Beyond Verification: Professional Duty and Workflow Governance in AI-Assisted Legal Practice

Authors: Donald Feaver
Paper Type: Working Paper

Working Paper Series Details
Series: *Leo Cussen Working Paper Series*
Working Paper No. 9: Donald Feaver, 'Beyond Verification: Professional Duty and Workflow Governance in AI-Assisted Legal Practice' (Working Paper No 9, Leo Cussen Centre for Law, 2026)

For more information, please contact:
Professor Donald Feaver
Leo Cussen Centre for Law
dfeaver@leocussen.edu.au

Copyright and Disclaimer

Copyright © Leo Cussen Centre for Law and the authors. This working paper is circulated for discussion and may be revised. The views expressed are those of the authors and do not necessarily reflect the views of Leo Cussen Centre for Law.

Beyond Verification: Professional Duty and Workflow Governance in AI-Assisted Legal Practice

Professor Donald Feaver
Leo Cussen Centre for Law

Abstract

Generative artificial intelligence can produce documents that resemble completed legal work without performing the professional workflow through which legal work is ordinarily framed, reasoned, checked and authorised. This creates a structural problem for professional responsibility. Existing duties of competence, diligence, confidentiality, independence, candour and supervision remain applicable, but do not fully identify where those duties must be exercised when legal work reflects divided human and technological contributions.

This article develops **workflow governance** as an organising framework for AI-assisted legal practice. It argues that there remains one normatively primary professional legal workflow, within which AI may perform computational subprocesses such as retrieval, extraction, classification, summarisation and drafting. Where AI materially affects that workflow, task-specific human legal skill remains necessary. **Convergence** occurs when technological contributions are transformed through human legal judgment into professionally usable legal work. **Friction** occurs where the required human skill, control or authority is absent, inadequate or cannot be evidenced.

Connecting this framework with Australian professional regulation, court guidance and emerging case law, the article reframes AI failures as workflow-friction problems. Solicitor conduct law supplies the duties; workflow governance identifies their operational location.

1. Introduction: the problem of apparently completed legal work

Generative artificial intelligence can produce outputs that take the form of artefacts that look like finished legal work without performing the professional processes that would normally occur if the work was produced by a human author.[1] An AI generated research

[1] In this article, AI output refers to material generated by the system before professional evaluation. "Artefact" refers to the document or work product that is ultimately communicated, filed, relied upon or recorded. The professional risk arises because AI output may be incorporated into an artefact that appears complete, even though the legal work required to make it professionally usable has not been performed.

output might bear a strong resemblance to tasks such as issue identification, source selection, authority parsing, the assessment of relevance and its application to facts. An AI generated letter might appear to address salient instructions, contain credible factual analysis, legal reasoning and even strategic judgment. However, generative AI can only produce an appearance of a finished legal artefact.

The distinction between what appears to be a legal artefact and what, in fact, is a legal artefact should not be underestimated. Generative AI produces text in structure and content that mimics professional work: paragraphs, headings, citations, reasons, conclusions and recommendations. As a result, there is a temptation to treat the material as though it has already passed through the professional process. However, without the professional process, the work remains incomplete. A block of text that reads like sound legal reasoning is not legal reasoning. A citation that looks authoritative is not an authority. A draft that looks finished is not professionally usable legal work. If that simple distinction is kept firmly in mind, the status of AI output becomes clear. It is not a finished legal artefact only requiring verification and approval.

A great deal of emphasis is placed on the need to verify. However, verification is only one step in the much more nuanced interaction that occurs when AI is used to help complete legal work. Examined more closely, a lawyer using AI should also govern how the technology frames the task, organises information, shapes reasoning, influences judgment and contributes to the final artefact- actions that all require human skill judgement and expertise. The sequence of actions a lawyer should take are all embedded in what can be described as the broader workflow that leads to the production of legal artefacts.

Workflow is not usually treated by lawyers as a concept of professional responsibility, although it is a fundamental organising concept in information systems, management, healthcare and other professional fields.[2] In this article, a workflow means the systematic sequence of linked and logically ordered actions through which legal work is framed, performed, checked, authorised and brought to an output. It is the process by which the problem is framed, the facts and law are identified, competing possibilities are evaluated, risks are assessed, the work is checked, and a person with professional authority accepts responsibility for the result. A failure to exercise legal judgment at any of those points may allow the technology to frame the wrong problem, distort the factual basis, narrow the legal analysis, obscure uncertainty or produce an artefact that appears complete without having been professionally made so. [3]

---

[2] Thomas H Davenport, *Process Innovation: Reengineering Work through Information Technology* (Harvard Business School Press, 1993) 5–7; Wil van der Aalst and Kees van Hee, *Workflow Management: Models, Methods, and Systems* (MIT Press, 2002) 2–5; Marc Berg, 'Patient Care Information Systems and Health Care Work: A Sociotechnical Approach' (1999) 34(2) *International Journal of Medical Informatics* 87, 90–5.

[3] Rachel Reed, 'Can ChatGPT Replace Judges?' *Harvard Law Today* (Web Page, 30 September 2025); Thomson Reuters, *Tech, AI and the Law 2024 Report: Australian Edition* (Report, 24 October 2024); Deloitte Legal, 'How Generative AI Is Changing Legal Department Functions' (Web Page, 5 September 2024).

This article argues that a broadening of the scope of professional responsibility when using AI assistance to complete legal work is best understood through the concept of **workflow governance**. As mentioned above, workflow is the sequence through which legal work is framed, performed, checked, communicated and recorded.[4] Without AI, much of that sequence is tacit. Lawyers may not formally describe every stage, but legal practice routinely depends upon these ordered processes for taking instructions, identifying facts and issues, researching law, analysing consequences, drafting documents, advising clients, supervising others and maintaining an adequate file.[5]

AI does not create an alternative professional workflow of equal normative status. There remains one legal workflow, structured by client purpose, law, professional duties, institutional responsibility and human authority. Within that workflow, AI may perform or assist computational subprocesses: retrieval, classification, extraction, summarisation, comparison, prediction or generation. Those subprocesses may contribute to legal work, but they do not themselves make the contribution professional.

The central claim of this article is that professional responsibility attaches at the workflow points where AI assistance materially affects legal work. Workflow identifies **where** responsibility must be exercised. Legal skill identifies **what human capability** is required at that point. Convergence occurs when technological assistance is transformed through human legal judgment into professionally usable legal work. Friction occurs when the required skill, control or authority is absent, inadequate or incapable of being evidenced.

This article does not propose that AI creates a new field of professional duties. As will be discussed below, the *Legal Profession Uniform Law* and the Australian Solicitors' Conduct Rules already contain the normative foundations required to regulate AI-assisted practice. They require competence, diligence, confidentiality, integrity, independence, clear communication, candour and supervision.[6] The difficulty is not primarily normative deficiency. It is operational under-specification. Existing rules state what lawyers owe, but they do not always identify with sufficient precision where those duties must be exercised when legal work is technologically mediated.

The article makes four contributions. First, it offers an organising framework connecting reliability, verification, automation bias, legal ethics, agentic workflows and supervision. Second, it treats workflow as the normative location of professional responsibility rather than merely a method of automation or service design. Third, it interprets emerging AI

---

[4] Thomas H Davenport, *Process Innovation: Reengineering Work through Information Technology* (Harvard Business School Press, 1993) 5–7; Wil van der Aalst and Kees van Hee, *Workflow Management: Models, Methods, and Systems* (MIT Press, 2002) 2–5.

[5] Richard Susskind, *The End of Lawyers? Rethinking the Nature of Legal Services* (Oxford University Press, 2008) 28–38, 99–116; Richard Susskind, *Tomorrow's Lawyers: An Introduction to Your Future* (Oxford University Press, 3rd ed, 2023) chs 3–5.

[6] *Legal Profession Uniform Law* (NSW) ss 3, 296–8; *Legal Profession Uniform Law Application Act 2014* (Vic) sch 1; *Legal Profession Uniform Law Application Act 2022* (WA) sch 1; *Legal Profession Uniform Law Australian Solicitors' Conduct Rules 2015* rr 4.1.1, 4.1.3, 4.1.4, 7.1, 8.1, 9.1, 17.1, 19.1, 19.6, 22.1, 37.1.

decisions as workflow-friction cases, showing how existing duties become concrete when technological output crosses into legal work without the required human skill and control. Fourth, it provides a basis for proportionate professional guidance, firm governance, supervision and recordkeeping. The article proceeds as follows. Part 2 positions the argument within the existing literature and identifies the limits of verification-centred accounts. Part 3 develops the workflow-governance model. Part 4 locates the *Uniform Law* and Conduct Rules within the workflow. Part 5 analyses the emerging cases as different forms of workflow friction. Part 6 extends the framework beyond court-facing failure through the example of an AI-assisted letter of advice. Parts 7 and 8 examine supervision, institutional responsibility and regulatory reform.

2. Beyond verification: reliability, ethics and the missing workflow architecture

The literature examining the impacts of generative AI on legal practice has developed rapidly. Four overlapping streams are particularly relevant to this article. The first stream focuses on AI's reliability and the well-known concept of hallucination. Empirical work demonstrates that general-purpose language models can generate legally formatted answers, authorities and quotations that are inaccurate or fictitious.[7] Importantly, reliability problems are not confined to general-purpose systems. Evaluations of specialist legal research tools indicate that legal-domain systems and retrieval-augmented tools may also produce incorrect, incomplete or misleading answers.[8] These studies establish a necessary professional starting point: legally fluent output is not the same as legally usable work.

The second stream concerns litigation misuse and human overreliance. Legg, McNamara and Alimardani analyse the entry of fabricated authorities and inaccurate legal material into proceedings and connect that conduct with automation bias and verification drift.[9] Automation bias may tempt users to defer to a system's apparent authority. This can occur when lawyers are under pressure due to time and resource constraints. Verification drift describes the weakening of checking practices over time, particularly as users become accustomed to outputs that appear reliable. These phenomena explain why AI-generated errors may survive even where the lawyer understands, in the abstract, that the technology is fallible.

---

[7] Matthew Dahl, Varun Magesh, Mirac Suzgun and Daniel E Ho, 'Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models' (2024) arXiv:2401.01301.

[8] Varun Magesh, Faiz Surani, Matthew Dahl, Mirac Suzgun, Christopher D Manning and Daniel E Ho, 'Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools' (2024) arXiv:2405.20362.

[9] Michael Legg, Vicki McNamara and Armin Alimardani, 'The Promise and the Peril of the Use of Generative Artificial Intelligence in Litigation' (2025) 48(4) *University of New South Wales Law Journal* 1196, 1216–35.

The third research stream considers AI assisted practice and legal ethics and professional responsibility. Perlman applies established duties of competence, confidentiality, communication, supervision and candour to generative AI use.[10] Murray similarly argues that the movement from peripheral tools towards more integrated and agentic systems requires renewed attention and focus on competence, professional autonomy, diligence, confidentiality and supervision.[11] This literature correctly rejects the suggestion that technological assistance displaces conventional professional obligations- a topic that is considered more deeply later in this article.

The fourth stream concerns workflow and human–AI collaboration. O'Grady and O'Grady treat agentic workflows as a site of ethical design and governance.[12] Han and others describe legal fact verification as a collaborative but professionally accountable activity requiring auditable systems.[13] Yuvaraj's verification-value paradox argues that the efficiency promised by generative AI may be reduced or negated by the cost of the manual verification required to make its outputs professionally safe.[14] Linna and Linna demonstrate why legal reasoning cannot be reduced to fluent textual synthesis: it requires selection of the relevant legal framework, treatment of authority, management of ambiguity and application of law to fact.[15]

Together, these streams establish that AI-generated legal material may be unreliable, that lawyers may over-trust it, that verification is necessary and costly, that professional duties continue to apply, and that AI is increasingly integrated into legal workflows. What remains underdeveloped is an architecture connecting those insights within the process of legal work.

Reliability studies identify the technological risk but do not explain the professional location of responsibility. Ethics analyses identify the relevant duties but often remain at a level of generality. Verification accounts identify an essential safeguard but can imply that the professional problem begins only after an output has been generated. Workflow scholarship recognises that AI systems increasingly operate through sequences of

---

[10] Andrew M Perlman, 'The Legal Ethics of Generative AI' (2024) Suffolk University Law School Research Paper No 24-17.

[11] Michael D Murray, 'Algorithmic Ethics in an Era of Agentic AI Advocacy: An Analysis of AI's Impact on the Model Rules of Professional Conduct and the Model Code of Judicial Conduct' (2026) 16 *St Mary's Journal on Legal Malpractice and Ethics* 279.

[12] Catherine Gage O'Grady and Casey Sims O'Grady, 'Agentic Workflows in the Practice of Law: AI Agents as Ethics Counsel' (2025) 38 *Georgetown Journal of Legal Ethics* 247.

[13] Sirui Han et al, 'Reimagining Legal Fact Verification with GenAI: Toward Effective Human-AI Collaboration' (2026) arXiv:2602.06305.

[14] Joshua Yuvaraj, 'The Verification-Value Paradox: A Normative Critique of Gen AI in Legal Practice' (2026) 52 *Monash University Law Review* (forthcoming); final title, volume and pinpoints to be confirmed.

[15] Eljas Linna and Tuula Linna, 'Judicial Requirements for Generative AI in Legal Reasoning' (2025) arXiv:2508.18880.

actions but do not distinguish technological workflow from the normative structure of legal work.

A verification-centred approach is necessary, as mentioned above, but is incomplete. AI may cause professional harm well before there is a final output to verify. The lawyer may select an inappropriate task, use an unsuitable tool, disclose confidential information, accept an inaccurate summary of instructions, allow AI-generated categories to narrow the legal problem, overlook missing facts, adopt a defective reasoning structure or supervise only the final artefact. A final citation check cannot necessarily cure those upstream failures. Conversely, should a lawyer master the workflow interactions at critical risk points, a more accomplished use of AI tools through the workflow is likely to improve the quality of the output and thereby reduce the verification burden.

This point can be illustrated using a simple legal task as an example. A conventional advice workflow can be expressed as:

> Frame the problem → establish the facts → identify the issues → research → analyse → draft → advise → record.

Where AI assists, additional professional control points arise:

> **Assess task suitability → select the tool → control the information → frame the instruction → test the output → verify sources and facts → examine reasoning → revise for purpose → exercise judgment → authorise and record.**

These control points do not replace the legal workflow. They overlay or augment it. They identify the additional decisions required because AI has been introduced into the production of legal work. On the surface, this appears to be a significant shift, or even complex one. However, with proper education and training, it provides the architecture for legal AI use best practice and the development of good habits when using AI tools.

Whereas, the prior literature has identified key themes and problems, a better understanding of workflow and their governance connects those themes. For example, hallucination becomes dangerous when false output enters a professional workflow without legally skilled verification. Automation bias becomes professionally significant when technological fluency displaces independent legal judgment. Verification drift becomes friction when checking weakens at the point at which technological material must be converted into legal work. Agentic systems matter because they redistribute task performance and alter the points at which human intervention, supervision and authority are required. The contribution of workflow governance is not to replace existing scholarship. It is to provide a middle-level framework linking general professional duties with AI-assisted tasks and failures.

3. Workflow governance and the human–AI professional interface

3.1 One professional legal workflow

The concept of workflow has long been used to describe the organisation and sequencing of complex work.[16] In legal practice, workflow is more than administration. It is the architecture through which legal judgment is exercised. The legal workflow is normative because it is structured by professional purpose. A legal research task is not complete because information has been retrieved. The relevant law must be identified, evaluated and applied to the matter. A draft isn't complete merely because it is fluent. It must be accurate, appropriately qualified, suited to its audience and consistent with the lawyer's duties. Advice is not complete because options or client guidance has been generated. A lawyer must decide what can responsibly be recommended and communicated.

AI systems operate differently than analogue human legal workflows. AI responds to inputs and system design through processes such as retrieval, ranking, prediction, pattern recognition and generation.[17] They may produce outputs resembling the endpoint of a legal task without reproducing the legal reasoning, ethical judgment or institutional accountability through which that task would ordinarily be performed. As a result, we have a collision, of sorts, between two different professional workflows- on endemic to informational technologies and the other legal. A question then arises as to whether these two different workflows are compatible?

The relationship should therefore not be understood as the combination of two different or coequal professional workflows. AI performs technical subprocesses within a human-governed legal workflow. The technological subprocess may be sophisticated and may materially shape the result, but the legal workflow must be the normatively primary workflow. This is a critical point. Legal primacy avoids the risks inherent with treating AI as a professional actor. AI cannot owe loyalty to the client, protect the administration of justice, be disciplined, accept legal responsibility, or exercise professional judgment in the institutional sense. It can only contribute to task performance and can't authorise the contribution as legal work. In order to better understand how the human/technological relationship can be better understood and governed, clarification of several key terms and concepts is required.

3.2 The professional interface

The **professional interface** is the point in the legal workflow where an AI-mediated subprocess begins to affect work that is subject to professional accountability. It is not the screen, the prompt box or the moment at which output appears. It is the decision point at which an interaction with an AI tool that generates a technological contribution

[16] Davenport (n 3); van der Aalst and van Hee (n 3).

[17] Harry Surden, 'Machine Learning and Law' (2014) 89(1) *Washington Law Review* 87, 99–104; Mireille Hildebrandt, *Smart Technologies and the End(s) of Law: Novel Entanglements of Law and Technology* (Edward Elgar, 2015) 23–37; Emily M Bender, Timnit Gebru, Angelina McMillan-Major and Shmargaret Shmitchell, 'On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?' in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (ACM, 2021) 610.

that must be understood, tested, revised, rejected, contextualised, supervised or authorised before it can properly be treated as part of the lawyer's work.[18]

For that reason, the professional interface shifts with the task. Where AI summarises client material, the interface arises when the lawyer checks that summary against the client's actual instructions and decides how the matter should be framed. Where AI suggests authorities, it arises when the lawyer determines whether those authorities exist, remain current, support the proposition advanced and apply to the facts. Where AI drafts advice, it arises when the lawyer tests the analysis, calibrates risk, revises the language and decides what can responsibly be communicated. The interface therefore brings together the purpose of the legal task, the nature of the technological contribution and the human legal skill and authority needed to convert that contribution into professional action.

3.3 Human legal skill

Human involvement in AI-assisted legal work should therefore be understood as a requirement of capability, not presence. A lawyer does not satisfy the professional requirement simply by remaining somehow "in the loop", viewing an output or approving a document at the end of the process. What matters is whether the lawyer exercises the legal skill required at the point where the technological contribution enters and shapes the work.[19]

Those skills vary with the workflow. At the intake stage, the relevant capability may be listening, instruction-taking and diagnosis of the client's real problem. During factual analysis, it may involve assessing relevance, uncertainty and evidentiary support. At the research stage, it requires source selection, authority evaluation and judgment about jurisdiction, currency and weight. During analysis and drafting, it requires problem-solving, application of law to fact, communication, risk calibration and ethical judgment. Across the whole matter, it also requires professional scepticism, supervision, responsibility allocation and a record sufficient to show how the work was checked and authorised.[20] Workflow identifies the point at which responsibility must be exercised. Skill identifies the human capability required at that point. The practical question is therefore not whether a human has been involved in some formal sense, but whether the right human legal capability has been exercised at the right place in the workflow.

---

[18] Thomas H Davenport, *Process Innovation: Reengineering Work through Information Technology* (Harvard Business School Press, 1993) 5–7; Wil van der Aalst and Kees van Hee, *Workflow Management: Models, Methods, and Systems* (MIT Press, 2002) 2–5; Catherine Gage O'Grady and Casey Sims O'Grady, 'Agentic Workflows in the Practice of Law: AI Agents as Ethics Counsel' (2025) 38 *Georgetown Journal of Legal Ethics* 247; Sirui Han et al, 'Reimagining Legal Fact Verification with GenAI: Toward Effective Human-AI Collaboration' (2026) arXiv:2602.06305.

[19] Legg, McNamara and Alimardani (n 8) 1216–35; *JNE24 v Minister for Immigration and Citizenship* [2025] FedCFamC2G 1314, [20]–[25], [34]; Han et al (n 12); Linna and Linna (n 14).

[20] Donald Feaver, Craig Collins and Shirley Southgate, "What Are Legal Skills? Legal Capability, Professional Formation and the Public Assurance of Competence" Working Paper: Leo Cussen Centre for Law

3.4 Convergence

Convergence occurs when AI-assisted material isn't just reviewed. It must be taken all the way back in the workflow to pints where the lawyer's own reasoning contributes to the transformations that result in the production of professionally usable legal work- or legal artefacts. The point is not that the lawyer adds a name, signature or final approval to a generated artefact. It is that the technological contribution is tested against the matter, revised for legal purpose and absorbed into an accountable professional judgment.[21]

What convergence requires will depend on where AI has entered the workflow. If the system has framed the problem, the lawyer must decide whether that frame captures the client's actual instructions and the legally significant issues. If it has extracted facts, the lawyer must compare the output with the source material and identify uncertainty or omission. If it has suggested authorities, the lawyer must verify not only their existence but their relevance, currency, weight and application. If it has drafted advice, the lawyer must test the reasoning, adjust the level of qualification, consider counterarguments and ensure that the communication is suitable for the client's purpose.

Convergence therefore explains why source checking, although essential, isn't enough. A real citation is not necessarily a usable authority, and a fluent draft is not necessarily sound legal work. The lawyer must supply the legal synthesis that the system cannot provide: judgment about relevance, application, risk, purpose, professional responsibility and the terms on which the work can be authorised.

3.5 Friction

Friction occurs when the human legal skill, control or authority required at the professional interface is missing, inadequate or unable to be shown. It is the point at which AI-assisted work ceases to be properly governed by legal judgment, even though the resulting artefact may appear fluent, complete or professionally acceptable.[22] Friction is therefore not just another name for technological error. The error may originate in the system, but the professional failure lies in the process by which that error, omission or distortion is allowed to shape legal work without the required human intervention. In that sense, friction is a human–technology failure: the system contributes material to the task, but the lawyer does not supply the judgment needed to test, redirect or authorise it.

---

[21] Wanda J Orlikowski, 'Using Technology and Constituting Structures: A Practice Lens for Studying Technology in Organizations' (2000) 11(4) *Organization Science* 404, 407–12; Wanda J Orlikowski, 'Sociomaterial Practices: Exploring Technology at Work' (2007) 28(9) *Organization Studies* 1435, 1437–44; Gerardine DeSanctis and Marshall Scott Poole, 'Capturing the Complexity in Advanced Technology Use: Adaptive Structuration Theory' (1994) 5(2) *Organization Science* 121, 125–31.

[22] Raja Parasuraman and Dietrich H Manzey, 'Complacency and Bias in Human Use of Automation: An Attentional Integration' (2010) 52(3) *Human Factors* 381, 381–90; Christopher D Wickens, Benjamin A Clegg, Alex Z Vieane and Angelia L Sebok, 'Complacency and Automation Bias in the Use of Imperfect Automation' (2015) 57(5) *Human Factors* 728, 728–39.

Some friction is visible. A fabricated authority filed in court exposes a breakdown in research, drafting and assurance. Other forms are much harder to see. An inaccurate AI-generated chronology may quietly distort advice. A compressed summary of instructions may redirect the matter before the substantive legal question is ever reached. A polished draft may conceal from a supervisor that the lawyer has not understood the reasoning, the sources or the limits of the advice. Friction can therefore arise at any point in the workflow. It may begin with an unsuitable task being given to an AI system, confidential information being entered into an inappropriate tool, or an AI-generated frame narrowing the client's problem. It may appear later as missing or misclassified facts, weak authority selection, unsupported legal reasoning or drafting that is fluent but not fit for purpose. It may also arise across the matter, where responsibility is diffused, supervision is only formal, or the file does not allow the lawyer or institution to reconstruct how the work was produced, checked, revised and authorised.

3.6 Five propositions underpinning a workflow governance model

The workflow-governance model can be summarised as five related propositions, each of explaining how technological assistance is brought within, rather than placed outside, professional legal responsibility.

First, professional primacy. AI does not create a parallel professional workflow. It performs technical subprocesses within a legal workflow that remains governed by human professional purpose, judgment and authority.

Secondly, skill-specific control. The need for human control is not generic. At each point where AI materially affects legal work, the lawyer must exercise the kind of legal skill required by that task, its context and its potential consequences.

Thirdly, professional convergence. AI-assisted material becomes professionally usable only when it is drawn back into human legal reasoning: framed for the matter, tested against facts and sources, revised for purpose, contextualised for the client or forum and authorised through accountable judgment.

Fourthly, workflow friction. Professional failure may occur wherever the required skill, control or authority is absent, weak or incapable of being demonstrated. That failure may arise before any final document exists, because AI can shape the task, the facts, the reasoning and the options long before an artefact is produced.

Fifthly, proportionate governance. The safeguards required should respond to the nature of the task, the sensitivity of the information, the degree to which AI shapes the work, the intended use of the output and the professional consequences of error. Governance should therefore be calibrated, not merely imposed in the same form for every AI-assisted task.[23]

Table 1 — The AI-mediated legal workflow

[23] Davenport (n 3); van der Aalst and van Hee (n 3); O'Grady and O'Grady (n 11); Han et al (n 12); Linna and Linna (n 14); Parasuraman and Manzey (n 16D); Wickens et al (n 16D).

Legal workflow stage	Professional purpose	AI-mediated contribution and control points	Human legal skill and judgment required	Where responsibility attaches and friction may arise
Frame and instruct	Define the client's problem, instructions, purpose, authority and the scope of the legal task.	Decide whether AI is suitable for the task; choose an approved tool; control the information entered; test an AI-generated summary or issue frame against the client's actual instructions.	Legal perceptual and professional awareness; interviewing; instruction-taking; problem framing; issue identification; confidentiality and authority judgment.	Responsibility attaches before substantive output is produced. Friction may arise where AI is selected before the task is understood, confidential information is entered into an unsuitable system, instructions are compressed or distorted, or an AI-generated frame narrows the actual problem.
Research	Identify the governing legal framework and obtain authoritative, current and relevant sources.	Define the research question; select an appropriate tool and source environment; constrain jurisdiction and date; verify material authorities through authoritative sources.	Knowledge retrieval and transfer; source selection; legal research; authority evaluation; jurisdictional judgment; recognition of binding, persuasive and adverse authority.	Responsibility attaches to the research frame, tool selection, source selection and authority verification. Friction may arise where fabricated, outdated or irrelevant authorities are accepted, adverse authority is omitted, or confirmation that a case exists is mistaken for confirmation that it supports the proposition advanced.
Analyse	Determine the legal significance of the facts, apply the law and develop a reasoned view of options, risks and likely consequences.	Use AI to organise facts, compare positions, identify possible issues or propose a reasoning structure; test assumptions, factual completeness, classification and application of law to fact.	Complex problem-solving and judgment; factual relevance assessment; analogy and distinction; management of ambiguity; evaluation of competing arguments; risk assessment; independent legal judgment.	Responsibility attaches to issue framing, factual assessment, legal reasoning and conclusions. Friction may arise where plausible reasoning conceals weak classification, missing facts, unsupported assumptions, omitted counterarguments or inadequate application to the client's circumstances.
Draft	Convert legal analysis into an artefact suited to its audience, purpose and legal consequences.	Use AI to generate, restructure or simplify text; check factual and legal accuracy, legal effect, tone, uncertainty, qualifications and	Legal drafting; communication and audience judgment; contextualisation; precision; qualification of uncertainty; translation of legal	Responsibility attaches to the content, form and legal effect of the draft, not merely its grammatical quality. Friction may arise where fluent text is mistaken for sound legal work,

Legal workflow stage	Professional purpose	AI-mediated contribution and control points	Human legal skill and judgment required	Where responsibility attaches and friction may arise
		consistency with the intended purpose.	analysis into an appropriate professional form.	limitations are omitted, the position is overstated, or the artefact is not fit for its intended use.
Advise and act	Reach and communicate a professional judgment about what the client should understand, decide or do.	Use AI to identify options, compare scenarios or produce a preliminary risk summary; independently determine what advice or action is legally, ethically and practically justified.	Strategic and ethical judgment; client-centred decision-making; risk communication; professional independence; matter awareness; decision and approval authority.	Responsibility attaches to the final recommendation, decision and communication. Friction may arise where the lawyer defers to an AI-generated recommendation, communicates work without owning the reasoning, fails to explain uncertainty or allows technological output to displace independent judgment.
Cross-workflow assurance, supervision and record	Establish that the work is reliable, authorised and capable of being professionally defended.	Conduct task, information, source, factual, reasoning and purpose assurance; use independent rather than merely AI-to-AI checking; supervise process and product; record material AI use, verification, revision and approval.	Professional scepticism; self-monitoring; verification; supervision; responsibility allocation; recordkeeping; final authorisation.	Responsibility attaches throughout the workflow. Friction may arise where review is confined to the final artefact, AI is used to validate its own output, responsibility is diffused, the supervisor cannot reconstruct the process, or no reliable evidence exists of the judgments and checks performed.

Table 1 makes two concepts more visible. First, AI affects the whole legal workflow, not merely a drafting stage followed by final verification. It may shape the task before the legal problem has been properly framed, influence the factual matrix before analysis begins, structure the research question before authorities are selected and affect the form of advice before the lawyer has reached a fully independent professional judgment. Secondly, assurance is transversal. It operates throughout the workflow rather than only at its end. The relevant checks therefore include task suitability, information control, factual comparison, source verification, reasoning review, purpose revision, supervision and recordkeeping. The table's purpose is to show that responsibility attaches wherever AI materially enters the legal workflow and that friction may arise whenever the necessary human legal skill, authority or evidence of control is missing. That account provides the

bridge to Part 4, which locates those workflow obligations within the existing architecture of professional duties.

4. Locating professional duties in the workflow

4.1 The Uniform Law foundation

The *Legal Profession Uniform Law* provides the statutory foundation that enables the mapping of professional duties to the workflow concept. Its objectives include promoting the administration of justice and an effective legal profession, ensuring that practitioners are competent and maintain high ethical and professional standards, protecting clients and the public, enabling informed client choice, and promoting efficient, effective, targeted and proportionate regulation.[24] For example, section 296 of the *Uniform Law* defines unsatisfactory professional conduct by reference to conduct connected with legal practice that falls short of the standard of competence and diligence the public is entitled to expect of a reasonably competent lawyer.[25] In AI-assisted practice, that standard should not be assessed only by examining whether a final document was checked. The relevant inquiry is whether the lawyer exercised the judgment required at the workflow points at which AI entered, shaped or compressed the work.

A lawyer may fall short before producing a final artefact. The lawyer may select an unsuitable task, disclose confidential information, accept an inaccurate account of instructions, rely on incomplete factual extraction, adopt unsupported reasoning or fail to supervise the process. Each failure may bear upon competence and diligence. Section 297 becomes relevant where failures are substantial, repeated or serious. Professional misconduct includes unsatisfactory professional conduct involving a substantial or consistent failure to reach or maintain a reasonable standard of competence and diligence.[26] It also includes conduct that would justify a finding that the person is not fit and proper to engage in legal practice.[27] An isolated error and a systemic pattern of unverified reliance, confidentiality failure or misleading AI-assisted work may therefore warrant different disciplinary characterisation.

Section 298 connects the *Uniform Law* with the Uniform Rules. A contravention of the *Uniform Law* or Uniform Rules may constitute unsatisfactory professional conduct or professional misconduct.[28] The Conduct Rules are consequently more than aspirational statements. Their purpose is to assist solicitors to act ethically and in accordance with professional principles established by the common law and the Rules.[29]

---

[24] *Legal Profession Uniform Law* (NSW) s 3.

[25] Ibid s 296.

[26] Ibid s 297(1)(a).

[27] Ibid s 297(1)(b).

[28] Ibid s 298.

[29] *Legal Profession Uniform Law Australian Solicitors' Conduct Rules 2015* rr 2.1–2.3.

4.2 The Conduct Rules

The Conduct Rules operate across the AI-mediated workflow. Rule 4.1.1 anchors the use of AI in the client's best interests. Rule 4.1.3 requires competent and diligent delivery of legal services. Rule 4.1.4 protects professional integrity and independence. Rule 7.1 requires clear and timely advice enabling the client to understand relevant legal issues and make informed choices. Rule 8.1 requires the solicitor to follow lawful, proper and competent instructions. Rule 9.1 protects confidential information. Rule 17.1 requires independent forensic judgment. Rules 19.1 and 19.6 address candour to the court and applicable authority. Rule 22.1 regulates false or misleading statements to opponents. Rule 37.1 requires reasonable supervision.[30] These duties are engaged at different workflow points. Confidentiality is engaged when information is entered, not only when an output is communicated. Independence is engaged when a lawyer decides whether to accept an AI-generated analysis. Candour is engaged when technological material is placed before a court. Supervision is engaged in the process through which AI-assisted work is created, not simply in the final document.

Established professional-responsibility cases also supply the broader foundation. *Clyne v New South Wales Bar Association* illustrates that the lawyer must independently assess whether allegations and forensic action have a proper basis rather than simply pursue the client's immediate objective.[31] *Giannarelli v Wraith* and *D'Orta-Ekenaike v Victoria Legal Aid* emphasise the lawyer's obligations to the administration of justice and the independent professional character of legal work.[32] *Mullins v Legal Services Commissioner* demonstrates the disciplinary significance of professional candour and accuracy in dealings connected with legal practice.[33] These principles remain relevant when technology contributes to the work.

4.3 Court guidance

Australian courts have begun to operationalise these duties through AI-specific guidance. The Federal Court, Supreme Court of New South Wales and Supreme Court of Victoria have addressed matters including task suitability, evidence integrity, confidentiality, verification, disclosure, meaningful human control and ultimate responsibility.[34] The guidance is important not because it creates an entirely separate

---

[30] Ibid rr 4.1.1, 4.1.3, 4.1.4, 7.1, 8.1, 9.1, 17.1, 19.1, 19.6, 22.1, 37.1.

[31] *Clyne v New South Wales Bar Association* (1960) 104 CLR 186; precise pinpoints to be added.

[32] *Giannarelli v Wraith* (1988) 165 CLR 543; *D'Orta-Ekenaike v Victoria Legal Aid* (2005) 223 CLR 1; precise pinpoints to professional independence and duties to the administration of justice to be added.

[33] *Mullins v Legal Services Commissioner* [2006] LPT 012; citation and pinpoints to be checked against the authorised or most appropriate report.

[34] Federal Court of Australia, *Use of Generative Artificial Intelligence Practice Note* (GPN-AI, 16 April 2026); Federal Court of Australia, *Notice to the Profession: Generative Artificial Intelligence Practice Note* (16 April 2026); Supreme Court of New South Wales, *Practice Note SC Gen 23: Use of Generative Artificial Intelligence* (21 November 2024, commenced 3 February 2025); Supreme Court of Victoria, *Practice Note SC Gen 25: The Use of Artificial Intelligence by Court Users* (14 May 2026).

ethical regime, but because it identifies particular workflow points at which existing obligations require attention. It reinforces that AI-generated material must be independently checked, that court users remain responsible for material submitted in proceedings, and that special risks arise where AI interacts with evidence, affidavits, witness material and confidential information.

Table 2 shows how the article's workflow-governance model connects and can be mapped to existing sources of professional responsibility. It maps each stage of the AI-mediated legal workflow against the relevant *Uniform Law* provisions, indicative Australian Solicitors' Conduct Rules, illustrative case law or court guidance, and the specific point at which workflow friction may arise. Its purpose is not to allocate each duty exclusively to a single stage, but to show that established duties of competence, diligence, confidentiality, independence, candour, supervision and client communication become operational at points in the workflow. In this way, the table demonstrates that AI-related professional failure is not confined to final outputs or hallucinated authorities. It may arise earlier, when the task is framed, information is supplied, factual material is classified, legal reasoning is shaped, advice is communicated or supervision and recordkeeping fail to make responsibility visible.

Table 2 — Uniform Law, Conduct Rules and workflow friction

Legal workflow stage	Uniform Law foundation	Indicative ASCR duties	Illustrative case law and guidance	Point of friction made visible
Frame and instruct	Sections 296–298; the competence-and-diligence standard applies to task selection, instructions, confidentiality decisions and the scope of legal work.	rr 4.1.1, 4.1.3, 4.1.4, 7.1, 8.1, 9.1	*Clyne* illustrates the need for independent professional assessment of the basis and purpose of proposed action. Court guidance treats the nature of the task, information involved and intended use as matters requiring prior assessment.	AI may be adopted before the client's problem is understood; instructions may be compressed or distorted; confidential information may be disclosed; or a technological frame may displace the lawyer's diagnosis of the legal task.
Research	Sections 296 and 298; s 297 may apply where reliance on defective or unverified research is substantial, repeated or systemic.	rr 4.1.1, 4.1.3, 4.1.4, 9.1, 19.1, 19.6	*Valu*, *Dayal* and *JNE24* illustrate the consequences of carrying fictitious or inadequately checked authorities into legal proceedings.	Retrieval or generation is mistaken for completed legal research. Authorities may be fabricated, outdated, irrelevant, incorrectly characterised or unsupported by authoritative sources.
Analyse	Sections 296–298; competence and diligence extend to interpretation, relevance, reasoning	rr 4.1.1, 4.1.3, 4.1.4, 7.1, 8.1, 17.1	*JNE24* makes clear that verifying existence is insufficient: the lawyer must assess	Legal fluency is mistaken for legal reasoning. Classification, assumptions, factual

Legal workflow stage	Uniform Law foundation	Indicative ASCR duties	Illustrative case law and guidance	Point of friction made visible
	and application to fact.		whether authority supports the proposition, remains current and applies to the factual matrix.	completeness, counterarguments or application may not be properly tested.
Draft	Sections 296–298; responsibility extends to the accuracy, legal effect and professional adequacy of the artefact produced.	rr 4.1.3, 7.1, 19.1, 19.6, 22.1	*DPP v GR* and *Valu* illustrate how generated errors become professionally significant when incorporated into submissions and other filed material.	Review may focus on grammar, style or the overall conclusion while failing to test quotations, factual propositions, authorities, legal effect and qualifications.
Advise and act	Sections 296–298; final recommendations, decisions and professional actions must meet the required standard.	rr 4.1.1, 4.1.3, 4.1.4, 7.1, 8.1, 17.1, 19.1	*Giannarelli*, *D'Orta-Ekenaike* and *JNE24* reinforce that professional judgment and responsibility cannot be displaced by client pressure, delegation or technological assistance.	The lawyer may communicate an AI-assisted position without independently owning the reasoning, explaining uncertainty or exercising decision and approval authority.
Cross-workflow assurance, supervision and record	Sections 296–298; recurring or serious failures of assurance and supervision may affect disciplinary characterisation.	rr 4.1.3 and 37.1, supported by duties of confidentiality, candour and independence	*Murray*, *JNE24* and *DPP v GR* expose failures of process supervision, distributed checking and responsibility allocation. Court guidance reinforces ultimate human responsibility.	Review is confined to the final artefact; each reviewer assumes someone else checked the foundations; AI is used to validate AI; or no reliable record exists of use, verification, revision and authorisation.

Tables 1 and 2 should be read together. Table 1 identifies the work and skills required at each stage. Table 2 demonstrates how existing law and authority attach to those stages. The mapping is interpretive rather than exhaustive. It does not suggest that each duty belongs exclusively to one stage. Competence and diligence operate throughout. The purpose is to show why general professional duties require workflow-level application in AI-assisted practice.

5. The cases as workflow-friction cases

The emerging decisions are frequently described as hallucination or fake-citation cases. That description is accurate but incomplete. Their deeper significance lies in the professional process through which inaccurate technological material was allowed to enter legal work.

5.1 Research and source friction

In *Valu v Minister for Immigration and Multicultural Affairs (No 2)*, the applicant's representative filed material containing non-existent case citations and purported Tribunal quotations.[35] The Court held that the conduct fell short of the competence and diligence expected by the client and of the practitioner's duty to ensure that the Court was not deceived or misled, even unintentionally. The matter was then referred by the Court to the professional regulator.[36] The relevant failure wasn't that an AI system generated false material. AI entered at the research and drafting process. The material acquired professional significance when the practitioner incorporated it into submissions without verifying the authorities and quotations.

The same structure appears in *Dayal*, where a practitioner tendered an AI-generated list and summary of authorities and acknowledged that the authorities did not exist.[37] Subsequent regulatory action reportedly imposed significant restrictions on the practitioner's practising certificate.[38] The movement from court-facing error to professional regulation illustrates how friction at a workflow point can affect the lawyer's broader status and authority. Comparable overseas decisions reinforce the same gatekeeping principle. In *Mata v Avianca Inc*, lawyers filed fictitious authorities and quotations generated through ChatGPT and were sanctioned.[39] In *Park v Kim*, counsel was referred to a grievance panel after citing a non-existent authority generated by ChatGPT.[40] Those cases emphasise that technological assistance does not remove the human obligation to establish the accuracy and legal tenability of filed material.

5.2 Reasoning and relevance friction

JNE24 v Minister for Immigration and Citizenship is especially important because it distinguishes formal citation checking from substantive legal analysis. The practitioner used Claude for research and Microsoft Copilot to validate submissions.[41] The Court referred the practitioner to the Legal Practice Board of Western Australia and made a

---

[35] *Valu v Minister for Immigration and Multicultural Affairs (No 2)* [2025] FedCFamC2G 95, [1], catchwords.

[36] Ibid [18], [38].

[37] *Dayal* [2024] FedCFamC2F 1166; *Valu v Minister for Immigration and Multicultural Affairs (No 2)* [2025] FedCFamC2G 95, [30]–[32].

[38] Victorian Legal Services Board and Commissioner, 'Statement: Mr Dayal Matter' (Web Page, 2025).

[39] *Mata v Avianca Inc* 678 F Supp 3d 443, 448–9, 453–6, 461–6 (SDNY, 2023).

[40] *Park v Kim*, No 22-2057 (2d Cir, 30 January 2024).

[41] *JNE24 v Minister for Immigration and Citizenship* [2025] FedCFamC2G 1314, [14].

personal costs order.[42] The Court explained that confirming that a case exists is not sufficient. The practitioner must determine whether the authority supports the proposition for which it is cited, whether it has been overturned or distinguished, and how it applies to the factual matrix.[43]

This reasoning closely reflects the concept of convergence. A citation does not become professionally usable because it can be located. The lawyer must evaluate its legal significance. *JNE24* therefore shows that AI-related friction may occur even where an authority is real. The failure may lie in interpretation, relevance, weight or application. The case also undermines the assumption that one AI system can reliably validate the output of another. AI-to-AI checking may be useful as a preliminary control, but it is not independent professional assurance.

5.3 Factual and documentary friction

Murray on behalf of the Wamba Wemba Native Title Claim Group v State of Victoria demonstrates that friction extends beyond legal authorities. A filed summary contained numerous references to anthropological and historical materials, many of which could not be located or were incorrectly cited.[44] Considerable time was expended attempting to identify the sources. The Court required affidavit evidence explaining how the document had been prepared and how the junior solicitor responsible for it had been supervised.[45] The workflow failure concerned factual and documentary support in a matter of substantial significance. The relevant human skills included source evaluation, factual verification, evidentiary awareness and supervision. *Murray* is important because it demonstrates that AI risk is not confined to legal research. Similar problems may affect chronologies, due-diligence reports, documentary summaries, witness preparation, expert-adjacent material and transactional work.

5.4 Drafting, candour and distributed assurance friction

In *Director of Public Prosecutions v GR*, written submissions contained non-existent cases, inaccurate quotations and incorrect legislative references.[46] The material had passed through several professional participants. The failure was therefore not attributable only to the person who first used the technology. It reflected a broader failure of review and assurance. The Court emphasised the importance of the ability to rely upon

[42] Ibid orders 1–2, [35]–[43].
[43] Ibid [20]–[25], [34].

[44] *Murray on behalf of the Wamba Wemba Native Title Claim Group v State of Victoria* [2025] FCA 731, [3]–[4].
[45] Ibid [5].
[46] *Director of Public Prosecutions v GR* [2025] VSC 490, [61]–[80].

the accuracy of submissions made by counsel.[47] The case demonstrates that several layers of review do not necessarily amount to meaningful supervision or verification. Each person may agree with the overall conclusion while assuming that another person has checked the foundations.

This is distributed assurance friction. Responsibility is present formally but diffused operationally. The same concern appears in the English decision in *Ayinde v London Borough of Haringey; Al-Haroun v Qatar National Bank QPSC*, which considered suspected or established reliance on false authorities, wasted costs, possible contempt and regulatory referral.[48] The Court emphasised independent verification and candid response once concerns arise. The professional obligation therefore extends beyond preventing the initial error. It includes how the lawyer investigates, discloses and corrects it.

5.5 What the cases establish

The cases support three propositions. First, courts are not treating AI error as a separate technological accident. They analyse the conduct through familiar obligations of competence, diligence, candour, supervision, costs responsibility and regulatory accountability. Second, verification is substantive rather than formal. It requires more than locating a citation. The lawyer must assess authenticity, currency, relevance, weight, factual application and the accuracy of quotations or propositions. Third, supervision is process-based. Where AI has materially assisted the preparation of legal work, meaningful supervision may require an understanding of how the material was produced, what sources were used, what checks were performed, what was revised and who authorised the final action.

The case law remains predominantly court-facing because filed documents are visible, preserved and capable of judicial scrutiny. That does not limit the model to litigation. It means that courts are exposing the points at which workflow friction becomes public. A fabricated authority in submissions is visible friction. An inaccurate chronology used in private advice may be invisible friction. A distorted AI summary of instructions may be upstream friction. A negotiation letter containing an unsupported factual assertion may

[47] Ibid [79].

[48] *Ayinde v London Borough of Haringey; Al-Haroun v Qatar National Bank QPSC* [2025] EWHC 1383 (Admin); pinpoints to independent verification, wasted costs, contempt thresholds and regulatory referral to be added.

be adversarial friction. The professional structure is the same even where the failure never reaches a court.

6. Beyond litigation: the AI-assisted letter of advice

A letter of advice demonstrates the broader reach of workflow governance. The final letter is only one artefact. The professional work includes instructions, fact clarification, issue identification, legal research, analysis, strategic judgment, drafting, verification, communication and recordkeeping. AI may contribute to task completion at each stage.

Table 3 — AI-assisted advice workflow

Workflow point	Possible AI contribution	Required human legal skill	Possible friction
Instructions and problem framing	Summarise intake notes, emails or documents; suggest an initial issue frame.	Interviewing, listening, instruction-taking, problem diagnosis, authority and confidentiality judgment.	The AI summary omits or distorts the client's actual concern; the lawyer accepts a narrowed or incorrect problem frame.
Fact clarification	Extract facts, identify dates, generate a chronology or classify documents.	Factual relevance assessment, evidentiary awareness, source checking and identification of uncertainty.	Material facts are omitted, duplicated, confused or treated as established without adequate source support.
Issue identification	Suggest claims, defences, statutory provisions or legal categories.	Issue spotting, legal classification and recognition of missing information.	The generated categories direct the analysis towards an obvious but non-decisive issue or omit the controlling one.
Legal research	Suggest authorities, provide summaries or generate propositions of law.	Source selection, authority evaluation, jurisdictional judgment and legal research.	False, outdated, weak or irrelevant authority is relied upon; a secondary AI synthesis is treated as authoritative law.
Legal analysis	Generate an initial reasoning structure, compare positions or suggest risks.	Complex problem-solving, analogy, distinction, management of ambiguity and application of law to fact.	Plausible reasoning conceals unsupported assumptions, missing counterarguments or inadequate factual application.
Advice strategy	Identify options, compare scenarios or generate a preliminary risk matrix.	Strategic judgment, client-centred advice, ethical judgment and practical matter awareness.	Advice is technically plausible but impractical, disproportionate, inconsistent with instructions or inadequately qualified.
Drafting and communication	Generate a first draft, restructure text or simplify language.	Drafting, audience judgment, precision, risk communication and professional tone.	Fluent text is sent without adequate revision; qualifications, limitations or uncertainty are obscured.

Workflow point	Possible AI contribution	Required human legal skill	Possible friction
Assurance and record	Identify inconsistencies or generate an AI-use summary.	Professional scepticism, independent verification, supervision, authorisation and recordkeeping.	AI is used to validate its own work; the file does not show what was generated, checked, changed and approved.

The example demonstrates why "AI drafts; lawyer approves" is inadequate. The lawyer's responsibility is not confined to editing the final letter. It extends to governance of each interface through which AI affects the work. It also shows why upstream friction may be more difficult to detect than a false citation. If an AI system incorrectly frames the client's problem, every later stage may be performed competently in relation to the wrong question. If a chronology omits a decisive fact, otherwise sound legal analysis may be built on an incomplete factual matrix. If the AI-generated structure causes the lawyer to overlook an alternative legal characterisation, final proofreading will not reveal the omission. Workflow governance therefore requires more than error correction. It requires maintenance of professional purpose throughout the task.

7. Supervision, law-practice governance and authority

The workflow model has direct implications for supervision. Rule 37.1 requires the solicitor with designated responsibility for a matter to exercise reasonable supervision over solicitors and other employees engaged in providing legal services on that matter.[49] In AI-assisted practice, reasonable supervision cannot always be confined to reviewing the final artefact. AI may also conceal weaknesses in underlying capability. A polished document can mask the fact that the lawyer has not identified the relevant law, distinguished authority, understood the factual basis or internalised the reasoning. A supervisor who sees only the finished text may therefore be unable to tell whether the junior lawyer performed the legal work, understood the technological contribution, checked the material independently or merely accepted a fluent output that appeared professional.

Meaningful supervision in AI-assisted practice must therefore look beyond the final artefact and examine the process by which the work was produced. The supervisor may need to know whether AI was used, where it entered the workflow, what information was supplied, what the tool contributed, which sources were checked independently, what assumptions or reasoning were tested, what was revised or rejected, and whether the lawyer can explain and defend the work without relying on the AI output.[50] This does not mean that every minor or mechanical use of AI requires a detailed audit. Supervision

[49] *Legal Profession Uniform Law Australian Solicitors' Conduct Rules 2015* r 37.1.

[50] Megan Ma and John Nay, 'A Supervisory AI Agent Approach to Responsible Use of GenAI in the Legal Profession' *CodeX: The Stanford Center for Legal Informatics* (Web Page, 25 January 2024); Dazza Greenwood, 'Task Force on Responsible Use of Generative AI for Law' *MIT Computational Law Report* (Web Page, 28 February 2023); Matthias Braun, Hannah Bleher and Patrik Hummel, 'Human Control of AI Systems: From Supervision to Teaming' (2024) *AI and Ethics* 1.

should remain proportionate. Its depth should depend on the nature of the task, the sensitivity of the information, the experience of the person using the tool, the extent to which AI shaped the work and the consequences of error.[51]

Institutional responsibility must likewise move beyond generic statements that AI may be used "responsibly". A law practice requires systems capable of deciding which tasks may be assisted by AI, which tools are approved, what information may be entered, what assurance is required, who has authority to review or approve the work and what record should be kept.[52] This is usefully understood as a question of authority. Four forms of authority are relevant: authority to perform or assist the task, authority to review the work, authority to make the substantive legal or strategic decision, and authority to approve communication, filing or other professional action.

AI may assist performance, but it cannot hold decision or approval authority in the professional sense. Those forms of authority remain human and institutionally allocated. The distinction matters because a legal task may appear supervised if a person is somewhere in the process, even though the person has not exercised the authority or skill required at the point where the technology shaped the work.

The issue becomes more acute in complex or agentic workflows. As systems retrieve information, compare materials, draft text and pass outputs between components, the point of human intervention may become less visible. Governance must therefore specify not only that a human is involved, but what authority that person exercises, what legal skill is required, and what evidence exists that the work was checked, revised and authorised. Recordkeeping supports that allocation of responsibility. Not every use of AI should require a formal certificate or extensive file note. But where AI materially contributes to substantive legal work, a proportionate record should make the process reconstructable. It may identify the tool used, the purpose of the use, the type of information supplied, the material output relied upon, the independent checks undertaken, significant revisions or rejections, the person who reviewed the work and the person who authorised the final action.

The present law does not necessarily impose an express AI-specific recordkeeping duty in every context. The point is more modest. An adequate record may be necessary to demonstrate competence, diligence, supervision, confidentiality compliance and responsibility where the process is later questioned. It also helps prevent responsibility from becoming formally present but operationally diffuse: everyone may have touched the work, but no one may be able to show who exercised judgment, what was checked, what was changed and who authorised the result.

8. Regulation and reform

[51] Yuvaraj, 'The Verification-Value Paradox' (n 13)

[52] O'Grady and O'Grady (n 11); Murray (n 10); Thomson Reuters Institute, *2025 Generative AI in Professional Services Report* (Report, 2025); Law Society of New South Wales, *A Solicitor's Guide to Responsible Use of Artificial Intelligence* (publication details to be confirmed); Cathryn Urquhart, 'Generative AI: Supervision and Risk Management' *Law Society Journal* (Web Page, 27 May 2025).

As discussed above, existing professional duties appear capable of governing AI-assisted legal practice. Their generality is part of their strength. The introduction of a new technology does not require abandonment of competence, diligence, candour, confidentiality, communication, supervision or independence.

However, the problem is more operational than regulatory. General duties do not always specify how they apply where technological systems perform or shape parts of legal tasks. Current court guidance and case law have understandably focused on court-facing failures: fabricated authorities, false quotations, inaccurate evidence and misleading submissions. These are the most visible and institutionally urgent problems. But they represent only part of the workflow.

Friction may arise in client intake, factual extraction, issue framing, due diligence, research, advice, contract drafting, negotiation, settlement, internal review and recordkeeping. Many such failures will never become visible to a court or regulator. They may nonetheless affect client interests and professional responsibility. Future guidance may need to move beyond general statements that lawyers remain responsible for AI use and seek to operationalise responsibility through workflow.

8.1 Distinguish workflow points and task types

Guidance should begin by distinguishing the workflow point and task type involved. AI use for administrative formatting is not equivalent to AI-assisted factual extraction, legal research, evidence preparation, advice drafting or court submissions. The professional purpose of the task, the sensitivity of the information, the degree to which AI shapes the result, the opportunity for independent checking, the intended audience and the consequences of error all affect the level of control required.[53]

This approach would support targeted and proportionate regulation consistent with the objectives of the *Uniform Law*. It would also avoid two unhelpful extremes: treating all AI use as though it presents the same professional risk, or assuming that risk only arises when an output is filed in court. The better view is that guidance should identify the workflow conditions that make AI assistance material to professional responsibility.

8.2 Require meaningful human legal control

Guidance should also move beyond a bare requirement that a human remain involved. Human presence is not the same as meaningful human legal control. The relevant question is whether the person involved possesses and exercises the legal skill and

[53] National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (Report, January 2023) 20–6; ISO/IEC, *ISO/IEC 42001:2023 Artificial Intelligence — Management System* (Standard, 2023); European Commission High-Level Expert Group on Artificial Intelligence, *Ethics Guidelines for Trustworthy AI* (Guidelines, 8 April 2019) 16–20; Legal Profession Uniform Law (NSW) s 3.

professional authority required at the point where the technology affects the work.[54] Meaningful human legal control should require, as appropriate:

- framing the professional task;
- controlling information inputs;
- evaluating sources and facts;
- testing reasoning;
- revising the output for legal purpose;
- exercising independent judgment;
- supervising the process; and
- authorising the professional action.

The required form of control will differ across tasks. The governing principle is that technological assistance must not displace the human capability required by the legal purpose.

8.3 Develop differentiated assurance standards

Verification should be understood as a family of assurance practices rather than a single final check. The different forms of assurance can be organised as a form of AI use record that might appear as follows:

Assurance type	Question addressed
Task assurance	Whether AI is appropriate for the work.
Information assurance	Whether confidentiality, privilege, accuracy and permitted use of inputs have been protected.
Source assurance	Whether sources are authentic, current, jurisdictionally appropriate, weighty and relevant.
Factual assurance	Whether factual propositions correspond with source material and uncertain facts are properly identified.
Reasoning assurance	Whether assumptions, classifications, counterarguments and application have been tested.
Purpose assurance	Whether the artefact is appropriate for its audience, legal function and professional consequences.
Responsibility assurance	Whether supervision, authority, approval and the evidentiary record of the process are sufficient.

This recording framework reflects the reasoning in JNE24: confirming that a source exists is only one part of professional legal verification. In AI-assisted practice, verification must be understood more broadly as assurance that the material has been made professionally usable. A real case may still be irrelevant, outdated, distinguishable or

[54] Braun, Bleher and Hummel (n 43A); European Commission High-Level Expert Group on Artificial Intelligence, *Ethics Guidelines for Trustworthy AI* (Guidelines, 8 April 2019) 16–19; NIST, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (Report, January 2023) 20–6; American Bar Association, *Formal Opinion 512: Generative Artificial Intelligence Tools* (29 July 2024) 2–7.

incapable of supporting the proposition for which it is cited. A factual summary may be grammatically clear but unsupported by the underlying material. A draft may be coherent but based on an assumption the lawyer has not tested. The point is that professional verification is not a binary question of whether an output is "true" or "false". It is a structured legal inquiry into whether the material is suitable for the task, grounded in proper sources, factually accurate, legally reasoned, fit for purpose and capable of being responsibly authorised.[55]

8.4 Update supervision standards

Supervision guidance should address AI-assisted process as well as product. The degree of supervision should depend upon the junior lawyer's capability, the nature of the task, the system used and the consequences of error. Supervisors should not be expected to reproduce every step. They should, however, be able to determine whether the person performing the work understood the legal task, used an appropriate tool, checked material independently and can explain the reasoning. The cases demonstrate why. *Murray* exposes failures in the preparation and supervision of source-based work. *DPP v GR* shows that multiple reviewers do not provide assurance where each assumes someone else has checked. *JNE24* demonstrates that apparent validation through another technological system is not equivalent to independent legal review.

8.5 Clarify proportionate recordkeeping

Where AI materially assists substantive legal work, guidance should encourage or require a proportionate record of use. The record should not become a bureaucratic end in itself. Its function is to make responsibility and assurance reconstructable. The detail required should reflect the significance of the task and the risk posed by the use. Future guidance might provide that a solicitor who materially uses AI in legal work must take reasonable steps to:

1. identify the professional task and workflow point;
2. use an appropriate and authorised system;
3. protect confidential and privileged information;
4. ensure that the relevant human legal skill is exercised;
5. independently verify and revise material as the task requires;
6. supervise the process and allocate authority clearly; and
7. maintain a proportionate record of material use, checking and authorisation.

That formulation would not create a new ethical universe. It would translate existing professional duties into operational requirements for AI-assisted practice. Solicitor conduct law supplies the duties. Workflow governance supplies the location. Human legal skill supplies the capability. Assurance and authority supply the controls through which technological assistance becomes professional legal work.

[55] *JNE24 v Minister for Immigration and Citizenship* [2025] FedCFamC2G 1314, [20]–[25], [34].

9. Conclusion

AI-assisted legal practice is not best understood as a contest between lawyers and machines. It is a reorganisation of task performance within the professional legal workflow. Generative AI can produce material that resembles completed legal work while compressing or obscuring the professional process through which that work would ordinarily be produced. The central risk is therefore not only that the system may generate an error. It is that the appearance of completion may displace or conceal the human legal work required to frame, evaluate, revise and authorise the result.

Workflow governance provides an architecture for addressing that risk. There remains one normatively primary legal workflow. AI may perform or assist technical subprocesses within it. At each point where that assistance materially affects legal work, task-specific human legal skill and authority are required. Convergence occurs when technological material is transformed through professional synthesis into legally usable work. Friction occurs when the required skill, control or authority is absent, inadequate or incapable of being evidenced. The Uniform Law and Conduct Rules already provide the normative foundation. The emerging court guidance and decisions show those duties becoming concrete through requirements of competence, independent verification, candour, supervision and ultimate human responsibility. Properly understood, the cases are not merely warnings against hallucinated citations. They are examples of workflow failure.

The model also extends beyond litigation. Friction may occur in instructions, factual analysis, advice, drafting, negotiation, supervision and recordkeeping. Court-facing cases reveal the failures that become public. Workflow governance identifies the broader structure through which less visible failures may affect clients and professional responsibility. The future direction is therefore not a general prohibition on AI or a vague insistence that humans remain in the loop. It is the development of task-specific, skill-based and proportionate human legal control. Technology may assist performance. Human legal judgment, assurance, authority and responsibility remain the conditions that make the work professional.