



Queensland
Law Society[®]

AI Companion Guide

The QLS overview guide to AI for solicitors

Version: 1.0
September 2024

Executive summary

In 2024, proficiency with artificial intelligence ('AI') tools and legal automation will not determine a firm's commercial viability. Whether this will still be true in 2026 or 2030 is uncertain, so lawyers will need to watch developments carefully.

However, in the longer term it would be wise to consider AI as both a potential and disruptive technology in law. The work many smaller firms currently rely on as their bread and butter is most easily automated, so these firms will either need to utilise such technology or pursue other work to replace it.

Even if you don't use AI in your practice, you will need to adapt to an environment in which competitors and clients do. Adopting AI successfully is not (just) a matter of buying some software and turning it on. A change management and risk management process is required.

When using an AI tool, solicitors will need to:

- ensure that they get a return on investment, even if this is by acquisition of new skills;
- be mindful of client confidentiality and intellectual property;
- check automated output very carefully; and
- make sure that ethical duties to the Court, clients, colleagues and third parties are not forgotten.

Purpose of this guide

This introduction to generative AI is designed to be a primer for solicitors and firms, particularly small and medium-sized firms (SMEs) who want to understand more about the technology. Aspects of this guidance will also be relevant for solicitors working in-house.

AI for SME law practices is a catch-22 conundrum: SME Principals are time poor and urgently need a cost-effective way to automate and speed up routine work. At the other end of that rainbow is more time for clients, family and a fulfilling life. AI has that potential. However, to go from potential to reality a firm's leaders will need to devote significant time to selecting the right tool and using it safely.

The objective of this guide is to give solicitors (not just Principals) a starting point. It is intended to impart the following capabilities:

- to understand general discussion of AI technology;
- judge if and when it is time to consider AI in your own practice; and
- construct an appropriate governance and risk framework.

This guide should be read in conjunction with:

- [Guidance Statement No. 37 Artificial Intelligence in Legal Practice](#)
- [QLS Artificial Intelligence Policy template](#)

QLS gratefully acknowledges the Law Society of England and Wales for allowing us to use and adapt content from their webpage on introductory AI. Some blocks of text are used without specific attribution.

Contents

Executive summary	2
Purpose of this guide	2
Part A: Consideration of AI	4
1. Will AI impact my professional life?	4
2. Will AI make lawyers obsolete?	4
3. The Hype Cycle	5
4. Ensure staff are not using AI systems without approval	5
5. Consider AI for general office and administration work	6
Part B: Basic concepts	6
6. Glossary	6
7. Large Language Models	7
7.1. An LLM is really only good at one thing, although it talks a good game	7
7.2. “Hallucination” and inaccuracy	8
7.3. Other LLM limitations	8
7.4. LLMs combined with something else are potentially a lot more useful.	8
7.5. Accuracy is improving, but training wheels are still required	9
8. Subject to those limitations, what can AI do for lawyers?	10
Part C: Selection and Deployment	11
9. Where would we start?	11
10. Risk assessment and management	12
11. Protecting client interests	13
11.1. Intellectual property and clients	13
11.2. Intellectual property and third parties	15
11.3. Cybersecurity	15
Part D: Using AI successfully	16
12. Managing reliability limitations	16
12.1. Output supervision	16
12.2. Process supervision	16
12.3. Accountability structures	16
13. Disclosure when using AI systems	17
13.1. Disclosure by lawyers to the court	17
13.2. Disclosure by solicitors to their clients	18
13.3. Disclosure by clients to their lawyer	19
13.4. Disclosure by lawyers to their employer	19
14. ESG and bias	19
15. Remuneration	20
16. AI use by in-house and government lawyers	20
17. Business continuity	21
Annexure 1: Procurement and deployment checklist (COBIT) framework	22

Part A: Consideration of AI

1. Will AI impact my professional life?

Most solicitors and legal practices are likely to be affected by AI use within the next few years and will need to navigate a world in which AI use is commonplace. Even if you conclude that professional use of AI is not for you (or at least not yet), you will need to adapt to the fact it will increasingly feature in transactional work, litigation and decision making by large companies and governments that might affect your client's interests. An understanding of the basic concepts will be required.

Use or otherwise of AI should be an informed decision made with a pragmatic understanding of the potential for legal service automation in the areas your firm relies on. This potential (then reality) will change and evolve with the technology, so ongoing engagement is necessary to avoid unpleasant surprises.

The Queensland Law Society (QLS) does not seek to discourage use of AI tools in legal practice, nor does QLS view such use as incompatible with ethical duties. AI – used properly - has the potential to significantly enhance access to justice¹ and lessen the burden of routine work on solicitors and support staff.

Balanced against this is the need to ensure that innovation does not lead to negligence or professional censure. Appropriate use of AI requires guardrails and a realistic understanding of the limitations of the product.

It is important to keep each of these in mind when deciding how to future proof your firm and career, although for obvious reasons the majority of the immediate discussion will focus on the opportunities for efficiency gains in the near future.

2. Will AI make lawyers obsolete?

AI will not replace lawyers,² however law practices which adapt and use AI successfully³ are likely to out-perform those which can't. Understanding this technology's strengths and limitations will permit you to adopt it successfully or – if you decide that this is not for you - to ensure your personal skills and practice excel in domains where automated systems fall short.

An automated business can scale more rapidly than one which needs to hire and train a large team. Competition from alternative providers may therefore emerge quickly once they get the system working effectively, so you will need to be agile. When scanning the horizon for potential competition keep in mind that the structure of the legal services market may change. Competition may be from larger law firms reaching down to cherry-pick work which is newly profitable, or from businesses which are not traditionally competitors to lawyers at all.

Given the varied needs of individuals and small business, it is highly unlikely that all their legal work can be supplied by hybrid or automated services, no matter how good such systems become. However, to disrupt an industry you do not need to take all the work; you just need to take the most profitable portions. Automation works best in matters which are straightforward, predictable and can be broken into replicable workflows. That kind of matter is also the bread and butter for many SME law firms.

¹ Colleen Chien and Miriam Kim, 'Generative AI and Legal Aid: Results from a Field Study and 100 Use Cases to Bridge the Access to Justice Gap' (Research Paper forthcoming, Berkeley School of Law, University of California, 11 April 2024).

² It is likely that a significant number of functions a law firm currently does for its clients will ultimately be automated, and some law jobs will change significantly: World Economic Forum, *Future of Jobs Report 2023* (Insight Report, May 2023).

³ 'Technology and Innovation in Legal Services', *Legal Services Board of England and Wales* (Web Page, May 2023).

At the date of publication (2024), many AI tools are in their infancy. Assuming continued improvement,⁴ tangibly useful AI enhancement to legal practice is likely within 2 years, with significant structural change in the way that many consumers and business obtain legal advice by 2030.⁵

Important note: Software or tools mentioned in this guide are supplied as example only, not a recommendation. No risk assessment process has been applied to these products and no endorsement should be implied.

Until recently, most AI options were either generic consumer models or expensive bespoke systems designed and implemented in-house for large firms. AI systems for lawyers are now being released as stand-alone tools or integrated into practice management systems. Provided a cost / efficiency benefit can be achieved, these may soon be profitable for smaller practices although acquisition and deployment costs will be significant.

3. The Hype Cycle

Technology ranking and prediction consultancy Gartner observes that many new technologies [go through a predictable cycle](#)⁶ of Hype and apparent failure before emerging as an important productivity driver. AI is likely to be no different.

Neither the unrealistic expectation phase nor the inevitable reversal in sentiment that follow it represent the true picture. As the AI market matures beyond Large Language Models a number of products specific to the legal profession are emerging, either as stand-alone products or embedded into the software suites of practice management or legal research providers.

If an appropriate use-case can be identified for your practice, these products are worth considering for SME practices. Larger firms will consider both these and a more ambitious bespoke project more directly aligned with their existing practice requirements.

4. Ensure staff are not using AI systems without approval

Like other unauthorised use of IT platforms,⁷ “shadow AI” – or AI tools being used without approval is likely to be prevalent in law firms.⁸ This may be a deliberate choice or inadvertent. Many emerging AI functions look like a new feature of existing software. A tool to, for example, summarise a lengthy document, take meeting notes or suggest alternative wording in emails might be enabled by clicking a link or pop-up box in an existing program. Activating this feature may install third party software or a browser add on. Even if it is an additional feature of the same software the activation process often grants additional user data access permissions.

The reason this is a problem is self evident. But simple prohibition is insufficient. Firstly, without education, staff can quite reasonably say that using a new feature of software already supplied by the firm is fair enough. Secondly, there is a common misconception that a work system will block anything which is not approved.

The appropriate control measure is to develop and effectively communicate an AI usage policy. Firms should also educate users that as soon as they are agreeing to terms and conditions approval is required.

⁴ Continued linear improvement in existing large AI Models is unlikely, however there will be significant improvements in the way that generalist tools can be used in conjunction with other systems. There are several fundamental technological problems for AI that remain unsolved, and a significant debate is under way as to whether these are speed bumps or major hurdles.

⁵ London School of Economics Student Union and AI Society, ‘AI in Law & The Legal Profession: Industry Insights Report’ (Conference Paper, 2024 London School of Economics Law Summit, March 2024).

⁶ ‘Gartner Hype Cycle’, *Gartner* (Web Page, 2024).

⁷ Data from Australia and elsewhere consistently show very high rates of unauthorised IT systems in organisations of all sizes. See, eg, Robert Nikolouzos, ‘Shining a light on Shadow IT’, *PWC Digital Pulse* (Article, 12 December 2001).

⁸ Research by Microsoft shows that large numbers (up to 75%) of knowledge workers are using self-funded AI for work. See, ‘AI at Work Is Here. Now Comes the Hard Part: 2024 Work Trend Index Annual Report’, *Microsoft WorkLab* (Article, 8 May 2024).

A template AI policy for small firms is available on the QLS website.⁹ This is intended as a starting point rather than something which can be universally adopted. The basic tenor of the document is that only authorised use is permitted, including a pathway to request approval. This last is important, as without some mechanism to seek access to potentially useful tools some staff may conclude it is better to ask for forgiveness rather than seek permission.

5. Consider AI for general office and administration work

AI tools for legal tasks have been available for some time, but most required complicated and expensive implementation. The most mature are only really practical for larger firms.¹⁰

A lot of the immediately useful products for an SME are not law-office specific (at least not for our jurisdiction) but aimed at office work more generally.¹¹

As an example, see this demonstration of Copilot for Office 365: Teacher's Tech, 'Don't Miss Out on Microsoft 365 Copilot Features' (YouTube, 4 March 2024)

<<https://www.youtube.com/watch?v=AhywEEHg6Es>>.

Part B: Basic concepts

A basic understanding of the terminology is a necessary underpinning to discussing the issues arising from AI, so please read this short glossary section. As with many evolving technologies there is often no universally agreed definition of terms used in this guide, but for our purposes the general idea is sufficient.

6. Glossary

Artificial Intelligence (AI): AI enables machines to perform tasks that typically require human intelligence, such as interaction using natural language, sorting data and solving problems. However, intelligence and understanding are only simulated, not replicated.

Generative Artificial Intelligence: A type of AI technology that can create new content, such as text, images, or music. The content generated is novel, in the sense that the exact combination may never have existed before, but the system will only be able to produce content that is a synthesis of the training data ingested. Generative AI output is very different to a search engine (which only locates pre-existing information) but can be useful in conjunction with search by summarising, mapping and synthesising search output. A Generative AI system is predictive based on the training data and therefore subject to limitations in that data (see: *Bias, Hallucination*).

Chatbot: A digital tool designed to simulate conversation with users, primarily via text or synthesised speech. Some chatbots operate based on predefined responses, determining the enquirer's question and then matching this to pre-prepared material. More advanced versions use AI (including generative AI techniques) to provide more dynamic and contextually relevant interactions, reducing the need for immediate human intervention.

Machine learning: Programming a computer using a large body of relevant data ("training data"). The machine is told the objective but not how to achieve it. The machine uses trial and error to develop a model, which solves the problem it was set. The model self-corrects by checking whether it is "right" in comparison to outcomes reflected in the dataset ("model training"). The resulting model can often improve over time through further interaction and data ingestion ("model refinement" or "model tuning"). Each model will tend to excel at one task, with significantly limited capability in other dimensions.

⁹ 'Artificial Intelligence Policy Template', *Queensland Law Society* (Template, 19 April 2023)
<<https://www.qls.com.au/Content-Collections/Template/QLS-Artificial-Intelligence-Policy-template>>.

¹⁰ E-discovery tools utilising AI have been used in law firms for some years, but require careful data calibration for each new matter. The calibration and testing cost has meant that – until fairly recently – it was cheaper and easier to undertake discovery manually unless the case involved large volumes of material.

¹¹ This is changing fast. See Section 7.4 of this Guide for examples.

Large language model (“LLM”): A type of AI tool (made possible by the advent of Machine Learning) that interacts naturally with humans using human language rather than traditional computer code or instructions. It will typically be very good at simulating communication, but performance in undertaking research or performing calculations can be erratic.

Foundation Model: The term ‘foundation model’ is often used synonymously with LLM but this is a conflation of two concepts. A foundation model is the core AI model from which others are developed. An LLM is a foundation model tuned for conversational interaction.

Hallucination: The propensity of a generative AI system to create output in which certain components appear plausible but have been invented to fill gaps in the surrounding material. Current AI does not understand the content it generates, and cannot be relied upon to validate its own output. For example, a generalist LLM might produce superficially convincing legal submissions, but fill it with fictional case references and invented judicial statements. Hallucination risk is likely to reduce as systems become more refined and specialised but is unlikely to be completely eradicated.

Bias: A tool created using machine learning will inevitably reflect limitations in training data. Where the training data does not represent the population or problem the tool is applied to, “bias” or higher rates of inaccuracy may occur. For example, facial recognition software trained using images largely representative of one racial group may be less accurate in differentiating between members of a racial minority. Software trained to negotiate contracts in one jurisdiction or language *may* be effective when used in another context, however this cannot be assumed without careful validation.

7. Large Language Models

The systems that have recently transformed the public perception of AI are mainly Large Language Models (“LLMs”). To some extent these are becoming synonymous with AI generally, and many of the widely publicised “AI failures” affecting lawyers arose from misuse of LLM’s to do something they were not designed for.

A wide range of paid, free and freemium LLMs are available. Some popular examples include:

- OpenAI’s ChatGPT (at time of publication version 4.0)
- Google’s Bard
- Anthropic’s Claude
- GitHub Copilot
- Meta’s Llama

7.1. An LLM is really only good at one thing, although it talks a good game

An LLM is a specialist interaction tool, not a general AI agent that can be expected to solve novel problems by assembling facts then reaching a conclusion from first principles.

An LLM is trained on hundreds of millions of pages of text, video and recorded sound to communicate with humans using natural human language. They are capable of parsing complex human language including nuances such as context, allusion and sarcasm. They can also completely misunderstand the user’s intent, sometimes failing unpredictably after a long series of accurate interactions.

An LLM can also be requested to change the tone and content of output to suit a particular audience, varying from light and playful marketing content to academic papers or legal submissions.

However, the proficiency of the language interaction can create a misleading impression that the output is factually accurate. While there are many facts in an LLM’s training data, the model has been primarily trained to communicate not to analyse. As a stand-alone artefact it will not be good at weighing up competing positions, allocating relevance and providing a reasoned answer from first principles. It can, however, present an excellent summary of prior explanations and “pro and con” arguments that are already in existence.

This is why a stand-alone LLM can be very effective in supplying answers to undergraduate assignment questions but much less reliable when dealing with a legal question that is not commonly analysed in publicly available material.

7.2. “Hallucination” and inaccuracy

An LLM is quite capable of producing material that looks very much like something that a human expert would create, but without the accuracy necessary for professional use.

An LLM will identify and reproduce patterns in language and structure, assembling facts that appear to be relevant. Where a jigsaw piece is missing it may fill in the void with invented (but plausible-sounding) material.

There are now many examples of lawyers coming to grief by misconstruing the authoritative tone of AI generated content for a genuinely accurate statement of the law.

7.3. Other LLM limitations

A generalist LLM such as Chat GPT can produce legal documents, submissions and judgments. These will be an amalgam of material in the public domain.¹² The more examples of particular content that are available to the model the better the output is likely to be. In specialized domains it can be a significant challenge to source precisely relevant training data, so it is quite possible the data used was a composite of material from different jurisdictions, drafting styles and specific purposes. This will be inevitably reflected in the output.

A further difficulty is that the same query can lead to a different result each time. The tool will likely draw on a limited subset of relevant content in the training data for each query. The more examples available to the AI, the better the likely quality but the higher the variability. While the generated output may be quite similar, it is important to understand that a document drafted by an LLM alone is not like a precedent with variable content inserted. The whole thing is drafted anew each time. Scrutiny therefore needs to extend to the whole document, which quickly erodes any time savings.

As a visual example the following are synthetic illustrations produced in response to identical prompts¹³ by the same system a few seconds apart:



7.4. LLMs combined with something else are potentially a lot more useful.

If you have experimented with an LLM to produce legal documents but been un-impressed you are not alone. However, an LLM harnessed to a specialist system or data set is potentially very useful, reducing the impact of an LLM's inherent limitations.

¹² Or, more accurately, data that someone has acquired and made available for training an AI. Whether or not this was really “public domain” is subject to extensive litigation.

¹³ Prompt: a cat in a farmyard, in the style of Norman Rockwell. Tool: Discord / Midjourney build #9282. Fundamentally, an LLM producing text and an image generator works in much the same way, so image generation can be a useful starting point to explore the potential and limitations of generative AI.

As an interpretation and communication layer it can allow non-specialists to use complicated software products far more intuitively, find relevant information from a disorganised set of records and put it into a format that the dumb-but-reliable software you have used for years can understand.

From the user's perspective the outcome is the same: you type or say what you want and the system generates an output. What is going on under the hood is very different, however. The LLM is not the primary generation mechanism for all output but only the tool by which the specialist system is engaged.

An example of this kind of use is Microsoft's Copilot for Office 365,¹⁴ adding a natural language layer by which a user can specify output without being proficient in the use of such tools. Examples include:

- excel tables and formulae
- data analysis and illustration in PivotTables
- presentations in .ppt.

Similarly, an LLM could be harnessed to:

- a CRM tool to maintain a [real-time dashboard](#)¹⁵ of a client's matter status
- [practice management software](#)¹⁶ to extract data for use in precedents
- [project management software](#)¹⁷ to assign and track workflows in a complex transaction
- an accounting package to [improve costs estimation](#).¹⁸

7.5. Accuracy is improving, but training wheels are still required.

AI lacks the capability to understand its output and meaning in the same way that humans do. In fact, as it currently stands it has no judgement or common sense whatever.

As a general starting point, any generative AI system in 2024 will be unable to undertake automatic self-validation. An error rate must be assumed.¹⁹

Errors can be *reduced* by:

- using specialized systems
- integrating unpredictable LLMs to drive other systems
- ensuring the task is specified in an appropriate way ("[prompt engineering](#)").²⁰
- designing the model to research and check the answer from a specified body of reliable external information, ("[Retrieval Augmented Generation](#), or RAG"),²¹
- training the system on appropriate content and linking it to jurisdiction specific legal libraries and precedent sources.

¹⁴ Copilot will offer assistance such as "let me draft a response to this email" which is primarily using an LLM, so the usual limitations would apply to that element of the task.

¹⁵ See, eg, 'Legal Case Management Template', *Retool* (Web Page).

¹⁶ See, eg, 'Meet CoCounsel – the world's first AI legal assistant', *Thomson Reuters* (Blog Post, 1 March 2023).

¹⁷ See, eg, Mary Pratt, 'How is AI transforming project management', *Tech Target* (Article, 7 March 2024).

¹⁸ See, eg, Smart PM Hub, 'Unlocking efficiency: A deep dive into AI-enabled project cost management', *Medium* (Blog Post, 20 June 2023) ('AI-enabled project cost management').

¹⁹ Stanford University's school for Human Centred AI undertook research into Legal-Tech AI error rates and was not complimentary in their assessment: Varun Magesh et al, 'AI on Trial: Legal Models Hallucinate in 1 out of 6 (or More) Benchmarking Queries', *Stanford University HAI News* (Article, 30 May 2024). Note, this research has been criticised, see (n 22).

²⁰ See, eg, 'Building the future of legal services with legal AI', *LexisNexis* (Blog Post, 6 February 2024).

²¹ See, eg, Priyanka Vergadia, 'Prompt engineering and you: How to prepare for the work of the future', *Google Cloud* (Blog Post, 27 September 2023). Note that "searching the internet" for relevant material is very resource intensive for an LLM. For example, the provider of a tool using GPT to analyse legal cases for relevance to a particular problem would be charged per word in all cases analysed. The cost mounts quickly. To manage this, the tool must either (1) truncate the number of cases analysed, which will dramatically reduce accuracy, (2) have access to a library which permits keyword searching or (3) pass the full cost – perhaps as much as \$10 per query, onto the user.

Some AI systems will also fine tune themselves to enhance domain-specific accuracy, learning from users and the documents that are uploaded and therefore improving over time.²²

Despite rapid improvement, as at the date of publication no combination of these strategies has produced a system that could be used in a law firm without careful and ongoing supervision of all output.²³

8. Subject to those limitations, what can AI do for lawyers?

It should be no surprise that AI tools designed only for lawyers are less refined than general office productivity applications. Despite this, AI powered legal tech offerings are being released on a regular basis and there is no shortage of choices.

An early example of generative AI used in law is Harvey AI. This is a multi-tool platform that can assist lawyers in many practice areas with their daily workflows in multiple languages. Dentons has also launched FleetAI, a client-secure version of ChatGPT and many larger firms have their own variant in use as beta systems.

Larger firms are creating their own AI tools by licensing a suite of AI models and training them to use their own firm's data. An example would be the preparation of a deed or lease using the firm's existing precedents and past examples of special conditions. Third party documents could also be analysed, with the system preparing draft letters of advice and requests for alteration based on past examples of such material in the firm's data store.

Specialist AI legal tools are not just the domain of big law firms. Familiar legal software and library providers are also introducing AI tools into their offerings;

Legal research:

- Thomson Reuters;²⁴
- Leap²⁵

Workflow automation:

- Lexis + AI²⁶

Integration of third party AI tools:

- Actionstep integration with Smith.ai²⁷

Numerous aspects of legal work can be automated to some extent using existing software:

- analysing contracts, redlining clauses which may be of concern and suggesting changes
- undertaking legal research and summarising evidence
- extracting data from large volumes of unstructured information and loading it into practice management software
- preparing first drafts of correspondence
- drafting or summarising transaction documentation
- facilitating e-discovery and large scale merger document review²⁸

²² This propensity to “ingest” prompt information and material the system is trained on is the main reason AI must be carefully assessed before access to client information is allowed. See Section 11 of this Guide.

²³ See AI-enabled project cost management (n 18). This Stanford study [has been criticised](#), but the basic point remains; these tools are improving but are not yet good enough to use without careful scrutiny.

²⁴ ‘Westlaw Precision Australia’, *Thomson Reuters* (Web Page) <<https://www.thomsonreuters.com.au/en-au/products/westlaw-precision.html>>.

²⁵ ‘LawY’, *Leap* (Web Page) <<https://www.leap.com.au/exclusive-integrations/lawy/>>.

²⁶ ‘Lexis+AI’, *LexisNexis* (Web Page) <<https://www.lexisnexis.com/en-us/products/lexis-plus-ai.page>>.

²⁷ ‘Smith.ai’, *Actionstep* (Web Page) <<https://www.actionstep.com/integrations/smith.ai/>>.

²⁸ For example, extracting key data from hundreds of leases/ franchising agreements and inserting these into a comparison table.

- powering chatbots to answer FAQ's²⁹
- enhancing internal knowledge databases
- tracking information received from clients and following up anything missing
- predicting case outcomes.

Part C: Selection & Deployment

9. Where would we start?

The answer to this is highly specific to your firm's individual practice, however a few generalisations can be made. The procurement journey requires analysis of what you need, and at least a basic understanding of whether that objective can be met within budget and acceptable risk parameters.

Step 1: scope & try a discrete offering

Complex procurement decisions often leads to paralysis for small firms as you do not have the human resources. A small scope project can be a worthwhile starting point.

A workforce (or partners) that see a new technology as a way to free their time for more rewarding work will be a lot easier to bring on the journey.³⁰ Selecting and using a few test products is a good way to build this awareness and acceptance. This will give you the opportunity to develop a feel for the very real limitations in AI technology and learn how to use it.

Consider using the new system within a sandbox: deploying it for a single task or on part of the network that does not allow general data access and where errors do not flow directly into client work. This solves the three biggest problems lawyers have with AI: supervision, data access and data security.³¹

Step 2: determine what you will use it for

Only large firms really have the resources to build bespoke AI solutions. SME's are likely to be limited to purchasing "off the rack", albeit with customisation.³² One downside to this is the propensity to purchase a tool because it is available rather than because it is likely to really address an existing need.

The key to avoiding [FOMO driven purchasing](#)³³ is analysis of your existing and likely future needs. Be realistic in budgeting: early adoption of AI tools is likely to be cost neutral at best. The additional supervision and management workload may well offset any time savings in client file work. Getting the best from most tools may require quite a bit of adaption and technical expenditure.

However, you don't need to wait until the perfect tool emerges to start the adaption strategy. Automation and generative AI have the highest potential to increase efficiency where processes are already standardized and broken into workflows, or clusters of tasks.

Breaking up work into structured workflows or process maps will identify where time is spent, where the savings can be made³⁴ and then assist you to predict and measure return on investment.

²⁹ This would be regarded as a higher risk application as the interaction is difficult to supervise. For now, using the Chatbot to link your client or prospective client to pre-prepared answers would be best. Eg, rather than relying on the Chatbot to tell the client how long they need to live apart from their spouse to qualify for a divorce, you would use it to determine that information was being sought and direct the query to an appropriate FAQ response. For an example of what can go wrong if you allow free-form answers see: Maria Yagoda, 'Airline held liable for its chatbot giving passenger bad advice – what this means for travellers', *BBC* (Article, 23 February 2024).

³⁰ A popular change management system for SME's is Prosci's ADKAR model.

³¹ Of course, limited use such as this is unlikely to have a significant commercial impact, so it is really only an interim step.

³² Smaller enterprises can realistically purchase tuned models based on an existing LLM. Chat GPT has a secondary market for such personalized models, but you need to be realistic about the degree to which this will be adapted to your firm. It will still be stock – GPT under the hood.

³³ Gary Drenik, 'How SMBs can avoid generative AI FOMO', *Forbes* (18 April 2024).

³⁴ This kind of business process mapping can lead to many benefits, not just as part of an AI project. Identifying fraud risk, cybersecurity weak points, and duplication to name a few.

Don't just restrict this analysis to your internal processes either. Look at clients as well. A supervised legal pipeline that operates straight from their own network is very attractive option for a business client, and raises significant barriers to exit if they build their own systems around it.

Step 3: *plan for the client mix you will need*

If you invest heavily in systems to do specific work more efficiently you need enough of it to justify that investment. It will take time to position the firm to attract the right mix of clients³⁵ or manage existing clients to embrace any changes. Test any assumptions about this before being fully committed to a specific product.

If the plan is to service your existing clients more efficiently, work through a business model in which both clients and firm perceive and share in the benefit.

Using AI to expand into new areas of work rather than speed up work you already do introduces extra risk. If you are not expert in an area of law, do not have reliable precedents and do not have a system for identifying and tracking deadlines you will be far more vulnerable to errors in the auto-generated output.

Step 4: *choose the right vendor as well as the right tool*

Working with the right people will be an important factor in the success or failure of the project; probably as important as the features of the product itself. If your solution is to build not buy, the development team is even more critical.

Ensure the vendors fully map out what they will do and scope any additional work and budget necessary to deliver a fully operational system. An orphaned product left unsupported after hand-over will be very limited.

Step 5: *establish a risk management and governance framework*

Ongoing monitoring, adjustment and adaption will be essential.³⁶ Ensuring you are comfortable with the risk and the resources that will be needed to manage it is important before making the investment, rather than as an after-thought. It may be that you determine a less ambitious project is appropriate as an interim step until the technology and market matures.

10. Risk assessment and management

In broad overview, the risks that will need to be managed include:

- intellectual property: potential infringements of copyright, trademarks, patents and related rights, and misuse or disclosure of confidential information supplied by your clients;
- reliability: the potential for generative AI to produce misleading, inaccurate or false outputs that can be misconstrued or misapplied (See Sections 7.2 - 7.4 of this Guide);
- cash flow impact: initially, the supervision and governance requirements for the untried systems are likely to exceed time saved by use. Only some of that can be billed to the client;
- cybersecurity: vulnerabilities to hacking, data breaches, corruption of data sources and other malicious cyber activities;
- legal ethics concerns: competence, negligence claims, the potential erosion of professional independence, conflict of interest and fiduciary issues;³⁷

³⁵ In any event being strategic about the work you pursue is also likely to pay tangible dividends even if AI never enters the equation.

³⁶ See Section 10 of this Guide.

³⁷ See Queensland Law Society, 'Guidance Statement No.37 Artificial Intelligence in Legal Practice' (Guide, 30 May 2024).

- ESG and bias concerns: the possibility of AI models reflecting or amplifying societal biases present in their training data, leading to unfair or discriminatory results;
- reputation: if the use of generative AI may result in negative consequences for firm or clients brand damage could ensue;
- business continuity – the need to keep manual systems operating side-by-side in case of failure of either the new technology or (equally as likely) the commercial failure/acquisition of a vendor.

Not every problem needs to be solved before purchasing the system, but you should be reasonably assured that they can be solved before investing too much in the project.

The most critical pre-purchase issues are confidentiality, reliability, cybersecurity and ethics.

11. Protecting client interests

11.1. Intellectual property and clients

There are a number of significant intellectual property and confidentiality concerns with AI technology. Ensuring client interests are protected will be one of the most challenging aspects of the selection process. The more data access a system will have the more stringent the data use analysis will need to be.

There are two basic issues relating to clients:

- Maintaining *confidentiality*; and
- Ensuring that the *expertise or value* contained in client / firm data is not extracted by the tool.

Confidentiality is at risk if the client information is copied and exported or otherwise lost.

Data value is reduced if it is used to train a system that will be available to competitors.

Hypothetical Illustration:

Bank South³⁸ is accused by the ACCC of forcing junk mortgage insurance on customers. The bank has a large repository of property valuations and loan documentation. Bank South's solicitors propose to analyse the records using an AI tool for discovery and to generate insights into the disputed issues. Automating this analysis will potentially save hundreds of thousands of dollars and permit enquiries that would not otherwise be available to the defence team.

Confidentiality consideration:

The loan applications/approvals contain a great deal of confidential financial information. The Bank is subject to a wide range of regulatory obligations with respect to that data.³⁹ The solicitors will need to ensure that *any* system they entrust the data to (not just an AI system) is consistent with those regulatory obligations and their own duties.

Data value consideration:

If an AI model is given access to a big enough library of valuation and loan data it could be trained to auto-generate valuations or process loan applications without human involvement. This potential applies even to completely anonymised data, provided that it tracks the characteristics and outcome of the interaction.

Note: use of data for such training does not necessarily infringe IP rights as they are currently understood, so the protection regime needs to be contractual.

³⁸ A fictitious entity.

³⁹ Potentially: SOCI, GDPR, CPS 234, *Privacy Act 1988* (Cth), PCI-DSS, maintaining ISO 27001 certification requirements.

Practical issues to look at before using this specific tool:

- The most significant impact on the client would occur if the provider was allowed to re-sell the data. This is surprisingly prevalent.⁴⁰ This fact can be obscured by framing the privacy and data use elements of the user agreement around “confidential information” then adopting the position that if the information is de-identified it is no longer “confidential”.
- Some use of interaction data to make improvements (“tuning”) of the AI model can be expected. In fact, with a detailed project like this tuning to improve the output will likely be essential.
- Ideally, changes to the model should not be retained after the engagement ends.
- Even if the data is completely anonymised (which is recommended), the use of it to train someone else’s AI system can damage the client’s interests. This damage may be diffused if the data is only one drop in a large bucket, but it is still an important consideration.

The only way to find out what the vendor can and can’t do with data access is a careful analysis of the contract. Whether the vendor complies with the agreement is not an issue that an end user will be able to verify.

For the purpose of this overview, the key initial question will be whether the AI tool or provider is prepared to supply detailed information about data permissions and, if so, what that usage will be.

Despite early concerns⁴¹ about user information reaching the public domain through use of AI systems (with possible impact on client legal privilege),⁴² it would appear that this is not a regular occurrence. Most reported cases⁴³ where an AI tool disgorged confidential information appear to be the result of such information being scraped from a public web-page rather than harvested from user interaction and available due to ingestion into the model.⁴⁴ Even then, it took specialised methods to bypass the AI’s ethics wall to extract information which would not usually be included in output.

Other considerations include:

- use of prompt data and user interaction to train or tune the AI model
- permitted use of other data to which it has access
- whether it processes data on your network, on theirs or uses third party providers (any tool “powering” their system is likely to have its own data use regime)
- “ownership” of the prompt (input) and output data, and if so, what they can do with it
- whether the provider disclaims any responsibility to keep data confidential or secure at their end
- the degree to which anonymisation is an element of their privacy regime, and if so, whether they are transparent about what use can be made of anonymised data.

Absent clear answers on these points the default position should be that the product not be used at all or should only be used in a limited way.⁴⁵

⁴⁰ This may seem unlikely, but there is high demand for access to technical information to train AI with many data brokers emerging that are not overly scrupulous where the data came from: see, Australian Competition and Consumer Commission, *Digital Platform Services Inquiry – March 2024 report on data brokers* (Issues Paper, 10 July 2023). Unfortunately, this risk not only applies to AI providers. Any storage or online service provider might be tempted to commercialise data access, especially if providing a free service.

⁴¹ David C and Paul J, ‘ChatGPT and large language models: what’s the risk?’, *National Cyber Security Centre* (Blog Post, 14 March 2023).

⁴² Isabel Gottlieb, ‘Generative AI Use Poses Threats to Attorney-Client Privilege’, *Bloomberg Law* (Blog Post, 23 January 2024). Note, the premise of this article is questionable: that by design a Gen AI model is likely repeat information from one user’s query to the next.

⁴³ See, eg, Jeremy White, ‘Personal Information Exploit With OpenAI’s ChatGPT’, *New York Times* (Article, 22 December 2023).

⁴⁴ An exception is an incident in which Open AI’s chat GPT showed some users the prompt headings submitted by other users. What the prompts contained or who had asked those questions was not disclosed. This error arose from misconfiguration and was not specific to AI and any cloud provider could have done something similar. Note, however, the observations as to cybersecurity under Section 12 of this Guide. There are also reports of [account phishing](#), again not specific to AI providers, and it has been reported that bespoke GPT’s created soon after these became available in beta [were not secure](#), probably because there was no standard tool used to upload information and additional risks were introduced [via the third party ecosystem](#).

⁴⁵ Unless there is some reason to think that the data practices will change, a tool that you cannot analyse properly is not worth the risk.

For example, it may be appropriate to only submit queries which do not identify your client, and not upload documents containing sensitive data. Certainly the tool should not be given broad access to your network or cloud storage system but only used as a wholly external service or applied to a limited data set that does not contain confidential information.

This kind of restricted use is a good way to start using a tool without getting bogged down in detailed analysis that is required for more comprehensive integration into your firm's work, but double handling and inefficiency will detract significantly from productivity gains.

Point to watch: An AI tool might be supplied by a third party and integrated with the platform you are already using. As such, it might have a completely separate privacy regime. Watch out for additional click-through user agreements when activating the tool.

11.2. Intellectual property and third parties

Legal practitioners should be wary regarding the subsistence or ownership of copyright and other intellectual property rights in any AI-generated outputs, and the risk of infringement of third-party intellectual property. Current case law is unclear as to the extent to which machine-generated content will be protected by copyright or other intellectual property rights, and there is the potential risk that use of AI-generated content may infringe third party intellectual property rights.

While it is possible that an end user might be implicated in an AI provider's misuse of copyrighted information, this is a low risk unless the user is training their own model rather than using someone else's. The more likely risk in this context is business continuity. If a significant copyright claim successfully impacts the financial position of the provider

What intellectual property rights does the firm acquire in the output? Can it convey this to the client? Can the provider alter these rights unilaterally, either prospectively or retrospectively?

11.3. Cybersecurity

Many AI tools operate in the cloud, or transfer data internationally for processing. This may apply even in the case of "walled off" products specific to your firm that are not available for public access.

AI providers are as likely to be subject to data loss and intrusion as any other online service; in fact somewhat more so due to the rapidly evolving technical framework, intersection between multiple tools and breakneck commercial development in a resource constrained environment.

The same considerations that apply to any other cloud provider should be considered when entrusting client data to AI:

- Where is the data processed and by whom?
- Do they have a meaningful data protection certification?⁴⁶

Points to note:

- Just because a service runs on a recognized platform (such as AWS) does not mean it is secure, any more than a poorly maintained fleet of trucks is safe on a good road.
- Any IT infrastructure is subject to these concerns. Although AI providers may be less secure than more mature providers (see above) there is no special characteristic of cloud based AI services that make them inherently more dangerous than other cloud systems. While AI Cybersecurity is a concern, it is no more a concern than it would be for (say) a new backup system.

⁴⁶ Appropriate standards are ISO27001 (and/or ISO27017 – cloud) and SOC2. Frameworks such as NIST, the ISM & Essential 8 are not certifications but indicate the control measures the entity will be using. Purchasing approval regimes for governments (Such as Fedramp, IRAP & NIST 80-53 or the various Australian State Government procurement guidelines) are a useful indicator that a demonstrated standard can be evidenced.

Part D: Using AI successfully

12. Managing reliability limitations

We have already examined some of the strategies used to improve reliability (see 7.1 *et seq*). No combination of these strategies as at mid-2024 is good enough to obviate the need for line-by-line checking of most output, especially when the system is new.

AI is very much like a newly graduated lawyer. Initially, it probably takes more effort to delegate a task than it would to do it yourself. Over time, improvements in the system and the way it is used should improve reliability and generate a better understanding of the strengths and limitations. At that point supervision can relax somewhat, but ongoing scrutiny will be essential.

Governance will require management of the immediate output, of the overall system and a way to benchmark accuracy and track problems.

12.1. Output supervision

In most cases, all automatically generated output will require someone to carefully scrutinize drafts line-by-line and approve them, thereby accepting responsibility for the accuracy of the content.

There may be some classes of material where that degree of accuracy is not required. It will be important for the organisation to decide how unverified material will be used and ensure that it is not misapplied later.

Example: AI translation of a large body of foreign language records might be appropriate to review the material and decide what needs to be translated by an accredited translator. It would not be appropriate to annexe AI translations to an affidavit in most circumstances. A system to mark the provenance of translations, explain appropriate / inappropriate use to staff and make sure such guidance is applied would be required.

The supervisor at this level will need to know enough about the applicable law and the facts of the matter to be able to identify errors. This requires the maintenance of a discrete matter file or record so the human supervisor can understand the context of any material they are looking at.

Example: checking a will draft requires an understanding of the instructions, and whether the instructions interpreted by an intake AI are consistent with the client's circumstances.

Once the practice is satisfied that the error rate is acceptably low, appropriate scrutiny may potentially be limited to audit and spot checks, but this will be a case-by-case decision. Unfortunately, updates to the model / process might result in a drop in accuracy so the supervision (review) system will need to regular.

12.2. Process supervision

Consider a supervisory committee or, in the case of a small firm, a senior person within the firm tasked to maintain an overview of any AI tools in use. This is a broader role than checking output accuracy, and the supervisors should be sufficiently senior to be able to intervene if a structural problem emerges. New, complex delivery systems carry with them an inherent risk of reputational damage. This risk can be mitigated if the practice identifies and addresses a problem early.

Co-ordination is important to ensure clarity of risk management roles and responsibilities, and a structure to allow lessons to be learned and improvement undertaken.

This supervisory and coordination role is a governance function, not a purely technical one.

12.3. Accountability structures

When failures occur, the organisation will need to demonstrate that these did not arise because of indifference to appropriate management or ethical responsibility. Factors which may be relevant include:

- whether a risk assessment and management framework⁴⁷ was applied during procurement and subsequently maintained during the implementation phase;
- whether risk management was documented and all relevant parties were aware of their roles;
- whether risks to third parties as well as the firm were considered;
- the degree to which persons potentially subject to adverse outcomes were identified, gave their informed consent, shared in the benefits or were told that AI generated content was used in their matter;
- audit records demonstrating that error rates for automated work were monitored, who took relevant decisions, whether there were changes to the system, when product versions and updates were effected;
- ensuring system logging is turned on;
- how problems (including errors and complaints) were responded to; and
- how accountability within the organisation is managed.

A record of the inputs (including prompts, what clients were asked and what they said in response) should be available. Ideally this should be on the client file as this might need to be considered years later, but at the very least should be capable of accurate reconstruction.

If a generative AI tool is used and the vendor does not provide a history of use, it is advisable that you document all inputs, outputs, and system errors to ensure that the use of the tool can be monitored as appropriate. Simply accepting vendor claims without ongoing verification is unlikely to satisfy the practical onus of answering critics if an issue arises.

13. Disclosure when using AI systems

13.1. Disclosure by lawyers to the court

In a notorious example from the United States, a judge imposed sanctions on two New York lawyers who submitted a legal brief that included six fictitious case citations generated by ChatGPT.⁴⁸ Numerous instances since then amply illustrate the basic precept: do not use an auto-generated document in legal proceedings without ensuring that it has been carefully checked for accuracy. Particular care is required with respect to submissions containing case citations, evidence and witness statements.

In response to similar incidents, Courts in Australia⁴⁹ and elsewhere⁵⁰ have issued AI specific practice directions emphasising the importance of maintaining a human in the loop or some other assurance framework to reduce error risk.

It is recommended that **all** AI generated material (either wholly or in part) be marked as such when produced. When subsequent quality assurance steps are completed, the watermark could be removed. Until the use of AI is more mainstream it may be prudent to maintain a disclosure watermark even once

⁴⁷ No standardised procurement and governance framework has emerged in AI as yet, especially not for a small / mid law firm purchasing a third party tool. For larger projects or enterprises there are a large number of useful guidelines available. Some examples of structured decision making systems include: from the US ([CIO.Gov](#)), Canada ([Impact Assessment Tool](#)), the OECD (Impact Assessment Tool), Finland (University of Turku: [AI lifecycle and governance](#)) and commercial providers ([IBM](#)).

⁴⁸ David Bowles, “Advo-cat’ for the GPT Age: US lawyer faces sanction for filing a submission by ChatGPT”, *QLS Proctor* (Article, 28 May 2023).

⁴⁹ In Queensland, see ‘Using Generative AI’, *Queensland Courts* (Web Page, 14 May 2024); in Victoria, see ‘Guidelines for Litigants: Responsible Use of Artificial Intelligence in Litigation’, *Supreme Court of Victoria* (Web Page, May 2024); although not AI specific, also see, Federal Court of Australia, *General Practice Direction ‘Technology and the Court Practice Note (GPN-TECH)’* (25 October 2016).

⁵⁰ Shane Budden, ‘Hands up who’s using AI?’, *QLS Proctor* (Article, 7 June 2023).

the material has been fully checked, but there is no consensus on this point.⁵¹ Formal “disclosure” to the Court of AI use is not required in all cases.⁵²

Citing AI generated opinion is required as an application of this principle. However, it is highly unlikely that ChatGPT or any of the other LLM’s would be a credible authority for an assertion of the law without specific reference to the source material that provides the foundation.

There may be some classes of document produced by AI in which line-by-line verification will not be undertaken. Examples might include:

- translations of background or secondary material
- disclosure prepared using an AI tool in accordance with a pre-agreed protocol or order
- a summary or index of a large document bundle
- a transcript of voluminous audio or video surveillance recordings

The provenance of all such material should be clearly marked on the document itself, and the practitioner must be in a position to answer questions about any quality assurance process that was followed.

13.2. Disclosure by solicitors to their clients

As fiduciaries, solicitors are not just service providers.⁵³ Clients trust us to be competent, to act in their interests⁵⁴ and include them in decision making as appropriate.

There are two aspects to that duty: (1) we must exercise judgment to protect client interests and (2) where appropriate we should disclose relevant issues to the client so they can make their own choices.

In theory, a solicitor must make,

“...full and frank disclosure to the client of all information known to the solicitor which the client should know...and if there be aspects of the (retainer) ... in which the solicitor is in a position of advantage vis-à-vis the client those matters should be brought ... to the attention of the client so that the client can decide whether she should enter into the agreement.”⁵⁵

In practice it is impossible for the client to be consulted about every part of the system that will be used to complete their work and a value judgment must be made.

- Is the use of the AI significant?
- What is the potential risk to the client?
- Does the risk management process leave a tangible (if acceptable) residual risk or is it fanciful or remote?

The further the proposed method of work deviates from the “usual” or traditional practice the more likely it is that disclosure would be required.⁵⁶ The retainer agreement (or website, if that is what the client is interacting with) should state what AI tools are used, but simply inserting a list of products somewhere in the retainer does not satisfy the professional obligation. For the most part, retail clients’ primary

⁵¹Theoretically, the process of checking and adoption erases the error risk that AI use may have introduced. However, especially if dealing with a junior practitioner, a judicial officer may prefer a warning that material being submitted was prepared with AI involvement. Several US jurisdictions which originally required “disclosure” are now winding that back in favour of a risk management and responsibility model: see, eg, United States Court of Appeals Fifth Judicial Circuit, *Court Decision on Proposed Rule Amendment 5th Cir. Rule 32.3 and Form 6*, 22 November 2023.

⁵² New Zealand guidelines state that disclosure is usually not required; Victorian guidelines that such disclosure should be made “where appropriate”.

⁵³ Queensland Law Society, *Australian Solicitors’ Conduct Rules* (at 1 June 2012) rr 4.1.1 and 4.1.3.

⁵⁴ *Ibid* r 4.1.4.

⁵⁵ *Law Society of NSW v Foreman* (1994) 34 NSWLR 408, [435] (Mahoney JA).

⁵⁶ *Re Morris Fletcher & Cross’s Bill of Costs* [1997] 2 Qd R 228, [244]-[255] (Fryberg J).

concerns with AI will be confidentiality, whereas business clients will also potentially be concerned about their IP (see hypothetical illustration in Section 11 of this Guide).

13.3. Disclosure by clients to their lawyer

When using client supplied documents, data or expert reports a solicitor should enquire whether AI was used in their production if this seems likely. Where the client supplies drafts of statements, timelines or summaries of business records the issue should probably be raised by default.

In *DPP v Khan*,⁵⁷ Mossop J criticised defence counsel who had tendered a character reference which appeared to have been drafted or translated (or both) using an AI tool. His Honour stated that using AI to draft or translate such a document was “undesirable” and Counsel should make appropriate enquiries and be in a position to warn the court if such use had been made.

It was unclear whether his Honour considered that such enquiries should now be made as a matter of course or only where eccentricities in phrasing or content should put the lawyer on notice that this may be likely.

In a case from the United States, a lawyer representing a disbarred colleague adopted research undertaken by the client without knowing that it was generated by Google’s Bard LLM.⁵⁸ The judge described the lawyer’s failure to check the provenance of the material as “negligent, perhaps grossly negligent” despite the fact that the client was an experienced lawyer themselves.

Legal practitioners will be in the front line in the war on deepfake⁵⁹ evidence in Australian courts, an issue that will become increasingly crucial over the next few years. Deepfake enabled fraud is now occurring,⁶⁰ and lawyers will need to be vigilant to detect forgeries both in defence of their own businesses and to avoid being the unwitting instruments of perpetrators.⁶¹

13.4. Disclosure by lawyers to their employer.

For the reasons discussed in Section 4 of this Guide, it is entirely inappropriate for law firm employees to use AI without approval. Fraudulently passing off AI generated material as bespoke work and charging an inflated number of hours to the client would be serious misconduct on the part of an employee solicitor.

14. ESG and bias

Machine learning output is dependent on the quality⁶² and suitability of the data used to train it. AI is also an inherently amoral artefact that reflects the world as contained in that data or the task set in the prompt.⁶³

Bias can arise from various stages of the AI development lifecycle, including data collection, algorithm design, and deployment. Bias in AI can perpetuate and even exacerbate existing inequalities, leading to unfair treatment of individuals based on race, gender, age, or other protected characteristics.

For instance, biased training data can result in [discriminatory hiring practices](#) or [skewed loan approval](#) processes, undermining the fairness and integrity of business operations. Moreover, the lack of diverse perspectives in AI development teams can further entrench these biases, as homogeneous groups may overlook critical ethical considerations. To address these risks, businesses must prioritize the use of

⁵⁷ [2024] ACT SC 19.

⁵⁸ Jonathan Stempel, ‘Michael Cohen will not face sanctions after generating fake cases with AI’, *Reuters* (Article, 21 March 2021).

⁵⁹ Audio, video, documents or images created using generative AI as forged representations of reality

⁶⁰ Jeannie Patterson, ‘So, you’ve been scammed by a deepfake. what can you do?’, *The University of Melbourne* (Article, 26 February 2024).

⁶¹ Douglas McGregor and Christy Foster, ‘Deepfakes, and how to avoid them’ (2020) 65(3) *Law Society of Scotland* 1.

⁶² In this context, “quality” includes both accuracy and whether the data is appropriate to train a tool which will be used in a particular way.

⁶³ For example, as at June 2024 GPT 4 would “happily” prepare an essay justifying the execution of convicted witches or homosexuals in jurisdictions in which these activities are criminalised, citing data from *Human Rights Watch*, *Amnesty International* and the United Nations.

inclusive and representative datasets, implement rigorous bias detection and mitigation strategies, and foster diverse AI teams. By doing so, they can enhance the fairness and reliability of their AI systems, ensuring that technological advancements contribute positively to society.

The computational power and electricity use of AI infrastructure is significant. On some estimates AI may be consuming between 5% - 16% of US power generation by the end of the decade.⁶⁴ Calculating net effect is complex, but there is no doubt this technology will contribute significantly to greenhouse gas emissions.

Like all new technologies there will be winners and losers. There is likely to be poor alignment between those who benefit most from AI adoption and those who bear the cost, either in terms of the training data being used, impact on employment or loss of access to human supplied services.

15. Remuneration

If a firm is charging on a time basis, only time spent on a matter can be charged to the client. Time savings belong to the client, although any time spent checking that the system has undertaken that client's work correctly could be charged.⁶⁵ For that reason heavily automated work is not a good fit with time based billing, although AI might be very useful to speed up work which solicitors typically don't fully time record such as preparing first drafts,⁶⁶ research and double-checking.

If your access to the AI system is billed on a "per matter" basis that fee can potentially be charged to the client as a disbursement but it should be noted with other such items specifically in the Costs Agreement. It is permissible to agree with the client that a particular cost shall be treated as an outlay provided it meets the usual criteria⁶⁷ and the client's agreement is "informed".⁶⁸

If access is calculated monthly or annually the charges cannot be properly characterised as a disbursement and must be absorbed by the firm as an overhead.⁶⁹

Realistically, the payback period for AI in a law firm may not be especially rapid.

16. AI use by in-house and government lawyers

Lawyers working in State Public sector bodies need to consider generative AI in context of the Queensland Government AI Enterprise Architecture guideline⁷⁰ which applies to any AI systems that public bodies are already using or that others may be developing or using on their behalf. In an enquiry submission, the Qld Information Commissioner has also provided detailed observations about public sector ethical AI use.⁷¹ See also the interim guidance from the Commonwealth Digital Transformation Agency.⁷²

⁶⁴ Jordan Aljbour and Poorvi Patel, *Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption*, *Electric Power Research Institute* (Report, 5 May 2024)

⁶⁵ This only extends to work done verifying the output in the client's particular matter, not general audit or accuracy calibration.

⁶⁶ While there is no specific limitation on this, a firm without access to a good set of precedents will struggle commercially to charge for the time wasted as a result.

⁶⁷ The three requirements for a disbursement are that the amount has been paid out to a third party, without undisclosed surcharge, and is capable of specific and accurate attribution to a specific client's matter. For further information, see, 'Legal costs, outlays and disbursements and billing', *Legal Services Commission* (Guideline, 2024). As an illustration, phone call outlays could be charged if the number of calls for a particular client were tracked by the billing system whereas a pro-rated percentage of monthly telephone costs could not.

⁶⁸ See *Council of the Queensland Law Society v Roche* [2004] 2 Qd.R 574, [32].

⁶⁹ *Equuscorp Pty Ltd v Wilmouth Field Warne* (a firm) (No. 4) [2006] VSC 28, [53].

⁷⁰ 'Use of generative AI in Queensland Government', *QGCDG Data and Information Services* (Guideline, August 2023).

⁷¹ Office of the Information Commissioner, Submission to CSRIO, *Australia's AI Ethics Framework* (31 May 2019).

⁷² 'Policy for the responsible use of AI in government', *Australian Government Digital Transformation Agency* (Web Page, 1 September 2024).

For public enterprise, ESG, bias and accountability considerations are likely to be more stringent than for private sector small / medium enterprises and managing these risks would be an essential foundation of any AI project.

When generative AI tools are used, it would be advisable to ask the generative AI vendors for more information about:

- what datasets they used
- how such datasets were acquired and fed into their systems for training
- who is tasked with data labelling and training the generative AI system.

As the [layers of data and generative AI platforms increase](#),⁷³ computational and bureaucratic systems become increasingly complex.

This means that, when harm does occur, figuring out why and who is responsible becomes more challenging, with further complication as the technology continues to evolve.

17. Business continuity

Consider whether you need or will have long-term support from a vendor, as well as an exit plan should a generative AI tool be adopted but the vendor exits the market.

⁷³ 'Ethical Prompts: Professionalism, ethics, and ChatGPT', *Harvard Law School Centre on the Legal Profession* (Article, April 2023).

Annexure 1: Procurement and deployment checklist (COBIT) framework

It should be noted that a less ambitious project requires a less rigorous analysis. However, even if the proposed use is to be limited, careful consideration of data protection and privacy issues is required.

What do we need?

1. Identify Business Needs and Requirements
 - List all the work the firm does, including administrative work.
 - Determine which tasks could most easily broken down into predictable workflows.
 - Check whether any existing tools can automate some or all of these tasks.
 - Gather input from stakeholders (partners, solicitors and paralegals, administrative personnel) about work that they find time consuming and would like to automate.
 - Define use cases for AI within the firm. Consider a smaller scope project either as a pilot or an experiment.
 - List essential and non-essential requirements.
 - Ensure selection team is aware of regulatory and ethics requirements.

Evaluation and Selection

2. Evaluate
 - Research available AI tools in the market.
 - Conduct vendor assessments and comparisons.
 - Review case studies, references, and user reviews.
 - Evaluate tools against predefined criteria (functionality, scalability, integration).
 - Conduct pilot testing with selected tools.
 - Select the most suitable AI tool based on evaluation and pilot testing results.
3. Select
 - Undertake detailed data use assessment of shortlisted vendor/s.
 - Assess their support offering, bespoke options, consulting services.
 - Prepare a budget, including time allocation for supervision and governance
 - Undertake detailed contract and offer analysis, including product deployment package, training options.

Risk Management

4. Risk Assessment and Management
 - Identify potential risks (data privacy, security, ethical concerns).
 - Develop risk mitigation process.
 - Implement a risk management framework.

Establish Governance Structures

5. Allocate responsibilities, create policies
 - Define governance roles and responsibilities.
 - Create a high level supervisor / steering committee.
 - Allocate responsibility for output supervision
 - Establish an AI governance policy tailored to the selected tool.
 - Develop procedures for compliance, implementing risk management framework, and performance monitoring.

Implementation Planning

6. Develop Implementation Plan
 - Create a detailed implementation roadmap.
 - Define timelines, milestones, and deliverables.
 - Allocate resources and assign responsibilities.
 - Develop a communication plan for stakeholders.

7. Conduct Training and Development
 - Identify training needs for staff.
 - Develop training materials and programs.
 - Conduct training sessions and workshops.
 - Provide ongoing support and resources.

Performance Measurement

8. Monitor and Measure Performance
 - Define key performance indicators (KPIs) for AI tool performance.
 - Implement monitoring tools and processes.
 - Regularly review performance against KPIs.
 - Collect and analyse user feedback.

Continuous Improvement

9. Continuous Improvement
 - Establish a feedback loop for continuous improvement.
 - Regularly update AI tools based on feedback and performance data.
 - Stay updated with new AI technologies and practices.
 - Conduct periodic reviews of AI governance and practices.