

AI in legal practice

Myth-busting for busy lawyers

By Prathna Tiwari
Head of Legal Content, LawY

LawY



Summary

01

Cutting through the noise

Let's explore what AI actually is and the root of "AI anxiety" for lawyers.

02

Myth-busting 7 assumptions

We'll sort fact from fiction to give you useful, jargon-free answers.

03


Putting AI into practice

Don't settle for theory. Explore practical tips for safe AI adoption.



 Shelley Burger

 LawY AI

 @meetlawy



FOREWORD

Shelley Burger

Somewhere between "AI will replace lawyers" and "AI is just autocomplete" lies the truth, and it is actually quite practical.

The legal profession is navigating a genuine shift, and for lawyers running lean practices without a legal ops team or a tech department, sorting fact from fiction can feel like yet another task on an already impossible list.

That is exactly why we wrote this.

At LawY, we're a team of lawyers & technologists who have spent the last two years working with over 80,000 lawyers to understand where AI genuinely helps and where it falls short.

This guide will not tell you to adopt everything immediately, nor will it tell you to wait.

It will give you seven clear, grounded answers to the assumptions we hear most often, along with practical steps you can take at your own pace, proportionate to your practice.

Shelley Burger

Global CEO of LawY



01

Cutting through the noise

Chapter one

Cutting through the noise

Exploring AI anxiety

Most of the anxiety about AI in law comes from treating it as either:

1. a magic lawyer-in-a-box; or
2. a career-ending asteroid.

In reality, today's AI is closer to a fast junior who never sleeps, sometimes misunderstands instructions, and occasionally invents case citations with heroic confidence.

Used well, AI can remove friction from legal work: first drafts, document summaries, checklists, email and client communications, and other admin tasks that must be done even when they are not billable. That matters most for practices with limited resources.

Used badly, AI can increase risk. It is undeniable that incorrect use poses risk of confidentiality leaks, fabricated authorities, biased outputs, and over-reliance by staff or clients.

The core theme of responsible guidance is consistent across jurisdictions and sectors: keep humans accountable, protect confidential information, verify accuracy, and document how you use the tools.

This paper myth-busts seven common assumptions, reducing AI anxiety and offering practical, proportionate steps for safe adoption.

Defining what we mean by "AI"

"AI" is a large umbrella. In practice management conversations, most people mean one of two things:

- automation inside existing software (for example, sorting, tagging, routing, or extracting information); and
- generative AI tools (chatbots and assistants) built on large language models (LLMs) that generate text and other content.

Generative AI chatbots are not "intelligent" in the ordinary human sense. LLMs work by predicting the next most likely word (and then the next, and the next) based on patterns in their training data and your prompt. That is why the output can be fluent, persuasive, and completely wrong.

So think of an LLM as the world's most confident autocomplete. Autocompletion is great for speeding up drafting and summarising. It is a dangerous substitute for authority checking.

Because LLMs optimise for plausible language, they can fill gaps with invented material (for example, fictitious cases, quotations, or legislation). The tone may sound authoritative, but the tool does not reliably distinguish facts from inferences or opinions.

In other words, if you ask it to write a submission with authorities, it may confidently produce a beautifully structured argument supported by imaginary cases and legislations.

So what is AI good for?

Used within its limits, AI can be valuable as a drafting, organising, and communication aid. Think of it like using Google for quick research, we don't accept search results as gospel. Instead, we critically evaluate the credibility and relevance of each source before relying on it.

The same principle applies to AI outputs: they require verification and professional judgment.

AI becomes significantly more reliable when combined with specialist datasets or software systems (for example, a purpose-built legal AI platform like LawY where its research features always link responses to authoritative sources, allowing you to verify the AI's answer against primary materials).

Let's move to myth-busting. By the end of this section, I hope you'll walk away with less anxiety about AI and some practical strategies to help you experiment with it safely and integrate it into your work.



02

Myth-busting 7 assumptions

Chapter two

Myth-busting 7 assumptions



Myth 1: AI will replace lawyers

I've heard this one too many times now. Some will joke about it; others will talk about it as if they are announcing the profession's obituary.

It's not true. And there's solid academic evidence for why.

AI will transform legal work, but it cannot replace professional judgement, ethical responsibility, or the relational aspects of legal service. We are the ones with the law degree and years of hard-won experience; that remains irreplaceable.

Research by Professor Michael Legg and Dr Felicity Bell of UNSW Law gives us a clear framework. AI reduces the cost of prediction: the ability to generate a probable outcome from data. That's useful. But law is not just prediction. It involves judgement: deciding what to do with predictions, weighing factors the data cannot capture, communicating reasons to a client, all the while exercising ethical responsibility. As AI automates prediction, human judgement becomes the scarce and therefore more valuable resource. Those things remain irreducibly human.

Let me give you three reasons why human judgement is irreplaceable.

01

AI cannot reason in context

A family law matter may yield a statistical prediction about likely asset division. But your judgement about what is actually best for that client (such as factoring in the relationships, the children, the long-term wellbeing) is something AI cannot replicate.

02

AI cannot communicate with reasons

Recognising when a client is withholding information, asking the questions that uncover the real issue, and explaining the tactical implication of a clause are all irreplaceable human skills.

03

Professional responsibility is inherently human.

Duties of competence, confidentiality, loyalty, and duties to the court cannot be outsourced to an algorithm.



Myth 2: AI will make young lawyers lazy

I hear this a lot from partners. And I understand the instinct. Nobody wants to see a junior produce a mediocre AI-assisted memo and claim it as their own work.

But here's the reality: juniors are already using AI. Right now. Whether you've told them to or not. The question is whether they're doing it safely, with oversight and standards, or doing it behind closed doors with no guardrails at all.

A better analogy than 'lazy junior' is 'junior with a very fast keyboard.' You still need to teach legal reasoning, issue spotting, client management, ethics, and professional communication. AI doesn't do that training for you, but it can create space for it.

Research by Bell and Rogers notes that generative AI will reshape how junior lawyers develop their professional identity. The response to that is not restriction, it's guidance. Teaching disciplined verification, critical evaluation of AI outputs, and the reasoning skills that remain distinctly human.

In practice, that means juniors who check authorities, confirm facts, and spot overconfident errors will be more valuable, not less skilled. The key skill is no longer 'running searches', it's 'knowing when the AI is bluffing.'

So instead of fighting the instinct to use AI or imposing a blanket ban, let's channel it into discipline.

You can do this by:



Building an 'AI second-pass' culture

The junior drafts first, uses AI to stress-test structure and clarity, then verifies everything before it goes anywhere near a client or a court.



Cultivating source discipline

If AI suggests an authority, the junior must locate it in a reliable database and read it. Their inherent legal research skillsets are at play here.



Myth 3: AI is basically Google with better manners

No. No, it isn't. This is arguably the most dangerous myth on the list.

Google retrieves. AI generates. Very different beasts. When you type into Google, it finds documents that exist. When you type into a public AI chatbot, it produces the most probable sequence of words based on what it was trained on. Sometimes correct. Other times it could be outdated. And other times it could just be persuasive nonsense.

Public chatbots may have limited access to current Australian law and can fabricate cases and legislation and therefore hallucinations are not a rare edge case; it is a known failure mode.

What's especially sobering is the psychological research on why lawyers fall for hallucinated content. And it's not just the lazy lawyers falling for this. Two overlapping phenomena explain most cases of AI misuse.



Automation bias

The tendency to trust technology recommendations more than they deserve. The GPS says turn right, so you turn right, even when you know the road is closed.



Verification drift

Users who know they should verify, but ultimately don't, drawn in by the tool's authoritative tone, pushed by time pressure. The verification step gets deferred. Then deferred again. Until it's filing day.

Together, these two are a dangerous combination.



Myth 4: If I avoid AI, I avoid risk

I wish this were true. It would make governance so much simpler. You would just say "No AI" and you're done.

But even if you personally don't use AI, your clients, opponents, staff, and self-represented litigants probably do, or soon will.

This was depicted in the Victorian Supreme Court Case from 2025 of Director of Public Prosecutions v GR [2025] VSC 490. The defence's submissions contained errors due to incorrect AI use and failure to verify. However later in the proceedings, the parties filed submissions on a joint basis. The prosecution failed to identify the very same errors. Justice Elliot noted that the prosecution had agreed with the defence submissions and did not check the references because it accepted the conclusion.

The cases we will discuss later and those you have probably already heard of generally involve lawyers who used AI carelessly and filed the result. This case however involves a lawyer who got caught in the consequences of someone else's AI errors.

This case therefore is a clear warning that simply not using AI yourself does not mean you can ignore its existence or avoid understanding its limitations. You need to build the skillset to recognise AI content.

The practical risk for a small practice is 'shadow AI': staff using tools embedded inside email apps, web browsers, or office suites or practice management software's, without approval, often just by clicking a convenient new button that appeared this week.

In Mertz & Mertz (No 3) [2025] FedCFamC1A 222 (a case from the Federal Court), a solicitor asserted that her paralegal had used AI when drafting a summary of argument that referred to fictitious cases. The solicitor had failed to check or provide supervision or training to that paralegal. That case is a cautionary tale for anyone managing staff who may be using AI without your knowledge.

Therefore avoiding the topic simply increases the risk and does not reduce it. A simple policy, training, and an approval pathway is usually safer than a blanket ban that nobody actually follows.

The Robodebt Royal Commission findings reinforce the point. The lawyers involved weren't personally running an AI tool. They were advising on an automated system. The lesson: AI does not need to be on your desktop to engage your professional obligations. It may arrive through a client's system, a government process, or a tool quietly embedded in your office software.

So avoiding AI cannot be your AI policy. What you should be doing is:

- Auditing what tools are switched on by default in your office.
- Creating clear rules for workflows in which you'll permit AI use and workflows that must remain fully human driven.
- Making sure everyone understands the risks of using public chatbots like ChatGPT.

Exploring the enterprise AI tools either within or outside your practice management systems.



Myth 5: Using AI means confidentiality is automatically breached

Now as we have discussed, confidentiality is a real risk when speaking about AI. It is. But it is a manageable risk, and therefore a blank prohibition on AI tools is not necessary.

The safest rule is also the simplest one: if you would not paste it into a social media post, do not paste it into a public chatbot.

Seriously. That is the test.

As we covered in the ethics section, public AI tools may use your inputs to train future models.

Enterprise tools, with the data-handling safeguards we discussed earlier, remain the ethical choice when client data is involved.

So confidentiality and AI can coexist — it just comes down to which tool you're using and what you're putting into it.

Practical Tips:

01

Adopt a 'no confidential data in public AI' rule. If you want a task to be done with AI and it has to involve client data, you can only rely on your legal AI tools.

02

Apply data minimisation: remove personal identification information from your prompts and ensure your prompt only provides the information required to complete the task. It is unlikely the client's full name, date of birth and address is ever necessary in an AI prompt for a simple task. Think before you prompt.



Myth 6: AI will let you do the same work, bill the same time, and keep the difference

I admire the optimism. I really do.

But the market is already moving and clients are paying attention. The Australian legal market is grappling with what commentators are calling a 'productivity paradox': the more efficient firms become using AI, the less they earn under traditional time-based billing models. And regulators are watching. It is foreseeable that the Legal Services Commission may see a rise in complaints against lawyers for failing to pass on fee reductions where AI has been used and regulators are likely to treat this as more than a commercial dispute, examining whether the lawyer has complied with their fiduciary, ethical, and costs disclosure obligations.

The smarter play is using AI to deliver more value not to quietly pocket the time savings. For smaller practices, the opportunity goes even further. AI doesn't just help you work faster. It means you no longer must turn clients away or take on more staff to keep up during a busy period. If you're at capacity, a smarter workflow creates room to do more with what you already have. That's not just efficiency: that's growth.

The market is already responding. Fixed fee arrangements, subscription-based legal services, and alternative billing models have quickly gained momentum across Australia, driven by client demand for transparency and predictability. Among mid-sized Australian firms, 64% now offer flat fees and 27% have adopted subscription-based pricing, a significant shift away from the traditional hourly billing model.

Firms that are adjusting their pricing in response to AI are broadly doing one of three things: some are increasing prices, reflecting the higher value of AI-augmented work; some are reducing them, passing efficiency gains directly to clients; and some are doing something more deliberate: introducing AI-specific fees as a transparent line item in their costs agreements. It is that third approach that deserves a closer look, because it represents a genuinely new way of thinking about how legal services are priced.

Two models are emerging:

01

The first is the flat AI services fee. Rather than absorbing AI subscription costs as overhead or quietly building them into hourly rates, the lawyer names the charge openly and explains what it delivers in their cost agreements. This is not a new concept in principle. Law firms have long passed on disbursements like photocopying. AI is simply the latest tool that carries a real cost and produces a tangible benefit. The difference is that unlike a photocopy, AI materially changes the quality and speed of the work: research that once took a junior lawyer half a day can be completed in minutes, documents that required multiple drafts can be turned around faster, and briefs that would have taken days to compile, are ready in hours.

When you're upfront about it, clients accept this fee without hesitation, especially when their overall bill comes down. The trade-off is simple: a small technology fee in exchange for faster turnaround and more thorough work. Everyone wins. For the lawyer, it also means the tools making your practice more efficient aren't quietly eating into your margin.

02

The second model is more ambitious and represents a fundamental rethinking of how AI fits into a legal practice. Rather than treating AI as an overhead to be recovered, this approach positions AI paralegal work as directly billable activity: work that is logged, tracked, and charged against the cost agreement just as any other professional service would be. The AI does not sit behind the scenes reducing the lawyer's time; it sits alongside the lawyer as a productive member of the team, generating output that has real value to the client and real revenue for the firm.

This model has begun to take shape in the Australian market, particularly in family law, where AI tools are now being used to autonomously source financial disclosure documents, prepare balance sheets, compile disclosure schedules, and draft consent order applications. Work that would otherwise fall to a human paralegal or the lawyer themselves is completed in hours rather than weeks, and that work is billed directly to the matter. Rather than the AI being a cost the firm absorbs, it becomes a contributor to the bottom line. The overhead concern disappears entirely, because the tool earns more than it costs and the value to the client is not abstract. It is visible, concrete, and moves the matter forward.

Both models can coexist, and both are preferable to the alternative, which is doing nothing and hoping clients don't notice. They will notice. And increasingly, they are asking.

Practical Tips:

01

Keep your time records honest, consider fixed or capped fees where AI is doing the heavy lifting, and update your retainer if your delivery model has changed. Most importantly — get ahead of the conversation. If AI is making you faster, clients will notice, and regulators are paying attention. It's far better to explain your pricing on your terms than to justify it on someone else's.



Myth 7: AI errors are the tool's problem, not mine

As the case law before us demonstrates, this assumption has serious consequences for those who fall afoul.

DPP v Khan (2024, Australia). *DPP v Khan*, decided by Justice Mossop of the ACT Supreme Court in February 2024. In that matter, a character reference from the offender's brother was tendered in sentencing proceedings. Justice Mossop found the reference 'consistent with an artificial intelligence generated document.' His Honour held that it was clearly inappropriate for personal references in such proceedings to be generated by, or with the assistance of, large language models. And he made clear that counsel has a duty to enquire how references are written before tendering them.

This is another case where someone else's AI generated content has had repercussions for the lawyer. Failure to recognise AI generated content is the real problem in this case. Given the increasing use of AI by everyone, discernment and identification of AI-generated material will be a necessary skillset as a lawyer.

Valu v Minister (2025, Federal Circuit Court). In *Valu v Minister for Immigration and Multicultural Affairs (No 2)* [2025] FedCFamC2G 95, a legal representative admitted using ChatGPT to identify cases. The practitioner said the summary 'read well' and so he incorporated it into his submissions without checking the details. He was referred to the Legal Services Commissioner. The consequences included adjournments, increased costs, and unnecessary additional work for the Court and the other party.

In ***Oberoi v Douglas [2026] VSCA 31***, the Victorian Court of Appeal found seven non-existent cases and twelve fabricated quotations in an applicant's written submissions. The document had been signed by a solicitor. The Court wrote to that solicitor. Twice. Asking for an explanation. No substantive response came. The solicitor was referred to the Legal Services Commissioner. The Court was clear — relying on fabricated authorities falls short of the standard the public is entitled to expect. And conduct that undermines confidence in the justice system breaches the rules of professional conduct. This wasn't a grey area. It was a career-defining mistake, made in minutes.

That duty has always been the lawyer's job: enquire, check, never just accept material and hand it to the court. What AI changes is this: for the first time, we have a category of material that looks completely authentic and can be entirely made up.

Gerrard J in a Federal Court judgement from 2025 expressed concern that the lawyer's in question did not sufficiently understand what professional verification actually requires. The Court stated, in substance, that it is not enough for a practitioner simply to check whether a cited case exists. Proper professional diligence requires more. A lawyer must also review whether the authority:

- actually supports the legal proposition for which it is cited,
- remains good law,
- has been overturned, distinguished, or limited by a higher court, and
- is relevant to the factual matrix of the case at hand.

The Court's closing observation in that case was that: legal principles are not slogans that can simply be attached to submissions without context or analysis.

(*JNE v Minister for Immigration and Citizenship* [2025] FedCFamC2G 1314)

Researchers have identified five recurring reasons AI goes wrong in litigation and I think these reasons can apply to use of AI in all legal matters. They are:



Limited AI literacy

Lawyers who don't fully understand how the tool works or what it can get wrong. Which is exactly why I bore you with some AI technical explanations earlier in this session.



The copy-paste culture of legal practice

The habit of lifting text directly from one document and dropping it into another without critically reviewing it first. It's how precedents get recycled, how boilerplate clauses travel from matter to matter, and how research gets moved from a draft into a final document. The assumption is that if it worked before, it's probably fine now.



Limited legal literacy

Relevant more for self-represented litigants.



Verification drift

Again something we have covered, which is where checking becomes less rigorous over time as the tool feels more reliable.



Automation bias

Which as we've discussed, is the tendency to trust output simply because a machine produced it.

Now none of these are new human failings. AI just gives them sharper consequences.

The global picture reflects the same theme. Jurisdictions across the world are issuing AI use guidelines for practitioners and judicial officers. Verification and disclosure is the consistent theme everywhere.

It is worth noting that these guidelines are not only directed at practitioners. The Queensland Courts' revised Guidelines for Judicial Officers (September 2025). Judicial officers apart from not using AI for decision making or writing, are required now to alert to the possibility that material filed before them could be AI-generated. The Guidelines even list the warning signs such as unfamiliar citations, submissions that don't align with the law as the judge understands it, or language that doesn't match how the person usually speaks or writes.

So to think that it's the AI tools problem and not yours, is completely wrong. When something goes wrong, the court will not be asking what the AI did. It will be asking what you did.

Practical Tips:

01

Before filing anything AI-assisted: verify every authority from a trusted database; verify every quote against the source; confirm factual assertions against your instructions and evidence.

02

Keep a record of what you checked because having a verification story is critical to discharging your disclosure obligations.

03

Ensure you are across the specific jurisdiction's AI guidance notes and directions before using AI in litigation.



03

Putting AI into practice

Chapter three

Putting AI into practice

Where AI actually helps

When we consider the practical side of what ethical and effective AI use looks like, it's helpful to laser in on specific use cases that lawyers will tackle daily.

High-Value, Low-Risk Use Cases

These tasks will thrive with a purpose-built legal AI tool where your matter details are integrated.



First-draft of simple correspondences

Think chasing that client for the executed contract, sending a post settlement update, or following up on outstanding instructions or documents.



Plain-English rewriting

You supply the law and the legal reasoning. Think translating the importance of the section 32 vendor statement to a first home buyer who is overwhelmed by the whole process or explaining restraint of trade to a client who just sold their business.



Chronologies organisation

Summarising a long document into a structured outline, extracting key dates into a chronology. Think: pulling key milestones out of a 200-page document.



Practice management content

Internal policies, precedent instructions, staff training materials, matter opening and file closing checklists. Think putting together an induction guide for a new paralegal joining next Monday.

Higher-Risk Use Cases: Possible, With Guardrails



Drafting substantive documents

Useful for structure and language, but every factual statement and every authority must be verified by you. AI can give you a solid first draft, then it's up to you to refine it to your instructions and current law. Avoid staring at that blank screen!



Summarising litigation material

Helpful first-pass, but prone to subtle factual distortion. Therefore, treat it like a draft from your junior, you will need to review it carefully.



Legal research assistance

Excellent for issue-spotting and structuring your analysis, but unreliable until verified by a lawyer who knows the matter. Think: ask AI to map out the elements of misleading and deceptive conduct under the ACL in a dispute. It will give you a useful starting framework, but you still need to confirm the cases actually exist, the principles haven't shifted, and the analysis holds up against your specific facts.

High-Risk: Recommended to avoid



Uploading confidential client documents into a public chatbot

I hope by now it is clear what a horrid idea this is and that your confidentiality obligations don't have a carve-out for convenience.



Using AI-generated citations without independent verification

We have covered the case law and highlighted the risks. AI cannot and should not replace your legal judgement.



Allowing AI to operate without meaningful human oversight

A chatbot on your website to field enquiries and bring in leads is a great idea, as long as all it is doing is gathering information and not providing unsupervised legal advice.



Using AI as a decision-maker

Think: letting AI determine whether to advise a client to settle, issue proceedings, or sign a deed. It does not hold a practising certificate. It does not have a duty to the court. It does not carry the ethical obligations that sit behind every piece of advice you give. You do.

A proportionate governance model

If your firm has no AI policy, you already have one. It's just unwritten, inconsistent, and enforced by vibes.

Beyond the policy, two practical tools make a real difference.

01

An AI register

A living list of the tools in use, how they're used, and who is accountable. Tool name, approved uses, data rules, accountable person, last review date, incident log. That's it.

02

Human oversight as a standing rule

AI can be fast and persuasive, but responsibility does not disappear because a tool was used. A competent lawyer must review any output that could mislead a client or court; this is the lawyer-in-the-loop principle we covered earlier.

For small firms, the main risks are over-reliance and unauthorised use. Your training pathway doesn't need to be a course. It needs to cover five things:

1. What AI is and isn't;
2. Confidentiality rules;
3. Accuracy rules;
4. Which tools are approved; and
5. How to label AI-assisted drafts so they're never mistakenly treated as final.

Become an AI-enhanced lawyer

Now at the start of this guide, I promised to reduce your AI anxiety, so I'll end with another attempt at doing just that.

AI is not here to steal your practising certificate. It's here to steal your admin time, if you let it. And used properly, with the right tools and the right guardrails, it is here to make you a better, more efficient, more client-focused lawyer.

Your ethical duties are the framework within which this enhancement occurs. Competence requires you to understand AI and its limitations. Confidentiality requires you to keep client information out of public tools. The duty to court requires you to verify everything you file. And your duty of independence requires that AI informs your judgement, but never substitutes for it.

The AI-Enhanced Lawyer is not someone who uses any tool available to them. It is a lawyer who chooses purpose-built enterprise tools designed with legal obligations in mind. Who keeps human judgement at the centre of every decision. Who verifies every output before it leaves their hands. Who maintains the guardrails of confidentiality, accuracy, and ethical accountability. And who invests the time AI saves into the high-value, irreplaceable work that clients actually need.

AI is not the end of lawyers, but it will change how you practise. The lawyers who thrive will be the ones who know when to use AI, understand how it works, and can critically evaluate what it produces. I hope this guide has made that feel a little more achievable.



ABOUT THE AUTHOR

Prathna Tiwari

Prathna Tiwari is a property and commercial lawyer with 11 years of experience. She was an early adopter of legal technology in practice, introducing tools like Dragon Dictate and PEXA before they became mainstream, and has since brought that passion for legal tech to LawY. As Head of Legal Content, she champions the ongoing growth of our Verifier panel and helps other lawyers discover the unique power of LawY in practice.

Prathna Tiwari

Head of Legal Content, LawY



Lawy AI



@meetlawY



hello@lawy.ai



About LawY

LawY is your trusted AI legal solution: built by lawyers, for lawyers. Speed up what slows you down with features that streamline repetitive drafting, research, contract review, and admin. Over 80,000 lawyers worldwide have already reclaimed hours in their week with LawY. See it for yourself by booking a personalised demo, led by one of the lawyers in our team:

[**Book a demo**](#)