assisted practice create new epistemic and ethical challenges for legal professional practice which heighten the importance of virtue ethics as part of the practitioner’s ethical toolkit. It develops this analysis through five substantive sections. It first provides, by way of context, a brief overview of AI developments in legal practice, before turning to the current state of genAI adoption by the legal industry, highlighting the increasing use of these tools for tasks such as process automation, document drafting, and legal research. In parts 3 and 4, its focus shifts to the role of ethics as part of the AI-assisted lawyer’s ‘epistemic expertise’. It examines, first, how genAI creates new challenges for professional competence, truthfulness, and compliance with professional standards, and then how the intermediation of genAI impacts the epistemic and ethical work of the lawyer, and particularly the lawyer’s capacity to offer ‘wise counsel’. It concludes by highlighting ways in which the ethical risks of genAI might be mediated by proper attention to the significance of virtue in legal practice.

2 AI in Legal Practice: The State of Play

The pursuit of an autonomic ‘law machine’ (Mehl, 1959) has been a goal of AI research from almost the inception of the field. Computer scientists – and some legal theorists—found law’s apparently formalistic rules-based structure, somewhat misleadingly, familiar and suggestive of a form of deductive reasoning that might readily be automated with the goal of providing either decision-support or substitutive tools in proving legal advice (Susskind, 1986; cp MacCarthy, 1990). The parallel rise in

¹ For the purposes of this paper, ‘legal AI’ is used to distinguish the use of AI by lawyers, or paralegals under the supervision of lawyers, to support or undertake legal knowledge-based tasks in practice settings. This encompasses a broader range of tasks than the direct provision of legal information and advice to a client. So legal document assembly or AI-assisted document review are within the purview of legal AI for this purpose, but not the use of AI to provide more general ‘back office’ services (such as business communications, hiring and supervision of staff, cybersecurity, automated billing, etc.). While some legal AI tools could be used to assist more directly with the adjudicative process, the ethics of ‘robo-judging’ fall outside the scope of this paper.

‘jurimetrics’ (Baade, 1963; Loevinger, 1949) as the pursuit of a modern (mathematical and behavioural) science of law also helped to create cross-disciplinary interest in what became the emerging field of legal informatics (see Paliwala, 2010). Early developments in legal informatics were, however, dominated by work on information search and retrieval (Webb, 2022) and there was little practical progress in legal AI until the emergence, first, of a wave of knowledge-based or ‘expert’ system research in the 1970s (see, e.g., McCarthy, 1977, 1990; Greenleaf et al., 1995; Francesconi, 2022) and, secondly, more successfully, by the development, since the early 2000s of data-led approaches to legal AI, characterised by modern deep learning and natural language (NLP) approaches (de Maat et al., 2010; Francesconi, 2022).

Law firms and legal tech suppliers have been using these increasingly powerful machine learning models to enhance and add value in existing legal tech use cases, such as legal research, document drafting and review. Generative AI (genAI) is the latest development within this space, whereby LLMs trained on very large datasets, comprising hundreds of millions or even billions of data points, are combined with a chatbot interface to generate novel content in response to natural language user prompts. Advanced neural network transformer models (GPTs) give genAI tools greater (semantic) contextual sensitivity in ‘understanding’ and applying prompts to its dataset. This is integral to operation in domains like law where meaning is heavily context-dependent. The other distinctive feature of GPTs is that they operate on their data in ways that are generative not just discriminative—put simply, they synthesise, rather than distinguish categories and select. This is because the algorithm is programmed so that the output is a product of the function, minimally, $p(x)$, or, maximally, $p(x,y)$, where p is the probability of the next word, x is the totality of data inputs, and y is the totality of prior outputs. It is this generative capacity that enables genAI to construct new or synthesised content from its dataset and stops it from being simply another search and reporting tool. At the same time, this is also the feature which explains genAI’s tendency to fabricate or ‘hallucinate’ outputs, a problem discussed in the next section.

The evolution of legal genAI has constituted a major leap forward in legal automation, certainly relative to the limited impact of earlier expert system technologies. The very significant improvements in Uniform Bar Exam performance by OpenAI’s GPTs from version text-davinci-001 to GPT-4 (Katz et al., 2024) shows that genAI can at least hit a baseline standard of legal competence, and the rapid rate of improvement in system performance across these versions may well indicate that we have not yet reached peak capability, though the sheer scale and costs of compute may yet become a barrier.

This growth is reflected in the development of new tools and use cases for the technology, and in increasing take-up (Webb & Paterson, 2025). Initially the focus had been on low risk activities such as simple document drafting (especially for practice support, eg web content, marketing, triage forms, etc.), and use as a first draft tool for legal documents, document summarisation, and process automation, particularly high-volume, lower value diagnostic legal tasks (Faulconbridge, Sarwar & Spring, 2024; Faulconbridge, 2025). Established ML-based contract and document review

tools, like Luminance AI,² have thus upgraded to exploit transformer technology. Higher risk uses, such as AI-assisted legal research and legal assistance are starting to become more mainstream, as large commercial research providers like Lexis and Westlaw, which have dominated the legal research and information retrieval space since the 1980s, have started to make use of transformer models.³ The availability of genAI chatbots, like Open AI's ChatGPT, Google Gemini, or Anthropic's Claude, potentially extends the reach of automation into smaller law firms and (in theory, at least) allows the public to access seemingly functional legal assistance. Inevitably there is significant commercial value for both tech suppliers and possibly law firms too in moving genAI up the value and risk chain into the development of tools that reliably augment the core legal advisory function.

Notwithstanding the 'transformational promise' (Armour & Sako, 2020) of AI, so far, these tools have had only a limited substitutive impact on legal labour. This reflects both the functional reality that the tools support a fairly limited range of legal tasks, and the extent to which reliability and capability concerns, discussed in the next section, have meant that regulatory standards and guidance (where they exist) currently place a premium on lawyer oversight as a risk management strategy (see, eg, Law Society of New South Wales et al., 2024; Solicitors Regulation Authority, 2022).

At the same time, significant cultural challenges remain for genAI. There is still a strong sense of caution, perhaps even resistance, around genAI adoption (eg, Centre for Legal Innovation, 2024; Caserta, 2022; Faulconbridge, Sarwar & Spring, 2023), reflecting specific doubts about its reliability, as well as concerns about confidentiality and privilege, privacy, data security and data governance, particularly with respect to generic AI tools. Moreover, traditional partnership structures and business models have also been seen as both a financial and cultural constraint on innovation (Caserta, 2022; Hodgkinson et al., 2023), particularly on the development of more effective multi-disciplinary business forms (Armour, Parnham & Sako, 2022). How easily these barriers may be overcome is an open question. There are significant economic drivers that weigh in support of AI adoption, in terms of both efficiency gains and possibly cost reductions once services are heavily automated. AI use is also seen by some as a critical next step in increasing access to justice (Susskind & Susskind, 2022). If the benefits are seen to justify or outweigh the risks, then the technology itself may become transformative of both professional activity and business form (Armour & Sako, 2020; Faulconbridge, Sarwar & Spring, 2023) and may even be seen to 'enact'⁴ role and expertise within organisations (Glaser et al., 2021, pp.9–10). This raises critical questions for practice, regarding the extent to which genAI itself

² See <https://www.luminance.com/>

³ See, eg, Lexis+AI at <http://lexisnexis.com/en-au/products/lexis-plus-ai>

⁴ That is, to operate as part of a complex socio-material assemblage of technology and human that goes beyond merely operationalising specific tasks or computations, though it should be noted here that Glaser et al. seem to have in mind a more narrowly functional conception of 'role' than we find in much of the professional ethics literature.

is agentic,⁵ and how such AI may, in turn, impact the ethical responsibilities and behaviours of lawyers.

3 Legal genAI: Agency Without Domain Competence, Truth, or Ethics?

The problems with genAI for law are not unique to the legal domain, but they take on greater salience insofar as legal advisory work can be seen as a relatively high stakes application for the tools, and one which adds to longer standing concerns about technology's deprofessionalising effects (eg, Remus & Levy, 2017; Webley et al., 2019; Thornton, 2021). At its most fundamental, the question here is how much agency should we give over to the AI in practice?

GenAI is already capable of having agentic effects; indeed it can, in a philosophical sense, be regarded as agentic in and of itself, even if not a fully human (moral) agent.⁶ This raises some important, ethical, considerations. As Floridi (2023, p.5) concludes, with the emergence of LLMs and genAI, we have potentially decoupled *effective* agency (or as he puts it, 'the ability to act successfully') from human intelligence.⁷ The two closely related questions for this section might thus be framed as, firstly, what does 'successful' acting look like in the legal domain, and, secondly, can genAI meet these criteria? For the purposes of this paper, it is proposed that we can answer the first question succinctly rather than exhaustively: 'successful' autonomous professional action requires three core qualities: domain competence, integrity, and an ability to advise consistently with professional conduct norms and values. We turn now to the question whether genAI can meet those criteria.

3.1 Domain Competence

Legal competence is a *sine qua non* of the advisory role and a central, consumer protection, justification behind the reservation of at least some legal acts to those with professional title. Competence itself is, of course, a complex and somewhat contested concept. Nonetheless, there is a long-standing recognition that professional competence needs to be understood as broad-based and highly contextual. At the risk of stating the obvious, it is widely accepted that competence goes well beyond some mastery of domain knowledge. Thus, in constructing a relatively early but internationally influential conception, the American Bar Association's *MacCrate Report* (1992) sought to define legal competence (beyond abstract legal knowledge) in terms

⁵As is apparent from Part 3, following, when this paper talks about AI as 'agentic', it is concerned with the ways in which any AI has agency in the real world, rather than 'agentic AI' in the narrow sense of autonomous AI agents based on large action models (LAMs).

⁶For purposes of the argument following, this paper accepts that genAI exhibits three qualities necessary to attribute agency to the tool: interactivity (ie responsiveness to stimuli), functional autonomy, and adaptivity (ability to learn) (Floridi, 2021; Floridi & Sanders, 2004).

⁷The larger question of whether we need to adapt our conception of *moral* agency to accommodate forms of artificial agency that are mind-less and not responsible (see Floridi, 2021), or whether allocating moral agency without (human) responsibility remains a category mistake lies beyond the scope of this paper.

of ten fundamental lawyering skills. These included skills of legal and fact analysis, research, communicating with clients and others, negotiating skills, managing workload, and recognising and resolving ethical dilemmas.

Does genAI demonstrate legal competence across this kind of breadth? Of course not.⁸ GenAI, particularly where it has had task-adaptive calibration and fine tuning,⁹ can be at least as effective as human agents at completing the discrete (diagnostic) tasks for which it has been trained. It *can* retrieve with a high degree of accuracy legal information held within its corpus, though whether it can do so consistently is another matter. It thus has the potential to augment human decision-making through these functions, but its substitutive capability remains limited (Faulconbridge, Sarwar & Spring, 2024). Moreover, functional analyses of what lawyers do and how they spend their time (Remus & Levy, 2017; Webb & Paterson, 2025) suggest that the kind of diagnostic work at which AI currently excels is in fact a relatively small proportion of day-to-day legal work, and disproportionately supports or substitutes for the labour of junior lawyers and paralegals. The inability of genAI tools to deliver the higher level and higher-value cognitive and affective capabilities or ‘soft skills’ required to tailor advice to the needs of a specific client, at least limits its current substitutional capability (see, eg, Lassébie & Quintini, 2022; Cantatore, Atwill and Field, 2022 p.7.). Longer term, we cannot rule out that the nature and scope of legal competence will be radically changed by the technology, but for the present five competence limitations are emerging.

First, hallucination appears to be an ongoing issue in practice and in testing. Hallucinations have been identified as the cause of several widely publicised failures, where lawyers or lay people have relied on genAI fabrications in proceedings, to their own detriment and at increased cost to the public resource that is the court and tribunal system.¹⁰ Hallucinations in the legal setting have been shown to fall into three kinds (Magesh et al., 2024). First, an AI can misrepresent its training data; secondly it can be ‘unfaithful’ to its prompted input, or, thirdly, it can misrepresent the ‘true facts’. This last category subdivides into two: what the authors describe as ‘correct-

⁸This statement needs to be subjected to methodological caveats. Inevitably, given the novelty of the technology, studies of genAI’s performance of legal tasks are still in their infancy. Most studies have examined the performance of generic platforms, like ChatGPT, rather than specialist legal genAI. The majority have not yet progressed beyond the pre-print stage; there is no standardised testing protocol, or common benchmark (though note ongoing work on benchmarking legal tools – eg Guha et al., 2023), so that a variety of methodologies are being applied, and this, combined with the rapid upgrading of models, has meant that there are no replication studies to speak of. Moreover, given the pace of change, any conclusions are subject to significant temporal caveats.

⁹Fine tuning uses a specialist dataset to build on the foundation model’s (general) knowledge acquired during pretraining, thereby improving the model’s performance on task-specific data. Fine tuning commonly entails (human) supervised learning, which has been widely used in legal contexts, for example, in training pre-GPT document review tools, though reinforcement learning, and forms of self-supervised learning are increasingly being developed; see eg Ohm, 2024.

¹⁰For examples of cases where lawyers have been sanctioned in court proceedings for presenting non-existent legal authority created by genAI, see *Zhang v Chen* [2024] BCSC 285; *Mata v Avianca Inc* 678 F Supp 3d 443 (SD NY, 2023). Likewise, see, eg, *Harber v Commissioners for His Majesty’s Revenue and Customs* [2023] UKFTT 1007 (TC); *Luck v Secretary, Services Australia* [2025] FCAFC 26, and *May v Costaras* [2025] NSWCA 178, where applications by litigants in person have also been rejected for relying on genAI-fabricated material.

ness' and 'groundedness' hallucinations (Magesh et al., 2024, pp.7–8). The former represent situations where the AI misleads by material inaccuracy, eg, it offers a false proposition of law, or denies a correct one, or identifies real but irrelevant material, whereas the latter may be more subtle and difficult to identify, insofar as the system may correctly represent a legal principle but misrepresent its source, or misrepresent what a case stands for. Fine tuning and RAG¹¹ techniques have undoubtedly improved the accuracy of specialist generative models, but recent research, using the same methodology as Dahl et al (2024), has still indicated error rates ranging from around 17%–34% (Magesh et al., 2024).¹² Not surprisingly, indications are that, notwithstanding techniques like RAG, commercial suppliers are having to invest significant resource into ongoing testing, fine tuning and recalibration to prevent performance degradation in their models.

Secondly, knowledge outputs may be significantly incomplete. Length and completeness of outputs appears to be a notable variable between systems, and there is also some evidence that longer outputs, while potentially providing greater depth, are possibly more prone to reliability issues (Peoples, 2025; Magesh et al., 2024; Tan et al., 2023).

Thirdly, genAI seems to struggle to recreate accurately the relationships between bits of information extrapolated from its dataset. This would appear to account for a proportion of its correctness hallucinations, ie, it is not just randomly making things up, but misrepresenting, for stochastic reasons, the true facts to be derived from its training data. Thus, for example, where ChatGPT-3 confuses and conflates the textual provisions in ss. 152(d)(2)(A) and 152(d)(2)(B) of the US tax code (Blair-Stanek et al., 2023) or wrongly states that Justice Ruth Bader Ginsburg's judgment in *Obergefell v Hodges* (2015) 576 US 644 was in dissent (Romoser, 2023).

Fourthly, and relatedly, while the more recent iterations of genAI can, perhaps, come close to or match human performance on relatively well-defined knowledge-elicitation and application tasks, such as Bar Exam questions (Katz et al., 2024; cp Martinez, 2025), where there is likely a substantial corpus of public and licensed information, their ability to deal with the nuances and messiness of real-world legal

¹¹ Retrieval augmented generation (RAG) increases search accuracy by taking user prompts and using these to run lexical and semantic search through a linked external expert knowledge base, retrieving and then running retrieved documents through the transformer model before providing a response. RAG also forms the basis for the next likely phase of legal AI development, since it potentially enables the integration of law firms' own knowledge banks with the commercial dataset, creating a more focussed and bespoke user experience. Problems have been found with domain-specific as well as generic models. In 2024 an Australian lawyer was referred to their regulator for relying on hallucinated filings generated by LawY, a RAG-augmented AI. In this case the AI service provider had also provided a human verification service, which sent corrected information to the lawyer ahead of the court hearing, but, for whatever reason, the uncorrected case list was presented to the court. See Josh Taylor, 'Melbourne lawyer referred to complaints body after AI generated made-up case citations in family court' *The Guardian*, 10 October 2024. Lawyers for the large US personal injury firm Morgan & Morgan have been sanctioned by the US Federal District Court for including fictitious case citations generated by an in-house proprietary AI. It is not clear from the filings what foundation model this was built upon or whether it was RAG-augmented: *Wadsworth v Walmart Inc.* 348 FRD 489 (D Wyo, 2025).

¹² One limitation of Dahl et al.'s methodology should be noted. It can be argued that it underutilises prompt engineering techniques in the tasks it measures, and to that extent may over-estimate the hallucination rate, relative to the outputs of a more prompt-engineered approach (cp Schwartz & Choi, 2023).

problem-solving remains more questionable. This reflects weaknesses in the systems' capacity for both rule comprehension, and fact analysis – both fundamental legal skills, particularly in relation to the more bespoke, high-level end of legal advisory work.

In terms of rule comprehension, there are at least two issues. The first is semantic. Studies suggest that the very formalistic structure of statutory rules creates challenges of comprehension that limit the AI's accuracy (Blair-Stanek et al., 2023; Peoples, 2025). The second is, again, more subtle, which is the extent to which AI tools can capture complex conceptual relationships between rules and rule-types. Critical here may be a system's ability to comprehend the 'meta' norms of legal analysis, what Hart (2012) famously calls the 'secondary rules' which enable the determination, operation and change of (primary) rules of substantive law in a legal system. For example, principles of precedent and rules of statutory construction thus help lawyers assess the relative weightings that might be applied to different legal sources, or set parameters around human interpretative discretion, but their application is often opaque and left implicit in the ways that judges and other lawyers talk about legal rules. It remains unlikely that, by themselves, genAI models can readily surface and apply this kind of tacit or hidden knowledge.

In terms of fact analysis, studies suggest that general performance seems to decline with increased real-world factual complexity (Bhambharia et al., 2024; Choi et al., 2022). More specifically, Niblett and Yoon (2024) make the point that the kind of factual specificity raised by real disputes tends to exist within a realm of private information that typically falls outside the (public) corpus of genAI tools, and this in itself limits the AI's utility in resolving factual disagreement. These studies suggest that genAI would particularly struggle with both novel and heavily fact-contingent scenarios. For example, in maintaining the high-tech theme, we might speculate over how genAI would assess whether a potentially novel duty of care in negligence could be owed by a government department to complainants harmed by its use of an automated decision-making system, or how it would evaluate a clinician's liability in negligence for relying on a specific diagnostic AI tool that has mis-prescribed certain drugs to their patient (cp Pricor, 2023).

Finally, ideas of competency also imply at least some knowledge of one's own limitations. Yet, as Dahl et al., (2024, pp.82–84) observe generic genAIs 'systematically overestimate' their own reliability, and may uncritically accept prompters' legal assumptions. Moreover, because of the 'black box' (Pasquale, 2016) nature of genAI's performance, interpreting why a particular LLM hallucinates on a particular task, or why different LLMs perform differently on the same task is something of an open problem (Ji et al., 2023; Zou et al., 2023). It also means that accuracy and reliability risks tend to require both machine *and* human mitigation.

3.2 Truth

Ethical standards place a premium on the lawyer's truthfulness and integrity. Lawyers are thus commonly under express obligations not to mislead the court,¹³ to advance claims only for which there are 'proper' legal grounds,¹⁴ and to communicate with other lawyers and third parties truthfully,¹⁵ and in good faith.¹⁶

These norms are significant when we come to consider the role of genAI. We have seen already that it has the capacity to mislead in a way that places its use in tension with these professional expectations. This is not because the technology is in any intentional sense a bad faith actor. It is not actively designed to lie or misdirect, but its capacity for truthfulness is heavily contingent on the truth-quality of its training data, and on the extent to which it is further aligned and trained so as to limit bias, hallucinations and misleading representations. GenAI models do not 'understand' their outputs and have no grounding conception of truth on which to assess their own outputs. To this extent, genAI thus fails Bernard Williams' test for mature moral agency, which is a commitment to truthfulness (see Thomas, 2024, p.2): taking commitment in its ordinary sense, it would be fallacious to talk of an AI having a commitment to something of which it has no real conception.

Moreover, notwithstanding any 'intent' regarding truthfulness, genAI design is still problematic insofar as it has features that could encourage a misleading over-reliance on the tools and their capabilities.

GenAI needs to be understood not so much as a neutral tool, but as an 'engagement machine'. The massive development and running costs of the technology means that genAI must rapidly become a ubiquitous (in a social and market sense), high volume, high use, technology. Its branding as a 'co-pilot', and bundling with other essential work-tools like browsers and word-processing apps is consistent with this process of creating ubiquity. Ubiquity requires user engagement. Employee engagement with AI is itself a complex and still relatively little researched phenomenon. It is likely shaped, for example, by a mix of organisational messaging (both formal and informal) about safety, security and appropriateness of tool use; users' perceptions of task-technology fit; individual's organisational roles, their personal attitudes to technological innovation, and attributes like age, gender and career stage. Engagement cannot be reduced to marketing, but that, and related design features, are important. In other words, the extent to which GenAI tools are themselves built to deliver a positive and engaging user experience, and the (ethical) risk that this in itself could create are thus relevant considerations.

While this requires further empirical testing (cp Du, 2025; Palminteri & Lebreton, 2022), the hypothesis advanced here is that AI tools are built in ways that encourage both authority and over-confidence biases in their use. This is a consequence, it

¹³ See, eg, Rules 3 and 17 *Australian Solicitors Conduct Rules* (2024); Rule 1.4, *SRA Code of Conduct for Solicitors, RELs, RFLs and RSLs* (2019); Rule 3.3, *American Bar Association (ABA) Model Rules of Professional Conduct* (2025).

¹⁴ Rule 2.4, *SRA Code of Conduct for Solicitors*; Rule 3.4, *ABA Model Rules of Professional Conduct*.

¹⁵ Rule 4.1, *ABA Model Rules of Professional Conduct*.

¹⁶ *Legal Services Commissioner v Mullins* [2006] LPT 012 (Qld).

is suggested, of the interplay between the features that make them both useful and attractive engagement machines, and their insufficient attention to the problem of truthfulness. Four particular design techniques are illustrative of the risk (Du, 2025; Shi et al., 2023; van Es & Nguyen, 2024).

Personalisation: the fact that genAI can welcome me back with a ‘Good Morning, <name>’ and a list of my recent ‘chats’, is a simple example of this feature, but it extends more deeply into the tools’ potential to analyse user data and preferences to create highly personalised content and interactions, increasing our sense of their utility and relevance to us.

Intuitive interface design: the simplicity of prompt interfaces, which also makes these tools largely indistinguishable from conventional search engines, their facility for natural language interaction, and the absence of (visible) user guides also adds to usability and engagement, but these features, particularly the human-like, comprehensible yet authoritative tone of the chat, combined with the way in which hallucinations can be buried within otherwise accurate and informative responses, add to the reliability risks.

Dynamic content creation: unlike earlier tools, genAI can easily create different versions of content based on the same or similar user prompts; its capacity to develop outputs quickly and iteratively, in response to successive, refining prompts likewise highlights its dynamic, adaptive quality. This obviously has significant utility and to that extent can be said to ‘fairly’ build engagement, but it is also here that the relative lack of truthfulness guardrails comes into play, insofar as the accuracy of response can itself be significantly shaped by the assumptions, style and direction of prompting (Du, 2025).

Lastly, and relatedly, *user-centricity*, what we might colloquially think of as a desire to please, can also be built into the AI, particularly through instruction tuning that disinclines the AI from engaging in debate or breaching perceived neutrality standards (Du, 2025).

These various biasing effects may be exacerbated where users are unskilled and under-informed about the risks of genAI, and/or under significant work pressure, and hence more likely to miss errors. The fact that there is generally limited explainability capability built into these tools to assist us in finding such errors adds to the consequential risks of bias and hallucination (see Schneider, 2024).

3.3 Ethics

A sense of ethical obligation is central to the professional ideal and this tends to be reflected in the language of the profession’s deontology. Australian solicitors in most jurisdictions are thus expressly reminded that complying with their own ethical rules is itself a ‘fundamental ethical duty’¹⁷; professional integrity likewise is said to connote and require ‘adherence to the ethical standards of the profession’ (Wingate

¹⁷ Under rule 4.1.5, *Australian Solicitors Conduct Rules*, adopted in all jurisdictions except the Northern Territory – see, eg, rule 4.1.5, *Legal Profession Uniform Law Australian Solicitors Conduct Rules 2015* (NSW, Vic), and *Legal Profession Uniform Law Australian Solicitors Conduct Rules 2022* (WA); rule 8(e), *Legal Profession (Solicitors’ Conduct) Rules 2020* (Tas).

& *Evans v Solicitors Regulation Authority* [2018] EWCA Civ 366, [100]) and is a pre-condition of the right to practice. To this extent, an inability to ‘do’ professional ethics can be seen as a major disqualification from the role.

This is a problem for law insofar as the academic literature to date adheres to the view that genAI cannot do legal ethics in any meaningful sense (see, eg, Remus & Levy, 2017; Wendell, 2019). The arguments disclose concerns that are relatedly conceptual, practical, and institutional. Together these suggest that there is something distinctive about not just the profession’s ethical rules, but its ethical *role*, and (thence) something ‘irreducibly human’ (Clark, 2018) about ethical lawyering.

The distinctiveness of professional rules as a moral code has long been acknowledged. Legal ethics is not just another branch of law, and its practices cannot be constrained by the kind of rule application theory that commonly shapes interpretation in substantive law fields. This is a problem, therefore, about both the form and substance of professional ethics. Professional rules are seldom ‘legislative commands’ (Law Council of Australia, 2024, p.8), being generally drafted as broad, often principles-based, statements requiring contextually sensitive interpretation against prevailing professional and social norms. Substantively too, they are frequently acknowledged to be indicative and non-exhaustive, as the Preamble to the ABA *Model Rules of Professional Conduct* put it, ‘[t]he Rules do not ... exhaust the moral and ethical considerations that should inform a lawyer, for no worthwhile human activity can be completely defined by legal rules’. Professional rules, it is said, cannot be a complete substitute for the lawyer’s personal moral compass,¹⁸ an idea that seems to emphasise both the important residual *role* for professional character and moral identity, and the risk that a heavily automated or augmented legal ethics might inhibit lawyers’ exercise of the moral virtues (cp Vallor, 2015).

More practically, it follows that (gen)AI cannot (yet) be seen as a reliable ethical decision-maker. It is poor at the underlying reasoning and judgement – it lacks a capacity for moral discernment and cannot autonomously reason from first principles which, in any event, often involve some element of incommensurability, or value or role conflict (Wallach & Vallor, 2020; Wendel, 2019). and, as we would expect, it struggles to apply ethical norms in practice given the messy, multi-step, multi-aspect thinking that such decision-making often requires (Jin et al., 2022). This does not mean that genAI cannot tell us *about* ethics and morality. At a relatively abstract level, it can generate helpful ethical statements, use moral concepts and thus potentially have moral effects as a decision support tool. But there is also an important institutional constraint on their further use.

Institutionally, the question of role also points us to deeper political-institutional, justifications for the legal function. Susskind’s (1986, pp.175–6) insistence that AIs ‘offer explanations for their lines of reasoning’ as a condition for ‘expertise’ is significantly relevant. The ability to find and give reasons on authoritative grounds is a fundamental feature of legal advice and decision-making. Practically-speaking, it is central to our ability, as citizens, to justify and enforce our use of the coercive powers of the law against another’s interests. But it also goes deeper than that, insofar as law’s authority and its legitimacy depends on its ability to provide reasons, and be

¹⁸ See, eg, *Public Trustee v By Products Traders PTE Ltd* [2005] 3 SLR(R) 499, [55].

accountable to the political community for those reasons (Wendel, 2019, pp.26–7).¹⁹ According to Wendel this process is necessarily relational – it occurs in the second person between moral agents ‘prepared to offer arguments that can be assessed for their soundness and accepted as reason-giving by other rational agents’ (id.). While this is most obviously central to the judicial role as currently conceived, it applies in a more limited fashion to lawyers. In many advisory settings, the matter goes no further than the lawyer’s office, or it is otherwise resolved without neutral adjudication. In such matters the lawyers effectively are the law and have a responsibility not to undermine public trust therein (Clark, 2018). If they do not give reasons, or, more likely, give reasons that are inconsistent with their professional obligations to the law and to the court, then they should be ethically accountable for such failures. This argument, it should be noted, is not special pleading in favour of the legal industry status quo. It doesn’t necessarily mean that the intermediating role has to be played by ‘a lawyer’ as currently understood by professional regulation. Whether that role could be fulfilled by a digital platform alone, however, raises significant questions of both social acceptability and philosophy (see, eg, Wendel, 2019; Davis, 2019). If those are to be resolved, the AI would first need to be able to make decisions that are explainable on proper legal and ethical grounds. Without this there is no real possibility of accountability.

In sum then, the deployment of genAI can be understood as one of a myriad ways in which society is potentially materialising agentic, including moral, effects in objects that, at a minimum, constitute ‘behaviour-guiding’ or ‘nudge’ technologies (Verbeek, 2011, 2017; Floridi et al., 2018) while lacking human moral capability in themselves. We have illustrated genAI’s moral agency limits with reference to three normative expectations that are core to conventional legal services delivery: competence, truthfulness, and compliance with professional ethical standards. Of these, genAI’s limited purpose competency is probably the least concerning. However, its lack of a meaningful standard of truth-telling, and its (associated) lack of explainability strongly suggests that genAI currently fails not just the truth and competence conditions set for professional advice, but the formal ethical condition as well. Taking these limitations as a totality, we can say that the technology so far lacks a foundation of ‘epistemic soundness’. But there is also a deeper sense in which we can say that genAI fails to do ethics, and this takes us into a yet more interesting question, so far relatively neglected in legal ethics, which is how the operations of technologies like genAI may undermine the *possibility* of not just meaningful oversight of the AI, but of its use as a tool to promote human flourishing through the law. That is the issue to which this paper now turns.

¹⁹ In some administrative decision-making settings, a failure to give reasons may be considered inconsistent with the Rule of Law: see, eg, discussion in *R. (on the application of Anufrijeva) v Secretary of State for the Home Department* [2003] UKHL 36, but it would obviously be a stretch, doctrinally, to extend that duty to give reasons to legal advice-giving in general.

4 Lawyers, Digital Transformation, and the Role of Virtue: Three Phases, Five Propositions and Four Proposals

In this section the paper turns to consider a further aspect of the lawyer's ethical role, which is the extent to which virtue plays a necessary part in 'good' decision-making (ie, decision-making that supports human flourishing), and how that might be impacted by legal genAI. The section divides into five parts which together seek to explain the (minimum) role of virtue in lawyers' ethics; to examine the potential impact of technological mediation on the capability of persons, and then specifically lawyers, to act as virtuous, 'wise counsellors'; to identify from that analysis a set of propositions regarding the impact of AI-assisted work on the practice of virtuous lawyering, and, lastly, to identify some (partial) solutions to the problems attending the practice of virtue in an AI-mediated legal world.

4.1 Virtue

As in many fields of applied ethics, underlying normative theory has, since the 1980s, been marked by a resurgence of interest in the virtue tradition, challenging the claims of both deontological and consequentialist accounts to be the only competing purveyors of a comprehensive normative ethics (Oakley & Cocking, 2001; Luban & Wendel, 2017, pp.363–4). The claim that certain moral (role) dispositions or 'virtues of character', such as courage, compassion, honesty, loyalty, a love of justice, and respect for dignity are not just good to have, but may have explanatory significance in relation to action, is widely shared across ethical traditions, though their centrality to the determination of right action is more contested (cp Oakley & Cocking, 2001, Dare, 1998, pp.11–15; Woolley & Wendell, 2010). This paper does not intend to take sides in what is, pragmatically-speaking, an irresolvable normative debate between these eminent schools of thought, rather, it starts from the modest premise that while formal legal practice norms themselves reflect a rule-centred paradigm, the realisation of that practice has features that are at least consistent with and may be supported by a non-principlist, virtue orientation (Arthur et al., 2012; Moorhead et al., 2012).²⁰ Moorhead et al. capture this quality when they describe ethical lawyering as an activity that requires sensitivity to the collective demands of 'character, context and capacity'.

In essence, there are three arguments at work here. The first is that actions in the Aristotelian sense of *praxeis* cannot be judged good or excellent in themselves without knowing something about the moral disposition of the agent. For virtue theorists, the habituation of character into virtuous dispositions thus plays an integral role in ethics. It is especially important in providing the basis for internalised standards of professional *excellence*, over and above the threshold legal standard of mere compliance (Van Domselaar, 2023, p.118). Secondly, standard rules and procedures are, in any event, often an insufficiently determinative guide to doing the right thing, so that qualities of character or 'virtue' are required to step-in and play a role in both

²⁰This argument is more fully acknowledged and worked through in the context of medical education and practice, than law: see, eg, Pellegrino & Thomasma (1993); Pellegrino (1995); Lyon (2021).

defining and motivating good action. There is, by definition, ‘no algorithm for [virtuous] ethical conduct’ (Tanesini, 2025, p.16).²¹ Thirdly, the mediating influence of organisational context and community norms on ethics, the importance of which is increasingly recognised in empirical studies of practice (eg, Mather, 2003) and is also captured in the contributions of virtue theorists (Annas, 2011; Beadle & Moore, 2006; MacIntyre, 1985) by the idea that situational sensitivity may in itself be a (minor) virtue,²² alongside consistency of character.

These ideas come together in the concept of *phronesis*, translated as ‘practical wisdom’. *Phronesis* can be understood as a kind of ‘meta-virtue’ that.

mediates over the character virtues; allowing someone to pick up on the morally and practically relevant features of a scenario, weigh them against one another, and decide how best to act to promote the ultimate end of human flourishing. (Zelny, 2025, p.3).

While anyone of reasonable intelligence might be able to produce good arguments (a pre-condition of *phronesis*), only a virtuous person can demonstrate practical wisdom, because only a virtuous person has the right rational motivations.

In legal ethics, the exercise of practical wisdom and civic-mindedness are at the heart of Anthony Kronman’s (1993) influential ‘lawyer-statesman’ or ‘wise counselor’ archetype. These are not merely skills, but central attributes of lawyerly identity. Practical wisdom requires not just what Legg (2024, p.276) calls ‘sound’ or ‘considered’ judgement, it requires a consistency²³ of judgement that is directed to a particular *telos*, a judging of the proper ends to which one’s capacities are to be directed. As Kronman observes, ‘a statesman aims at the good of the community to which he or she belongs and ...the statesman’s special virtue consists in an extraordinary devotion to this good and superior capacity for discerning where it lies’ (Kronman, 1993, p.54), while the emphasis on capacity to discern the public good of course mirror’s the lawyer’s ethical responsibility and commitment to the same, an idea that is at the heart of ‘social trustee’ professionalism (Fuller & Randall, 1958, p.1162; Brint, 1994), and the kinds of duties enumerated in Part 3.

Lawyerly practical wisdom seems to comprise three capabilities. It needs a deep sense of context: knowledge of the situation and the parameters of the conflict; a nuanced understanding of appropriate options that flow from the relevant virtue(s)

²¹ Virtue particularism (in this minimal sense) thus contrasts with the formalist’s reliance on rules and standards, which finds its equivalent in the technologist’s ‘craving for generality’ that Pasquale (2020, p.211) associates with ‘proponents of substitutive AI’.

²² Situational sensitivity is significant and in some respects a challenge across ethical theories. As Fletcher (1975, p.38) observes:

Many perimeters and parameters must be considered, including, inter alia, the people involved, their relationships, the objective circumstances, past as well as future factors, and remote as well as immediate consequences. The variables often complicate and obscure the entire picture. A misreading of a situation may cause the most thoughtful and sophisticated of “choosers” to make mistakes.

²³ This highlights that, while ‘sound judgement’ is a necessary component of practical wisdom, it is not sufficient, since the latter also requires the right motivation and attributes of character to make such judgements consistently.

(cp Russell, 2015, pp.37–38²⁴), and an appreciation of any competing values and incommensurabilities underlying the conflict and its resolution (Kronman, 1993, pp.69, 85). It requires sympathetic imagination – a capacity which is more than cognitive or even affective empathy (cp Steffanelo, 2024), entailing an openness and active engagement with others, such that one can richly imagine the experience and consequences that might flow from one's choices (Van Domselaar, 2023). Lastly, practical wisdom (for lawyers) also involves a 'suspended detachment', the 'power of moderating and confining feelings' (Kronman, 1993, p.74; Mortenson, 2023, pp.50–1; Legg, 2024, p.278) that comes from the practice of compassionate professional distance. These features thus capture how *phronesis* is both a careful balancing of reason and desire (Aristotle, 1985, *NE VI*, 1139a24–25) and a relational quality of judgement and of person(s)—it is, as Mortensen (2023, p.59) observes, a dialogical process of 'cooperative deliberation' between lawyer and client. Given these deeply humane and relational qualities, we might be sceptical from the outset that AI can have any role in the formation of the wise counsellor or *phronimos*. But is that scepticism misplaced?

4.2 The Technological Mediation of Virtue

There is an obvious question about how technology could impact, intermediate or even reflect these Aristotelian capacities (Tsai, 2020; Eisikovits & Feldman, 2022; Zelný, 2025). Legg succinctly points to AI's present inability to take the 'judgement step' required of the *phronimos* (Legg, 2024 p.276); while important, this gives an only partial, and somewhat negative answer. In engaging with this question, this paper therefore adopts the taxonomy of technological mediations proposed by Zelný (2025), drawing on the post-phenomenological framework for understanding technological mediation developed by Ihde (1979), Verbeek (2011) and others.

Zelný (2025, pp.8–12) argues that *phronesis* can be mediated in four distinct relations with technology, two of which are negative when it comes to nurturing virtue, and two of which may be more positive. He defines the negative processes as displacement, whereby human users are removed from situations that would properly or ordinarily require the exercise of practical wisdom, and diminution, whereby the need for practical wisdom within a situation is reduced, while keeping the human, to some extent, in the loop. The more positive modes are described as augmentation, that is adding or adapting a human user's ability to apply practical wisdom in a specific setting, and defining, which occurs when technology creates a new situation in which practical wisdom may be applied. How does this map onto legal practice?

4.3 Three Phases of Intermediation

This model, it is suggested, usefully highlights the (evolving) complexity of the legal technology-human ethics relation. We can illustrate this by looking at three relatively distinct 'levels' or phases of technology intermediation—in routine bulk process tasks; in 'intermediate' diagnostic work, and in augmenting bespoke advisory work.

²⁴ On algorithmic externalisation more generally, see, eg, Clark (2001); Chiodo (2020).

The last of these, bespoke advisory work, is perhaps the easiest to answer. The key marker of bespoke work is that it is traditionally human capital-intensive and valued as such (Armour & Sako, 2020; Legg & Bell, 2019). It will generally have some substantive legal and situational complexity, such that it needs a tailored, researched solution, and/or raises issues of factual or ethical uncertainty, and is conducted often in circumstances where there are thicker interpersonal connections between lawyer and client—eg, because the length, complexity or intensity of the matter itself builds some greater relationality, or just because the client is a familiar and valued repeat player. Performing well here requires a subtle and sophisticated interplay of practical, intellectual and ethical virtues that are reflective of *phronesis* as a process, not just an outcome: it involves both an excellence of information handling (Henning & Schweikard, 2015, p.3) and ‘an excellence of deliberation concerned with the very specification of determinable but as yet indeterminate ends’ (Russell, 2009, p.11). Critical here are the human capacities for (rich) situational understanding and discernment, including intersubjective abilities that are absent from the perception of technological tools, such as the ability to evaluate ‘what other people say, particularly about their own, or someone else’s, actions and/or feelings’ (Hursthouse, 2006, p.293). Such discriminatory capabilities enable the *phronimos* to ‘extract relevant information from one’s surroundings that might well be lost on others’ (Russell, 2015, p.25). Affective capabilities that are beyond AI, notably the capacity for sympathy-with-detachment highlighted by Kronman, and perhaps even the distinctively human capacity to feel both epistemic (Cabrera, 2021) and ethical anxiety (Kurth, 2015) about one’s deliberative processes and ultimate responsibilities—an idea that reaches its apogee in an ethics of alterity (see Nicolson & Webb, 2000, pp.46–8)—could also be said to distinguish human wisdom-in-action from the ‘mere’ epistemic soundness that genAI might yet achieve.

At the other end of the scale, we are already seeing, as discussed in Part 2, legal genAI’s use in support of triage, basic process functions, and straightforward legal writing. Importantly, this shows that processes of knowledge externalisation are, empirically-speaking, already underway.²⁵ On one view, such work routinely requires little ethical judgement and hence offers little opportunity (or risk) for *phronesis*. It might be argued that, unlike higher level, creative, autotelic work, the norms and relations governing routine technical work (more akin to the Aristotelian craft of ‘making’) might be more safely objectified and standardised within technological operations. Alternatively, however, there may still be subtle questions to consider around the moral significance of routine work for not just the client but the lawyer and the institution of legal practice. Does habituation to certain routine tasks inhibit reflective, ethical, professional practice? Or is the routine more important, and less amenable to standardisation, than we have assumed (cp Annas, 2011)? There is an argument that routine work also needs ethical judgements that are particular, responsive to ambiguity, and emergent (rather than ‘static’) responses to complexity (Nyberg, 2008). In this light, there is at least an empirical question over whether or how much the classification, allocation and doing of ‘routine’ work is something that

²⁵ Though, as noted above, these questions are being held somewhat in abeyance by current professional regulation.

both shapes and reflects the ethical culture of an organisation in important ways. For example, since much of the routine work that is the current target of genAI tools is work largely delegated to trainee and junior lawyers, there may be unforeseen risks that opportunities for social, ethical and contextual learning may be lost by automation of those tasks. The use of AI tools by novice lawyers, moreover, might in itself increase and create risks for the profession in two ways. First, risk may be created by enabling less skilled and experienced workers to ‘mask’ or disguise a lack of underlying substantive competence (cp Strich et al., 2021, p.319), thereby potentially undermining the skills and knowledge of the profession over time. Secondly, the hybridisation or outsourcing of expertise to the AI may lead to less responsible use of the technology and a loss of ethical mastery from the profession, a risk that is already present insofar as the professional habitat can be difficult, if not positively infertile, ground for ethical flourishing (Arthur et al., 2012; Vaughan & Oakley, 2016; Delacroix, 2018, 2022). Such risks beg critical questions not just of practice, but of legal education and training generally as to how they might build not just a better, more reflective, understanding of AI tools in use, but a deeper appreciation of the nature of both legal problem-solving and professional judgement and ethics (see, eg, Freid, 2024; Legg 2025). These are all ultimately empirical questions, but the point is that we at least need to ask the virtue question, and not assume that low level task automation does *not* involve a displacement or diminution of human practical wisdom.

It follows that the automation of ‘mid-level’ work perhaps represents the most significant foreseeable challenge for genAI and legal practice, since this likely contains an at best variable, at worst indeterminate, mix of the features of both the routine and the bespoke. The economic imperatives to instrumentalise and commoditise legal processes, discussed briefly in Part 3, are going to be important in shaping how and how much of this ‘mid-level’ work is automated. As noted, there are significant expectations that legal (gen)AI will move up the value chain of law firm offerings, creating ongoing questions for both service providers and regulators, about where the lines between automated and human decision-making should be drawn.²⁶ The use of technology assisted review in litigation discovery is a case in point. So long as those tools are sifting and flagging documents for human review, this is perhaps routine and (as discussed) less problematic. Indeed, the technology might even be seen as augmentative if it more objectively and reliably flags for human review documents over which, say, legal professional privilege might be claimed. However, the claiming of legal professional privilege itself is a highly complex decision, involving overlapping questions of law, strategy and ethics, particularly where the documents and the circumstances surrounding their making, cut close to the boundary of claims that the law will allow.²⁷ Attempts to provide more substantive decision support or even to make recommendations around privilege involves some diminution of human in the loop oversight, and a surrendering of some knowledge and ethical capabilities to the

²⁶Note that there are mounting concerns internationally over how privilege claims may be abused by lawyers to facilitate and hide client wrongdoing: see, eg, Higgins & Moorhead (2025); Campbell (2023).

²⁷The critical interface in this process is between lawyers and technologists, though this is but one feature in the emergence of an increasingly complex environment in which professional knowledge work is both more contingent, specialised and segmented, and yet operationalised in a more (digitally) connected inter-professional world (see, eg, Kritzer, 1999; Thornton, 2021; Noordegraf, 2020).

technology. Knowledge externalisation at scale starts to become a more significant professionalism issue, not least because organisational theory and practice seems to indicate that successful augmentation (in terms of epistemic soundness, let alone *phronesis*) requires organisations to do structural and identity work that will enable a more hybridised and increasingly multi-disciplinary conception of professional service to flourish (Nordegraaf, 2020; Armour, Parnham & Sako, 2022). This raises important issues for the legitimacy and epistemic authority of not just the technology, but the legal profession as well.

The discussion of these different phases and modes of adoption suggests that several provisional propositions can be advanced regarding the impact of genAI on professional work and ethics.

4.4 Five Propositions

First, it is clear that the move to AI-assisted legal practice is already involving lawyers in the construction of a new professional hybrid: that is, the gradual evolution of the professional into a human–machine assemblage (see Verbeek, 2011; Webb & Paterson, 2025).

Secondly, the construction of such an assemblage involves some degree of knowledge externalisation to the technology and away from humans. Such externalisation processes potentially challenge the epistemic authority of human experts and likely imply a diminution of *phronesis*. We might at this point wish to assert the need for some empirical proof that the assemblage can consistently function no less effectively as an epistemic expert than the human alone. The noted lack of reliable benchmarks, and the current inability of genAI to give reasons, however, undermine its epistemic authority on both functional and legal legitimacy grounds, as we saw in part 3, and places the functional onus on human in/on the loop processes.

Thirdly, this points us to a potentially critical tipping point for professional competence: should the lawyer-genAI assemblage achieve greater (explainable) epistemic expertise than the human alone, then *failure* to use AI starts to look like a matter of professional incompetence (cp Walters, 2019, pp.1074–78; Knake Jefferson, 2021). In normative terms, we might say that if (and only if) the lawyer-genAI assemblage can indeed demonstrate epistemic superiority, then it is these hybrid agents that become the proper subject to hold ‘expertise’ (and hence epistemic authority). What counts as ‘expertise’, particularly in terms of ethics, obviously becomes a critical question in this determination.

Fourthly, the more widely we define the necessary ethical domain, then, empirically, the less likely this tipping point is going to be achieved (any time soon). Following Zelny’s taxonomy, technological mediation is, on the above analysis, more likely to constitute a diminution rather than an enhancement or augmentation of practical wisdom. This is the likely consequence of an interplay of effects: the analytical and affective limitations of genAI, the possible underestimation of the ethical significance of routine and lower-level diagnostic knowledge work, and market pressures to construct riskier use-cases higher up the value chain as a way of increasing efficiency/value for money. Cultural and regulatory constraints are for the time being likely to act as some restraint on the degree and speed of diminution.

Finally, given propositions three and four, and the fact that there is, presently, little or no indication that genAI has adequate substitutive *ethical* capability, either in terms of professional or deeper normative ethical evaluation, then—given the centrality of both professional ethics and practical wisdom to the legal role – the ethical authority of the role (assemblage) should on normative grounds continue to vest ultimately in the human.

These propositions in turn point to four initial ways forward for the maintenance of not just the ethical use of legal AI, but of professional virtue.

4.5 Four Proposals

It is submitted, firstly, that this analysis supports the approach so far taken by much professional guidance, in requiring professional oversight and verification of legal genAI outputs. The knowledge, reliability and ethical limitations already discussed emphasise that genAI lacks both epistemic and ethical soundness and should not be released into the professional ‘wilds’, other than on a relatively short leash. This is not to say that the use of genAI lacks efficiency, cost and hence potential access benefits, but these should not override its (broad) competency, honesty and ethicality shortcomings.

Secondly, ethical professional oversight needs to extend into the crucial decisions to adopt genAI, both in the sense of the organisational decisions over what tools are selected (or built), and in terms of the responsibility and accountability of lawyers for the specific uses of these tools. This has so far received less attention from professional regulation (outside of the specific threats to existing duties, such as competence, confidentiality, and not misleading the court). Yet it seems to be a critical point, and one that the technology calls on us to address: in other words, the very existence and operation of legal genAI throws us into the ‘deliberative deep end’ of a ‘defining technological relation’ (Zelny, 2025, p.13). This, I suggest, is for three reasons: (i) the current uncertainties around genAI performance properly require us to treat technological intermediation itself as always already in question; (ii) allowing intermediation has consequences, it may be that once we allow an activity to be displaced wholly or substantially to the technology, we give up some capacity to ‘care for’ and ‘care why’ about that activity in a human way, and particularly to mediate techno-solutions through virtues like compassion, or even justice. (iii) There is also the risk of growing path-dependency if we don’t maintain a critical relation to the tools. As Jones (2018, p.657) puts it, while it may be unlikely that AI will ever achieve a high level of competence in human advocacy, the more important question may be ‘whether technology will allow any spheres where advocacy still matters’. The decision whether or not to use AI needs first and foremost to be a human one, and one that demands, at each of the individual, organisational and regulatory policy levels, *phronesis*.

Thirdly, requiring professional oversight only has value if that oversight can be effective. This places an obvious premium on technological competence, user training and a background of effective regulation. This, however, is also something of a professional weak spot. While there is widespread acceptance that technological competence forms part of the general duty of competence, the profession so far

has done relatively little to specify what such competence entails, while tending to assume that its existing codes do an adequate job in guarding against the prevalent ethical risks posed by AI and other advanced technologies, and it has largely eschewed any responsibility for AI benchmarking or platform regulation (see Needham, 2025; Webb, 2022).

Fourthly, there is also a question about what we might ethically demand of those who design legal genAI. As noted, benchmark standards are now beginning to emerge, and Delacroix (2018) has already noted the need for professional decision-support systems to be ‘ethically aware’ in a manner that aligns with Zelny’s ‘augmenting relation’. As this paper has highlighted, developments in genXAI could be valuable here. Zelny (2025, p.11–12) also illustrates ways in which technology might facilitate the exercise of practical wisdom through the development of prompting and deliberation tools or ‘reflection machines’ that enable human inquiry and critical evaluation of both process assumptions and outputs developed not so much by, but with the AI.

These kinds of innovations could be extremely important in both preserving a space for (normative) ethics and facilitating genAI’s safer progression up the legal value/legal complexity chain.

5 Conclusion: Holding the Space for Virtue

In conclusion, this paper has sought to highlight how genAI represents a significant problem for legal ethics, one that needs to be considered and addressed at a conceptual level, not just by focussing on the limited, conventional, reference points of competence, confidentiality and the duty not to mislead the court. GenAI models have been shown to be ethically problematic in a number of ways. They can produce fabricated or incomplete outputs, which risk undermining the honesty and integrity expected of legal professionals. They have limited capacity to explain their own reasoning and limitations (even with the use of guardrails and fine tuning); this not only makes it difficult to assess their reliability and trustworthiness, but may also undermine lawyers’ institutionally critical capacity to give reasons and be accountable on that reasoning. While genAI can perform well on certain well-defined tasks, it struggles with the nuances of real-world legal problem-solving, particularly complex relational problems where practical wisdom and contextual understanding play a significant role, over and above technical domain expertise. This is because the tools currently lack the capacity for moral discernment and the ability to autonomously reason from first principles, and from virtue-based reasoning, both of which are significant for professional decision-making. It suggests that the deep deployment of genAI in legal practice may thus lead to a diminution of practical wisdom, rather than its augmentation or enhancement.

The paper thus demonstrates that these limitations are not only important in their own right, they also serve to highlight the role of virtue ethics in professional practice. Based on this understanding, the paper pinpoints some initial steps that might be taken to preserve or even create space for practical wisdom. It thus re-emphasises the necessity of developing digital competence as a core professional capability, with a clear focus on understanding the broad ethical limitations of genAI; it stresses the

need to maintain human in or on the loop protocols for the oversight of AI tools, and it calls on lawyers and their regulators to recognise the importance of initial decision-making processes, and designers to develop tools that support ethical engagement and the development of practical wisdom. The decisions that organisations make when selecting and making tools available are themselves ethical decisions; for a lawyer, selecting a specific tool as fit for certain purposes is an ethical decision, and the specific case-by-case decision to use AI as a (partial) substitute for human work product likewise involves careful ethical evaluation. Each of these would benefit from *phronesis* at not less than one of the individual, organisational, or regulatory policy levels.

Acknowledgements My thanks to Jason Legocki for research assistance, and to participants at the Symposium on Virtue, AI and Law held at the Amsterdam Law School in March 2024, for their comments on an earlier version of this work.

Funding Open Access funding enabled and organized by CAUL and its Member Institutions. Research support and travel funds provided by the University of Melbourne.

Data Availability Not applicable.

Declarations

Conflict of interest None to declare.

Ethical approval Not applicable.

Consent to participate Not applicable.

Consent for publication Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- American Bar Association. (1992). Legal education and professional development – an educational continuum. Report of the Task Force on Law Schools and the Profession: Narrowing the gap. American Bar Association.
- Annas, J. (2011). *Intelligent virtue*. Oxford: Oxford University Press.
- Aristotle. (1985). *Nicomachean Ethics*. T. Irwin (Trans.). Hackett.
- Armour, J., Parnham, R., & Sako, M. (2022). Augmented lawyering. *University of Illinois Law Review*, 2022(1), 71–138.

- Armour, J., & Sako, M. (2020). AI-enabled business models in legal services: From traditional law firms to next-generation law companies? *Journal of Professions and Organization*, 7(1), 27–46. <https://doi.org/10.1093/jpo/joaa001>
- Arthur, J., Kristjánsson, K., Thomas, H., Holdsworth, M., Badini Confalonieri, L., & Qiu, T. (2012). *Virtuous character for the practice of law: Research report*. University of Birmingham.
- Baade, H. W. (1963). *Jurimetrics*. Basic Books.
- Beadle, R., & Moore, G. (2006). MacIntyre on virtue and organization. *Organization Studies*, 27(3), 323–340.
- Bhambhoria, R., Dahan, S., Li, J., & Zhu, X. (2024). Evaluating AI for Law: Bridging the Gap with Open-Source Solutions. *arXiv preprint*, arXiv 2404.12349.
- Blair-Stanek, A., Holzenberger, N., & Van Durme, B. (2023). Can GPT-3 perform statutory reasoning? *arXiv preprint arXiv:2302.06100*.
- Brint, S. (1994). *In an age of experts: The changing role of professionals in politics and public life*. New Jersey: Princeton University Press.
- Cabrera, F. (2021). Is epistemic anxiety an intellectual virtue? *Synthese*, 199, 13471–13495. <https://doi.org/10.1007/s11229-021-03383-2>
- Campbell, L. (2023). Legal professional privilege and corporate wrongdoing. *Adelaide Law Review*, 44, 339.
- Cantatore, F., Atwill, T., and Field, R. (2022). The job readiness of law graduates and entry level solicitors in private practice, final report. Queensland Law Society.
- Caserta, S. (2022). New technologies and law firms—an uneasy relationship: A European perspective. *Law, Technology and Humans*, 4(2), 183–196. <https://doi.org/10.5204/lthj.2246>
- Centre for Legal Innovation. (2024). The impact of generative AI in legal practice: Roundtables 2024 report, College of Law (Australia). https://vfpsua.files.cmp.optimizely.com/download/assets/2024+Legal+GenAI+Roundtables+Report_v2.pdf/7a4e6be2807c11efb3c16eaf20c0b6ad?submissionGuid=2385a790-ce9e-4a4f-ba87-7deaca280f40
- Chiodo, S. (2020). The greatest epistemological externalisation: Reflecting on the puzzling direction we are heading to through algorithmic automatisations. *AI & Society*, 35, 431–440. <https://doi.org/10.1007/s00146-019-00905-y>
- Choi, J. H., Hickman, K. E., Monahan, A. B., & Schwarcz, D. (2022). ChatGPT goes to law school. *Journal of Legal Education*, 71(3), 387–400.
- Clark, A. (2001). Natural-born cyborgs? In M. Beynon, C. L. Nehaniv, & K. Dautenhahn (Eds.), *Cognitive technology Instruments of mind. CT 2001. Lecture Notes in Computer Science*. Berlin Heidelberg: Springer.
- Clark, S. J. (2018). Work only we can do: Professional responsibility in an age of automation. *South Carolina Law Review*, 69(3), 533–555.
- Dahl, M., Magesh, V., Suzgun, M., & Ho, D. E. (2024). Large legal fictions: Profiling legal hallucinations in large language models. *Journal of Legal Analysis*, 16(1), 64–93. <https://doi.org/10.1093/jla/laae003>
- Davis, J. P. (2019). Artificial wisdom? A potential limit on AI in law (and elsewhere). *Oklahoma Law Review*, 72, 51–89.
- Delacroix, S. (2018). Computer systems fit for the legal profession? *Legal Ethics*, 21(2), 119–135. <https://doi.org/10.1080/1460728x.2018.1551702>
- Delacroix, S. (2022). *Habitual Ethics?* Hart Publishing.
- de Maat, E., Krabben, K., & Winkels, R. (2010) Machine learning versus knowledge-based classification of legal texts. In: Proceedings of the Jurix Conference: Legal Knowledge and Information Systems, pp 87–96. IOS Press.
- Du, Y. (2025). Confirmation bias in generative AI chatbots: Mechanisms, risks, mitigation strategies, and future research directions. *arXiv preprint arXiv:2504.09343*.
- Eiskovits, N., & Feldman, D. (2022). AI and phronesis. *Moral Philosophy & Politics*, 9(2), 181–199.
- Faulconbridge, J. (2025). Trajectories of legal work in the context of machine learning AI: Conceptualising mediated evolution. *International Journal of the Legal Profession*, 32(1), 97–120.
- Faulconbridge, J.R., Sarwar, A., & Spring, M. (2023). How professionals adapt to artificial intelligence: The role of intertwined boundary work. *Journal of Management Studies*, 65(5), 1991–2024.
- Faulconbridge, J. R., Sarwar, A., & Spring, M. (2024). Accommodating machine learning algorithms in professional service firms. *Organization Studies*, 45(7), 1009–1037. <https://doi.org/10.1177/01708406241252930>

- Floridi, L. (2021). Artificial agents and their moral nature. In L. Floridi (Ed.), *Ethics, governance, and policies in artificial intelligence*. Berlin Heidelberg: Springer.
- Floridi, L. (2023). AI as agency without intelligence: On ChatGPT, large language models, and other generative models. *Philosophy & Technology*, 36, Article 15. <https://doi.org/10.1007/s13347-023-00621-y>
- Floridi, L., Cows, J., Beltrametti, M., et al. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds & Machines*, 28, 689–707. <https://doi.org/10.1007/s11023-018-9482-5>
- Floridi, L., & Sanders, J. W. (2004). On the morality of artificial agents. *Minds and Machines*, 14(3), 349–379.
- Francesconi, E. (2022). The winter, the summer and the summer dream of artificial intelligence in law. *Artificial Intelligence & Law*, 30, 147–161. <https://doi.org/10.1007/s10506-022-09309-8>
- Fried, A. (2024). Unstructuring for insight: The legal profession in an age of AI and social change. *The Law Teacher*, 58(1), 74–91. <https://doi.org/10.1080/03069400.2023.2289789>
- Fuller, L. L., & Randall, J. (1958). Professional responsibility: Report of the joint conference. *American Bar Association Journal*, 44(12), 1159.
- Glaser, V. L., Pollock, N., & D'Adderio, L. (2021). The biography of an algorithm: Performing algorithmic technologies in organizations. *Organization Theory*, 2(2), 1–27. <https://doi.org/10.1177/26317877211004609>
- Greenleaf, G., Mowbray, A., & van Dijk, P. (1995). Representing and using legal knowledge in integrated decision support systems: DataLex WorkStations. *Artificial Intelligence & Law*, 3(1–2), 97–142. <https://doi.org/10.1007/BF00877696>
- Guha, N., Nyarko, J., Ho, D., Ré, C., Chilton, A., Chohlas-Wood, A., et al. (2023). Legalbench: A collaboratively built benchmark for measuring legal reasoning in large language models. *Advances in Neural Information Processing Systems*, 36, 44123–44279.
- Hart, H.L.A. (2012). *The concept of law*. 3rd edn., Clarendon Press.
- Henning, T., & Schweikard, D. P. (2015). Introduction. In T. Henning & D. P. Schweikard (Eds.), *Knowledge, Virtue, and Action: Essays on Putting Epistemic Virtues to Work*. Abington: Routledge.
- Higgins, A. & Moorhead, R. (2025). Reforms to privilege laws. The Post Office Project Working Paper 9. University of Exeter. <https://postofficeproject.net/wp-content/uploads/WP9-Legal-Professional-Privilege.pdf>
- Hodgkinson, G.P., Gulati, S., Nokes, K., James, A.D., & Kununka, S. (2023). Attitudes toward lawtech adoption: Findings from a survey of solicitors in England and Wales. The Law Society. https://pure.manchester.ac.uk/ws/portalfiles/portal/280905249/Lawtech_study_report_web_version.pdf
- Hursthouse, R. (2006). Practical wisdom: A mundane account. *Proceedings of the Aristotelian Society*, 106, 285–309. <https://doi.org/10.1111/j.1467-9264.2006.00149.x>
- Ihde, D. (1979). *Technics and praxis*. Reidel.
- Jin, Z., Levine, S., Gonzalez Adatao, F., Kamal, O., Sap, M., Sachan, M., et al. (2022). When to make exceptions: Exploring language models as accounts of human moral judgment. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho & A. Oh (Eds.), *Advances in neural information processing systems* 35. NeurIPS. https://proceedings.neurips.cc/paper_files/paper/2022/hash/b654d6150630a5ba5df7a55621390daf-Abstract-Conference.html
- Jones, C. P. A. (2018). AI and the legal profession: Could artificial stupidity and responsibility avoidance prove to be the biggest agents of change? *South Carolina Law Review*, 69(3), 637–657.
- Knake Jefferson, R. (2021). Lawyer ethics for innovation. *Notre Dame Journal of Legal Ethics and Public Policy*, 35, 1–48.
- Kritzer, H. M. (1999). The professions are dead, long live the professions: Legal practice in a postprofessional world. *Law & Society Review*, 33(3), 713–759. <https://doi.org/10.2307/3115110>
- Kurth, C. (2015). Moral anxiety and moral agency. In M. Timmons (Ed.), *Oxford studies in normative ethics* (Vol. 5, pp. 171–195). Oxford University Press.
- Lassébie, J. Quintini, G. (2022). What skills and abilities can automation technologies replicate and what does it mean for workers? New evidence. *OECD Social, Employment and Migration Working Papers* 282. OECD Publishing. https://www.oecd.org/content/dam/oecd/en/publications/reports/2022/12/what-skills-and-abilities-can-automation-technologies-replicate-and-what-does-it-mean-for-workers_262a2818/646aad77-en.pdf
- Law Council of Australia, *Australian Solicitors Conduct Rules 2022: Commentary* (March 2024) <https://lawcouncil.au/files/pdf/policy-guideline/ASCR%20Commentary.pdf>

- Law Society of New South Wales, Legal Practice Board of Western Australia & Victorian Legal Services Board + Commissioner. (2024). Statement on the use of artificial intelligence in Australian legal practice. <https://www.lsb.vic.gov.au/news-updates/news/statement-use-artificial-intelligence-australian-legal-practice>
- Legg, M. (2024). Better than a bot – Instilling ethical judgement into the lawyers of the future in the age of AI. *Griffith Law Review*, 33(3), 273–293. <https://doi.org/10.1080/10383441.2025.2493493>
- Legg, M., & Bell, F. (2019). Artificial intelligence and the legal profession: Becoming the AI-enhanced lawyer. *University of Tasmania Law Review*, 38(2), 34–59.
- Loevinger, L. (1949). Jurimetrics. The next step forward. *Minnesota Law Review*, 33(5), 455–493.
- Luban, D., & Wendel, W. B. (2017). Philosophical legal ethics: An affectionate history. *Georgetown Journal of Legal Ethics*, 30, 337–364.
- Lyon, W. (2021). Virtue and medical ethics education. *Philosophy, Ethics, and Humanities in Medicine*, 16(1), Article 2. <https://doi.org/10.1186/s13010-021-00100-2>
- MacIntyre, A. (1985). *After virtue: A study in moral theory*. 2nd edn., Duckworth.
- Martinez, E. (2025). Re-evaluating GPT-4's bar exam performance. *Artificial Intelligence and Law*, 33, 581–604. <https://doi.org/10.1007/s10506-024-09396-9>
- Mather, L. (2003). What do clients want? What do lawyers do? *Emory Law Journal*, 52, 1065–1086.
- McCarthy, L. T. (1977). Reflections on TAXMAN: An experiment in artificial intelligence and legal reasoning. *Harvard Law Review*, 90(5), 837–893.
- McCarthy, L. T. (1990). Artificial intelligence and law: How to get there from here. *Ratio Juris*, 3(2), 189–200. <https://doi.org/10.1111/j.1467-9337.1990.tb00057.x>
- McGinnis, J. O., & Pearce, R. G. (2014). The great disruption: How machine intelligence will transform the role of lawyers in the delivery of legal services. *Fordham Law Review*, 82, 3041.
- Mehl, L. (1959). Automation in the legal world: From the machine processing of legal information to the 'law machine'. In National Physical Laboratory, *Mechanisation of Thought Processes: Proceedings of a Symposium held at the National Physical Laboratory, 24th, 25th, 26th and 27th November 1958*, (vol. II, pp. 755–787). Her Majesty's Stationery Office.
- Moorhead, R., Hinchly, V., Parker, C., Kershaw, D. and Holm, S. (2012). Designing ethics indicators for legal services provision. UCL Centre for Ethics and Law Working Paper No1. University College, London. www.legalservicesboard.org.uk/what_we_do/Research/Publications/pdf/designing_ethics_indicators_for_legal_services_provision_lsb_report_sep_2012.pdf
- Mortensen, R. (2023). The Lost Lawyer regained: Virtue, liberalism and citizenship in lawyers' ethics. In J. Webb (Ed.), *Leading works in legal ethics* (pp. 44–66). Routledge.
- Needham, C. A. (2025). Regulation of the use of generative artificial intelligence tools in the delivery of legal services. *Washington University Journal of Law & Public Policy*, 77, 184–211.
- Niblett, A., & Yoon, A. (2024). AI and the nature of disagreement. *Philosophical Transactions of the Royal Society a: Mathematical, Physical and Engineering Sciences*, 382, Article 20230162. <https://doi.org/10.1098/rsta.2023.0162>
- Nicolson, D., & Webb, J. (2000). *Professional legal ethics: Critical interrogations*. Oxford: Oxford University Press.
- Nordeggraaf, M. (2020). Protective or connective professionalism? How connected professionals can (still) act as autonomous and authoritative experts'. *Journal of Professions and Organization*, 7, 205.
- Nyberg, D. (2008). The morality of everyday activities: Not the right, but the good thing to do. *Journal of Business Ethics*, 81, 587–598. <https://doi.org/10.1007/s10551-007-9530-1>
- Oakley, J., & Cocking, D. (2001). *Virtue Ethics and Professional Roles*. Cambridge: Cambridge University Press.
- Ohm, P. (2024). Focusing on fine-tuning: Understanding the four pathways for shaping generative AI. *Science and Technology Law Review*, 25(2), 214–240. <https://doi.org/10.52214/stlr.v25i2.12762>
- Paliwala, A. (2010). A history of legal informatics: An introduction to the special issue. *European Journal of Law and Technology*, 1(1). <https://ejlt.org/index.php/ejlt/article/view/21/38>
- Palminteri, S., & Lebreton, M. (2022). The computational roots of positivity and confirmation biases in reinforcement learning. *Trends in Cognitive Sciences*, 26(7), 607–621. <https://doi.org/10.1016/j.tics.2022.04.005>
- Pasquale F. 2016 *The black box society: the secret algorithms that control money and information*. Harvard University Press. Cambridge
- Pasquale, F. (2020). *The new laws of robotics: Defending human expertise in the age of AI*. Massachusetts: Belknap Press.

- Pellegrino, E. D. (1995). Toward a virtue-based normative ethics for the health professions. *Kennedy Institute Ethics Journal*, 5, 253–277.
- Pellegrino, E. D., & Thomasma, D. C. (1993). *The virtues in medical practice*. Oxford: Oxford University Press.
- Peoples, L. (2025). Artificial Intelligence and Legal Analysis: Implications for Legal Education and the Profession. *arXiv preprint arXiv:2502.03487*.
- Pictor, M. (2023). Where does responsibility lie? Analysing legal and regulatory responses to flawed clinical decision support systems when patients suffer harm. *Medical Law Review*, 31(1), 1–24. <https://doi.org/10.1093/medlaw/fwac022>
- Remus, D., & Levy, F. S. (2017). Can robots be lawyers? Computers, lawyers, and the practice of law. *Georgetown Journal of Legal Ethics*, 30, 501–558.
- Romoser, J. (2023). No, Ruth Bader Ginsburg did not dissent in *Obergefell* – and other things ChatGPT gets wrong about the Supreme Court, SCOTUS Blog, (26 January). <https://www.scotusblog.com/2023/01/no-ruth-bader-ginsburg-did-not-dissent-inobergefell-and-other-things-chatgpt-gets-wrong-about-the-supreme-court/> [<https://perma.cc/A487-P3S7>].
- Russell, D. C. (2009). *Practical intelligence and the virtues*. Oxford University Press.
- Russell, D. C. (2015). Aristotle on cultivating virtue. In N. E. Snow (Ed.), *Cultivating virtue: Perspectives from philosophy, theology, and psychology* (pp. 17–48). Oxford University Press.
- Schneider, J. (2024). Explainable generative AI (GenXAI): A survey, conceptualization, and research agenda. *Artificial Intelligence Review*, 57(289), 1–38. <https://doi.org/10.1007/s10462-024-10916-x>
- Schwarz, D., & Choi, J. H. (2023). AI tools for lawyers: A practical guide. *Minnesota Law Review Headnotes*, 108, 1–39.
- Scott, W. R. (2008). Lords of the dance: Professionals as institutional agents. *Organization Studies*, 29(2), 219–238. <https://doi.org/10.1177/0170840607088151>
- Shi, J., Jain, R., Doh, H., Suzuki, R., & Ramani, K. (2023). An HCI-centric survey and taxonomy of human-generative-AI interactions. *arXiv preprint arXiv:2310.07127*.
- Solicitors Regulation Authority. (2022). Compliance tips for solicitors regarding the use of AI and technology. <https://www.sra.org.uk/solicitors/resources/innovate/compliance-tips-for-solicitors/>
- Strich, F., Mayer, A.-S., & Fiedler, M. (2021). What do I do in a world of artificial intelligence? Investigating the impact of substitutive decision-making AI systems on employees' professional role identity. *Journal of the Association for Information Systems*, 22(2), 304–324. <https://doi.org/10.17705/1jais.00663>
- Susskind, R. (1986). Expert systems in law: A jurisprudential approach to artificial intelligence and legal reasoning. *Modern Law Review*, 49(2), 168–194. <https://doi.org/10.1111/j.1468-2230.1986.tb01683.x>
- Susskind, R. (2008). *The end of lawyers?* Oxford University Press.
- Susskind, R., & Susskind, D. (2022). *The future of the professions*. Oxford: Oxford University Press.
- Tanesini, A. (2025). Non-idealised virtue epistemology as particularist virtue theory. *International Journal of Philosophical Studies*. <https://doi.org/10.1080/09672559.2025.2456822>
- Tan, J., Westermann, H., & Benyekhlef, K. (2023). ChatGPT as an artificial lawyer? In L.K. Branting (Ed.), *Proceedings of the ICAIL Workshop on Artificial Intelligence for Access to Justice (AI4AJ 2023)* 47. <https://ceur-ws.org/Vol-3435/short2.pdf>
- Thornton, M. (2021). Legal professionalism in a context of uberisation. *International Journal of the Legal Profession*, 28(3), 243–263. <https://doi.org/10.1080/09695958.2021.1901715>
- Tsai, C. (2020). Artificial wisdom: A philosophical framework. *AI & Society*, 35(4), 937–944.
- Vallor, S. (2015). Moral deskilling and upskilling in a new machine age: Reflections on the ambiguous future of character. *Philosophy & Technology*, 28, 107–124.
- Vallor, S. (2017). AI and the automation of wisdom. In T. M. Powers (Ed.), *Philosophy and Computing* (pp. 161–178). Berlin Heidelberg: Springer International.
- Van Domselaar, I. (2023). The fragility of legal ethics: on the role of theory, lawyerly virtues, and moral remainders in the life of a good lawyer. In J. Webb (Ed.), *Leading works in legal ethics*. Oxfordshire: Routledge.
- van Es, K., & Nguyen, D. (2024). “Your friendly AI assistant”: The anthropomorphic self-representations of ChatGPT and its implications for imagining AI. *AI & Society*. <https://doi.org/10.1007/s00146-024-02108-6>
- Vaughan, S., & Oakley, E. (2016). ‘Gorilla exceptions’ and the ethically apathetic corporate lawyer. *Legal Ethics*, 19(1), 50–75. <https://doi.org/10.1080/1460728x.2016.1189681>
- Verbeek, P. P. (2011). *Moralizing technology: Understanding and designing the morality of things*. Chicago: University of Chicago Press.

- Verbeek, P. P. (2017). Designing the morality of things: The ethics of behaviour-guiding technology. In J. Hoven & van den, S Miller, T Pogge, (Eds.), *Designing in ethics*. Cambridge: Cambridge University Press.
- Wallach, W., & Vallor, S. (2020). Moral machines: From value alignment to embodied virtue. In S. M. Liao (Ed.), *Ethics of artificial intelligence* (pp. 383–412). Oxford University Press.
- Walters, E. (2019). The model rules of autonomous conduct: Ethical responsibilities of lawyers and artificial intelligence. *Georgia State University Law Review*, 35, 1073–1092.
- Webb, J. (2022). Legal technology: The great disruption? In R.L. Abel, H. Sommerlad, O. Hammerslev & U. Schultz (Eds.), *Lawyers in 21st-Century societies. Vol. 2: Comparisons and theories*. Hart Publishing.
- Webb J and Paterson. J. M. (2025). The evolution of legal knowledge work in an age of brilliant(?) technologies: From robo-lawyer to digital law clerk. *ANU Journal of Law and Technology*, 6(1), 51–84. <https://anujolt.org/article/150388-the-evolution-of-legal-knowledge-work-in-an-age-of-brilliant-technologies-from-robo-lawyer-to-digital-law-clerk>
- Webley, L., Flood, J., Webb, J., Bartlett, F., Galloway, K., & Tranter, K. (2019). The profession's engagements with lawtech: Narratives and archetypes of future law. *Law, Technology and Humans*, 1(1), 6–26. <https://doi.org/10.5204/lthj.v1.i1.1314>
- Woolley, A., & Wendel, W. B. (2010). Legal ethics and moral character. *Georgetown Journal of Legal Ethics*, 23, 1065–1104.
- Zelny, A. (2025). Offloading wisdom: Four technological relations that mediate phronesis. *Philosophy & Technology*, 38(56), 1–23.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.