



Artificial Intelligence and Judging

Remarks by the Honourable Chief Justice Niall to the Judicial College of Victoria, Melbourne

16 September 2026

Good evening and welcome. Thank you to the Judicial College of Victoria for organising this event, and to Kate Clark for your introduction. I welcome fellow members of the judiciary to this seminar.

I would like to start by echoing Kate's acknowledgment the Traditional Owners and Custodians of the land on which we gather, the Peoples of the Kulin nations. I pay my respects to their Elders and Knowledge Holders and to any First Nations people joining us today.

When speaking about AI it is hard to know where to start. AI, or at least talk of AI is everywhere. No conference is complete without a session on AI and everyday there is something in the news or in social media about AI. Should I start with the upside or the downside?

There are risks associated with AI. From the cataclysmic which puts the control of biological and other warfare in the hands of many or gives every malign regime or actor the ability to control or effect public debate to the more prosaic but important localised effects, including those on the courts.

AI can bring significant benefits. Information and knowledge are crucial to human welfare and prosperity and AI is incredibly adept at marshalling information and presenting even complex things in a digestible form and at a breathtaking pace.

For our purposes, AI has the potential to reduce the time and cost of coming to court, make judging easier and more efficient and reduce the time we spend on things that matter less. Of course, in the context of our work most of the claimed efficiency involves reducing the time people, that is judges and our staff, spend on different tasks. There is no doubt that using AI involves a form of substitution: we get AI to do some or part of the tasks that we, or our staff, would have done. From an economic point of view, spending less human time on a task means that you need fewer people for the same output or more output from the same people.

In looking at the adoption of AI should we focus on those tasks that we don't like doing, those tasks that we could do more efficiently by using AI or those tasks that will increase the output of the courts? And how will we measure these things?

How we answer those questions depends very much on where one starts. For my part, the proper starting point is our function and role as judicial officers and the nature of judicial power. Critical to our conception of judicial power is a judicial decision according to law arrived at by the application of the judicial method. As everyone here would know, we would be a lot more efficient and faster in making a decision if we did not have to hear from each side. We might be a lot more efficient and quicker if we asked Claude or ChatGPT to write our judgments.

In looking at the costs and benefits of AI, it would be wrong to ignore the fact that the two greatest threats to the sustainability of the justice system are costs and delay. And that we have tended to see consistency in decision making and the avoidance of idiosyncratic decision making as important goals. AI may well deliver faster, cheaper and more consistent decision making.

It is easy to see the problem of AI as primarily being one of verification. An inaccurate summary of the evidence or legal principle is worthless. There is no saving in getting a judge to make a wrong decision faster. So the first metric of AI use must be accuracy. The cost of verification is a significant burden and at least in the early stages of adoption has balanced out many of the promised time and cost savings. But these things are improving and will continue to improve. The tools constantly increase in their capacity and most of the tools, including the legally specific ones, clearly identify the sources from which the text is generated to ease the burden of checking.

But to my mind a more fundamental problem occurs where the judge is replaced by AI regardless of the accuracy of the result. Judging remains a fundamentally human exercise. I should say that not everyone thinks AI has no role to play in decision making. Sir Geoffrey Voss, the Master of the Rolls has written extensively on the topic and has expressed the view that lawyers and judges have no choice but to adopt, indeed wholeheartedly embrace, AI for efficient access to justice. He regards the cost benefits as so great that AI will inevitably intrude and in some occasions, perhaps many, replace judicial decision-making with the key issue being the consent of the parties.¹

¹ [Sir Geoffrey Vos, 'The Implications of the Development of Artificial Intelligence on the Judiciary and the Legal Profession' \(John Lehane Memorial Lecture, Supreme Court of New South Wales, 18 August 2026\).](#)

My own view takes us beyond accuracy and efficiency. Courts are not only producers of legally correct outcomes. They are institutions through which public power is exercised under conditions of independence, procedural fairness, openness and reasoned justification. The legitimacy of a decision rests partly on its result, but also on the process by which that result becomes the decision of a court.

We should approach artificial intelligence as a question of institutional design. We should not begin with a product and search for work it can perform. We should begin with the judicial function, and ask what aspect of judicial work could be improved through AI while preserving judicial independence, procedural fairness, transparency and accountability. The task being to ensure that human responsibility and accountability is not lost nor diminished.

That approach is neither resistant nor defensive; it leaves substantial room for innovation. But it insists that artificial intelligence be introduced into courts as an instrument of judicial values, not as a force that quietly redefines them. Central to this vision is that we make decisions about the adoption of AI, including the tools that we use and how we use them, from an institutional not individual perspective.

The correct approach gives primacy to purpose and governance.

Judicial power is more than the production of an answer

The distinction begins with the nature of judicial adjudication. Kitto J described judicial power as an inquiry into the law and the facts, followed by the application of the law as determined to the facts as determined. Each part of that formulation matters. There is an inquiry, conducted according to law. There are facts found from admissible material. There is law identified

through a disciplined legal methodology. And there is an application of one to the other by the person entrusted with the power to decide.

A generative model does not participate in that process as a judge does. It has no duty to attend to the parties' cases, avoid an unargued basis of decision or apply a burden of proof. It can imitate legal reasoning without being situated within the relationships of authority and obligation that make reasoning judicial.

That is why the constitutional and institutional setting cannot be treated as an external constraint to be applied after a tool has demonstrated its capability. It supplies the criteria by which capability is judged. Section 24 of the Victorian Charter protects a decision by 'a competent, independent and impartial court or tribunal after a fair and public hearing'.² This right is to adjudication by an institution of a particular kind, acting through a process of a particular quality.

The latest Guide to Judicial Conduct accordingly states the boundary with useful directness: 'Judges must not use Generative AI for judicial decision-making'.³ That prohibition is not an assertion that judges reason without assistance. Judicial work has always been supported by associates, librarians, researchers, registrars and technology. The point is that assistance takes place within structures of responsibility that preserve the judge as the author of the decision. Authorship here does not mean typing every word. It means intellectual ownership of the path from issue to conclusion. The judicial officer must identify the questions that require decision, engage with the evidence and submissions, determine the relevant law, confront the

² *Charter of Human Rights and Responsibilities Act 2006* (Vic) s 24

³ [Council of Chief Justices of Australia and New Zealand, *Guide to Judicial Conduct* \(Australian Institute of Judicial Administration, 4th ed, March 2026\) 13.](#)

considerations that point in different directions, and accept responsibility for the orders and reasons. Those elements cannot be satisfied by approving a text when the reasoning was formed elsewhere. Verification is not a substitute for judicial decision making.

Judgment writing is often part of the process itself. A provisional conclusion may fail when exposed to reasons; a paragraph may reveal that two propositions cannot coexist; a peripheral argument may become decisive when answered. Writing is sometimes the means by which judgment is tested.

Sir Harry Gibbs said that reasons should demonstrate that ‘an honest, careful and conscientious effort has been made to do justice’.⁴ This describes a quality of attention. If artificial intelligence removes friction from writing, it may help expression. If it removes the friction through which a judge discovers a weakness in the analysis, it may diminish judgment. The same feature can be assistance in one context and displacement in another.

The first governing proposition is therefore one of judicial authorship: technology may support the formation and expression of judgment, but it must not interrupt the chain of responsibility between the judicial officer, the material, the reasoning and the decision.

Before coming to how we seek to protect judicial authorship through governance it is useful to dwell on some risks.

⁴ [Sir Harry Gibbs, ‘Judgment Writing’ \(Speech, The Judges’ Conference, 19 January 1993\).](#)

Risks

The first and most profound risk is to institutional integrity. Gleeson CJ said that the most important measure of the court system is the extent of public confidence in its independence, integrity and impartiality.⁵ It is part of the authority by which judgments are accepted and obeyed, including by those who may disagree with them. Chief Justice Bathurst has preferred to talk of trust rather than confidence but the underlying themes are the same.⁶

Trust and confidence describe a relationship between the branches of government and between the community and the courts. It depends on decisions being made by independent judicial officers who bring their own abilities and judgment to each matter they decide. This paradigm is threatened if it is known or suspected that a judgment has been written by AI.

In recent times, text published by consultants, lawyers, and newspapers have been identified as having been written by AI. In 2025, Deloitte was commissioned by the Federal Department of Employment and Workplace Relations to review the Departments compliance framework and IT systems. Subsequently it admitted that it had used AI in preparing the report which contained a number of errors and hallucinations. It provided a partial refund to the Commonwealth when the errors were exposed.⁷

It is not only the presence of hallucinations that might point to AI use. There are now a number of AI detectors such as Pangram, GPTZero, Winston AI and Copyleaks that claim that they can detect AI written text. Research suggests that AI detection tools still have considerable

⁵ [Murray Gleeson, 'Public Confidence in the Judiciary' \(Speech, Judicial Conference of Australia, 27 April 2002\).](#)

⁶ [T F Bathurst, 'Trust in the Judiciary' \(Opening of Law Term Address, Law Society of New South Wales, 3 February 2021\) \[11\]–\[13\].](#)

⁷ [Edmund Tadros and Paul Karp, 'Deloitte to Refund Government, Admits Using AI in \\$440k Report', *Australian Financial Review* \(online, 5 October 2025\).](#)

limitations, with studies finding that they can be overly sensitive and, as a result, are not consistently reliable or accurate.⁸ In 2025, Pangram was used to accuse Jamir Nazir the author of *The Serpent in the Grove* a prize-winning story in the prestigious Commonwealth Short Story Prize of having used AI to write the story. It caused a significant backlash and led the main sponsor to abandon the prize. He denies having used AI and observed that with AI it has become impossible to trust each other.⁹

The unconstrained use of AI by judges also presents challenges to the adversarial nature of litigation. The common law system involves independent judicial officers resolving the parties' disputes according to law. It is not inquisitorial and it is for the parties to identify the issues that they wish to litigate and the arguments they want to run in support of their case. The parties and their legal advisers are bound to assist the court reach a just outcome and are bound by overarching obligations. Subject to those important restraints, the parties chose the ground on which they want to fight. Using AI to identify the issues or potential disputes that arise in the material provided by the parties risks altering our role.

Finally, there are risks to confidentiality and privacy.

How then do we manage these risks while harnessing the benefits that the technology can bring.

The emerging Victorian settlement

The Supreme Court's initial guidance in 2024 responded to the immediate risks presented by generative AI in litigation. It placed responsibility for filed material on the litigant and

⁸ See eg, [Mike Perkins et al, 'Detection of GPT-4 Generated Text in Higher Education: Combining Academic Judgement and Software to Identify Generative AI Tool Misuse' \(2024\) 22\(1\) *Journal of Academic Ethics* 89.](#)

⁹ ['Did AI Write This? Its Getting Harder to Tell', *Financial Times* \(online, 29 August 2026\).](#)

practitioner, where it belongs. The Victorian Law Reform Commission then undertook the first law reform inquiry in the country devoted to artificial intelligence in courts and tribunals. Its report made 30 recommendations following consultation across the justice system.¹⁰

The Commission's most important contribution was locating AI within the values of the justice system. It proposed eight principles: access to justice; impartiality and fairness; accountability and independence; privacy and data security; transparency and open justice; contestability and procedural fairness; efficiency and effectiveness; and human oversight and monitoring.

Those principles now inform Supreme Court Practice Note SC GEN 25.¹¹ The relationship between the principles matters. Efficiency is bounded by fairness. Transparency enables contestability. Privacy protects individuals and proceedings. Human oversight connects accountability to a person capable of acting. The framework defines worthwhile innovation by reference to principle.

For court users, the Practice Note adopts a position of permission and responsibility. Artificial intelligence may be used, subject to law, professional obligations and any direction of the Court. But the user remains accountable for everything submitted. Meaningful human control requires verification against original authorities and source material, attention to currency and jurisdiction, and consideration of incompleteness and bias.¹² One model cannot verify another merely by repeating its confidence.

¹⁰ [Victorian Law Reform Commission, *Artificial Intelligence in Victoria's Courts and Tribunals* \(Report No 49, October 2025\) xix.](#)

¹¹ [Supreme Court of Victoria, *Practice Note SC Gen 25: The Use of Artificial Intelligence by Court Users* \(14 May 2026\) \[4.1\].](#)

¹² *Ibid* [6.1]-[6.4].

The Practice Note also recognises that confidentiality is not simply a matter of removing names. Court material may reveal intimate personal facts, commercial information, protected identities, investigative methods or the substance of judicial deliberation. Before information is entered into a system, the user must know where it will be processed, who may have access, how long it will be retained and whether it may be used to train a model. Public tools are not appropriate repositories for sensitive material.

The accompanying guidelines for judicial officers permit supportive uses such as locating and organising material, producing an initial chronology or summary, assisting research and proofreading, provided those uses do not replace reading the evidence and submissions, finding facts, applying law or exercising discretion.¹³ Case-related work is confined to Court-approved tools.¹⁴

The Courts Council has established an AI Committee to develop a court-specific framework for acquisition and deployment across Victorian courts and VCAT.¹⁵ That institutional work is indispensable. Questions about security, assurance, procurement, record-keeping, incident response, equality of access and disclosure cannot be resolved responsibly through individual experimentation. Judicial independence includes a capacity for courts to govern the systems through which judicial work is undertaken.

The Victorian position is therefore neither a moratorium nor an open licence. It is a structured allocation of responsibility: adjudication remains human; supportive uses are assessed by function and risk; case material remains within approved environments; court users remain

¹³ [Supreme Court of Victoria, *The Use of Artificial Intelligence by Judicial Officers* \(14 May 2026\) \[3.4\].](#)

¹⁴ [Ibid \[3.5\].](#)

¹⁵ [Richard Niall CJ, 'Remarks at the Victorian Justice Leaders AI Forum' \(Speech, Victorian Justice Leaders AI Forum, Melbourne, 26 May 2026\).](#)

accountable for their documents; and institutional deployment is subject to judicially led governance.

Governance as an aspect of judicial independence

These developments lead to a further proposition: AI governance is an aspect of judicial independence.

Independence is usually discussed in relation to appointment, tenure, remuneration and freedom from interference in individual cases. But it also has an institutional dimension. If the systems through which court information is organised, searched and presented are designed by vendors or executive agencies without effective judicial control, choices about relevance, visibility and workflow may be embedded outside the judiciary even though formal decisional authority remains with the judge.

Before acquisition, a court should define the problem and ask whether artificial intelligence is needed at all. A better form, clearer process, or conventional search tool may solve the problem with less risk. Procurement should address provenance, security, retention, model changes, subcontractors, accessibility, audit rights, incident reporting and the ability to exit without losing control of court information.

Testing should use difficult court material: poor-quality scans, unusual names, overlapping speech, minority languages and contested accounts. Aggregate accuracy can conceal systematic failure. Results should be examined by task and affected group, against benchmarks fixed in advance.

Deployment is not the end of assurance. Models and interfaces change. Work practices adapt around them. Users may extend a tool beyond the purpose originally approved. The National Institute of Standards and Technology describes AI risk management as a continuing process of governing, mapping, measuring and managing.¹⁶ For courts, that cycle must include monitoring workload, unequal effects, incidents, professional capability and the practical ability of parties to contest material uses.

The strongest case for artificial intelligence is that it may release capacity for the work that is irreducibly human: listening, engaging with a difficult argument, explaining procedure, recognising the significance of an unusual circumstance and writing reasons that a person can understand. But that public dividend will not arise automatically. Saved time must be deliberately reinvested in the quality and accessibility of justice.

Public trust, as I said at the Victorian Justice Leaders AI Forum, is earned through ‘transparency, understanding and experience’.¹⁷ Experience should make our boundaries more precise. Understanding should make us less susceptible to both technological enthusiasm and technological fear. Transparency should allow the community to see that the judiciary remains responsible for the conditions under which its power is exercised.

Introduction of Court Services Victoria artificial intelligence framework

Further to my comments about judicial governance, I am pleased to announce tonight that Courts Council has taken an important step towards the responsible adoption of artificial

¹⁶ [National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework \(AI RMF 1.0\)* \(NIST AI 100-1, January 2023\).](#)

¹⁷ Niall (n 15) 10.

intelligence, with the release of a Governance Model for AI, comprising the AI Charter, AI Assurance Framework and Responsible Use of AI Policy. This framework sets out the terms under which artificial intelligence tools and systems are used across the courts, VCAT and Court Services Victoria. In summary:

- The AI Charter establishes a shared strategic vision for the use of AI in the courts and VCAT. It outlines a common approach to innovation, balancing the opportunities of AI with the need to preserve judicial independence, protect public trust and manage risk.
- The AI Assurance Framework establishes the assessment protocols through which new AI tools and systems must pass before they can be approved for use. By providing a practical pathway for assessing and approving AI tools and services, we can ensure that all AI initiatives align with agreed principles and governance requirements.
- The Responsible Use Policy sets out the conditions under which individual use of approved AI tools or systems is endorsed. This sets clear expectations for staff and the judiciary, supporting the safe, transparent and secure use of approved AI tools.

These documents will be published shortly together with the register of AI tools that have been approved for use which will include two general purpose tools Copilot Chat and VicGov Chat. They are the product of strong collaboration between jurisdictions and the hub that has been created within Courts Services Victoria to support our commitment to using AI responsibly and in a manner that enhances justice delivery and preserves public trust and confidence. These documents sit alongside the judicial guidelines that have been issued by the jurisdictions and the overarching national approach articulated through the Guide to Judicial Conduct that prevents the use of AI for judicial decision making.

There are two key points that I want to emphasise arising out of this new governance model. The first is the importance of maintaining the distinction between augmentation and replacement of the judicial role. The second is the important role of pilots rather than the ad hoc or individual use of tools and the final is how we might measure success.

Augmentation not delegation

The first point concerns the distinction between assistance and delegation. This distinction is easy to state and harder to apply. A tool that locates every reference to a contractual term is plainly assisting. A tool that recommends how the term should be construed is moving closer to delegated analysis. Between those examples lies a wide field: summaries, chronologies, comparisons of evidence, legal research and revisions to draft reasons.

The boundary cannot be drawn by asking whether a human remains nominally ‘in the loop’. Review may be robust or little more than acquiescence in a persuasive output. Meaningful control requires expertise, time, access to sources and independence from the system's framing. To repeat the question is one of demarcation not verification.

I accept that the distinction can be an elusive one. The object is to ensure that the decision is that of the judge. That he or she brings an independent mind to the facts and law. The principle is broadly drawn to reflect the fact that we all approach the task of decision making differently. For some of us, preparing a chronology is our way into the case and informs our understanding of how things evolved. How a disputed fact fits in the timeline can often be a crucial hallmark of its authenticity and it is important as we hear or review the evidence that we have the timeline in our mind. For others summarising a report or the evidence of a witness is an essential way to make it stick in the mind. In my court, the differences in approach is also be reflected in how we utilise our associates and the tasks we give them.

Ultimately the application of the principle depends on the ethical and responsible use by judicial officers. How, then, should judges assess a proposed use? I suggest three dimensions: proximity, provenance and reversibility.

Proximity asks how close the task lies to the exercise of judicial power. At some distance are administrative functions such as indexing, duplicate detection, optical character recognition, workflow support and retrieval of published procedural information. Closer are the organisation of case material, preliminary chronologies, transcription, translation, legal research and comparison of evidence. At the centre are fact-finding, credibility assessment, the choice of legal principle, evaluative judgment, discretion, outcome and reasons. The closer the use moves to that centre, the stronger the presumption that the work must remain unaided by generative artificial intelligence.

Provenance asks whether the output can be connected to authoritative material. Source-grounded retrieval is generally more suitable than free generation. A chronology that links each proposition to a page is safer than one that narrates events. A research tool that identifies a judgment for a judge to read is safer than one that supplies an uncited synthesis. Provenance does not eliminate the need for checking, but it makes checking possible.

Reversibility asks what follows if the system is wrong. Can the error be detected before it affects a person? Can the process be repeated without the tool? Is there an audit trail? Can a party challenge the use? A typographical suggestion is reversible. An erroneous classification that causes evidence never to reach the judge may not be. The more difficult the consequence is to detect or undo, the less suitable the use.

Certain uses are inconsistent with the judicial role. Artificial intelligence should not decide facts, assess credibility, recommend sentence, determine outcome, exercise discretion or supply the reasoning that the judicial officer adopts. It must not introduce an unargued basis of decision. And confidential case material must not be exposed to a public system.

Proofreading reveals a difficulty at the boundary. A system may identify a grammatical error without affecting substance. But a request to improve clarity may reorder propositions, remove a qualification or strengthen a conclusion. This distinction is between assistance with expression and intervention in thought. The judicial officer must remain able to say, in substance as well as form: these are my reasons.

The Council of Europe's ethical charter describes artificial intelligence as remaining ‘under user control’.¹⁸ In courts, control must be epistemic as well as operational. A human must be capable of understanding, testing and rejecting the contribution the system has made.

Pilots as a method of institutional learning

The value of pilots lies in making the courts’ artificial intelligence framework concrete. A court pilot is a form of institutional inquiry. It should test whether a defined use improves the administration of justice, identify the new human work it creates, expose failure under realistic conditions and preserve the freedom to conclude that deployment should not occur.

The Coroners Court pilot, commenced in March 2024, illustrates this point.¹⁹ The Court has explored tools that may organise substantial investigative records, construct working

¹⁸ [European Commission for the Efficiency of Justice, *European Ethical Charter on the Use of Artificial Intelligence in Judicial Systems and Their Environment* \(Council of Europe, February 2019\) 7–12.](#)

¹⁹ Coroners Court of Victoria, ‘Technology at the Court’ (Web Page) <https://www.coronerscourt.vic.gov.au/technology-court>.

chronologies, identify gaps, summarise media and reduce exposure to confronting images and descriptions. Testing occurs within a secure, isolated environment. The purpose is to ask whether technology can reduce repetitive labour and vicarious trauma while preserving human control over investigation and decision.

Supreme Court pilots involving judicial officers and staff have yielded a related lesson: the strongest uses are usually narrow and grounded in familiar material.²⁰ A tool may help a judge return to a long file after an interruption, locate a passage in a closed collection, compare documents or construct a provisional chronology linked to source pages. In those circumstances, the user knows enough about the material and task to recognise whether the system is useful and where it may have failed.

This suggests calibrated reversibility. Early uses should occur where error is detectable, consequences are limited and the original position can be restored. Assurance should rise with the stakes, opacity and difficulty of correction. A useful pilot identifies kinds of error, who bears them and whether they can be found before they matter.

The Victorian transcription pilots demonstrate why context is decisive. In a Magistrates' Court trial involving eight judicial officers, 38 per cent assessed accuracy as unsatisfactory, and human assistance remained necessary to identify speakers.²¹ The finding does not establish that automated transcription has no place. It establishes that a system tested on clean audio cannot be assumed to perform in a courtroom containing overlapping voices, unfamiliar names, accents, technical language and variable acoustics. The same output might be useful as a navigational aid and unacceptable as the official record. More recently, our transcription

²⁰ Niall (n 15).

²¹ Victorian Law Reform Commission (n 10) ch 2.

service and the County Court have piloted automated transcription with greater success, reflecting advances in the technology over time as well as advances in how we use it.²² This experience illustrates a broader point: as technologies mature and their capabilities improve, applications that may once have proved unsuitable can warrant renewed consideration.

The pilots have also shown that user capability develops with experience. Judicial officers become more adept at specifying a bounded task, recognising false certainty and identifying where verification cost overwhelms the apparent saving. The object of education is to cultivate the judgment needed to decide when the technology should be used, on what terms and when it should be put aside.

A responsible pilot must also permit an adverse conclusion. Singapore's courts have combined ambitious experimentation with a willingness to stop using a public-facing generative chatbot when the conditions for responsible deployment were not satisfied.²³ That is evidence that evaluation retained authority over enthusiasm.

The proper measure of a pilot is consequently not adoption. It is institutional knowledge: knowledge of the task, the system, its failure modes, the human capabilities it requires and the public value it can, or cannot, deliver.

Measuring success

As we embark on the more formal adoption of AI in the courts we must be alive to how we will measure its cost and benefits. We should be under no illusion that the adoption of AI will

²² The pilot is delivering approximately 70% accuracy prior to human review.

²³ [Aidan Xu, 'The Use \(and Abuse\) of AI in Court' \(Speech, IT Law Series 2025: Legal and Regulatory Issues with Artificial Intelligence, 30 July 2025\).](#)

be free or easy. The investment in AI and the energy and environmental costs of data centers that run it are gargantuan. That investment will soon enough be looking for a return. The CSV framework and procurement policies will work to ensure that the investment is sustainable and provides good value for money.

In making that assessment we will have to take into account that the cost of artificial intelligence includes the work required to establish that an output is current, complete, accurate, jurisdictionally relevant and faithful to its sources. A summary produced in seconds may require an hour of checking. If confidence leads to abbreviated checking, visible work has potentially been converted into latent risk.

Traceability is therefore central to suitability. A system that points a judicial officer to the precise page, passage or authority may reduce the warranting burden. A system that produces propositions without reliable provenance increases it. For courts, the best answer may be the one that verifies a response clearly.

We must also appreciate the distinction between institutional efficiency and the transfer of cost. A court may process a task more quickly while imposing additional work on parties, registry staff or an appellate court. A self-represented litigant may produce a lengthy document at negligible cost, while the other party bears the expense of identifying fabricated propositions. An automated summary may save preparation time but require duplication because no one is willing to rely on it. Apparent efficiency at one point in the system may be displaced elsewhere. Efficiency is properly understood as an internal metric *and* of public value. It should be measured by examining whether the use of artificial intelligence reduces delay, improves the quality of engagement, protects staff wellbeing, makes procedure more navigable or releases

time for the aspects of justice that require human attention. Faster production is valuable only if the benefit survives across the proceeding.

Key takeaways and practical considerations

Let me conclude this part of the discussion with a few practical points about what this all means for judicial officers and staff in their day-to-day work.

Approved AI tools are now available to support us to work more efficiently and effectively. Copilot Chat and VicGov Chat can be used in accordance with the Responsible Use of AI Policy. They have been vetted through the Assurance Framework as being suitable for use when working with sensitive and confidential information. The prompts you use and the documents you load into these tools remain contained.

These tools join the AI assisted research tools available through the Law Library and the bespoke AI tools developed in jurisdictions such as the Coroners Court.

We continue to pilot other legal based tools and anticipate the register of approved tools growing over time.

For most judicial officers and staff, Copilot Chat will be the AI tool for everyday work. It operates within our protected Microsoft environment and can assist with a broad range of tasks. As a starting point basic guides for the use of the tools are being made available to judiciary. Staff will be required to complete an e-learn. The College and CSV will build on this in developing further training and support to help us get the most out of the available tools.

Judicial officers might use it for appropriate administrative and preparatory tasks, for example, refining the structure or expression of a speech or draft practice note, or producing a simple

draft summary or chronology from case materials. This is not as a substitute for reading those materials for the purpose of judicial decision making, but they can nevertheless be useful, for example as prompts to remind you about a matter across directions hearings.

Staff supporting you may use it in these ways as well and might also use to prepare a first draft of correspondence, or triage emails after returning from leave. It is important that you speak with staff performing tasks for you about the kinds of use you are comfortable with and your expectations of the verification and checking you require.

I encourage you to speak with staff about where an approved AI tool might safely remove some of the more routine work from their day and create more time for work that requires human expertise, judgement and attention.

That leads to perhaps the most important practical point. AI can assist with our work, but it does not assume responsibility for it. It can generate a draft, summarise information, suggest an approach or accelerate a routine task. It cannot replace professional or judicial judgement. The person using the tool remains responsible for checking its output, exercising their own judgement and ensuring the accuracy and integrity of the final work.

Used on that basis, AI offers us an opportunity to think carefully about where technology can add value while preserving the human judgement, accountability and institutional values on which the work of the courts depends.

Conclusion

Artificial intelligence confronts courts with a temptation familiar in periods of technological change: to treat what is measurable as what matters, and what is automatable as what ought to be automated. Judicial institutions must resist that compression.

Adjudication lies in the quality of inquiry, the opportunity to participate, disciplined attention to law and fact, and an identifiable judicial officer who gives reasons and bears responsibility. Those are not inefficiencies surrounding the decision. They help make it an exercise of judicial power.

That does not confine courts to existing methods. On the contrary, fidelity to judicial values should give us confidence to adopt tools that genuinely serve them. Secure, source-grounded and reversible assistance may help courts navigate scale, reduce avoidable labour, protect wellbeing and make procedure more accessible. The Victorian pilots are valuable because they allow those possibilities to be tested without surrendering the authority to define success.

The boundary is not between human work and machine work in the abstract. It is between uses that enlarge the judiciary's capacity to judge and uses that attenuate the judiciary's responsibility for judgment. Proximity to decision, provenance of output and reversibility of error provide a principled way to distinguish them.

Our task is therefore not simply to regulate a new tool. It is to preserve the institutional conditions in which judgment remains genuinely judicial. That requires courts to retain control over their information, their processes and the standards by which technological assistance is evaluated. It requires practitioners and court users to remain responsible for what they put before us. And it requires judicial officers to remain the authors of the decisions they make.

The advent of artificial intelligence need not diminish the human character of justice. Properly governed, it may create more space for it. But the sequence is important: first identify what justice requires; then ask what technology can contribute. Never the other way around.

Thank you.