

A SOLICITOR'S GUIDE TO RESPONSIBLE USE OF ARTIFICIAL INTELLIGENCE



Artificial intelligence (AI) is a form of science and engineering that is concerned with making machines work in an intelligent way, a way that mimics the problem-solving and decision-making capabilities of the human mind. Common forms of AI include automatic speech recognition, such as Apple's Siri and Amazon's Alexa, and messaging chatbots, such as ChatGPT and Google Bard.

AI systems are powerful tools and can be of great assistance in legal practice if used correctly and within the goal posts of solicitors' ethical and professional obligations. For solicitors to be able to gauge and assess whether AI is being used within these goal posts, it is recommended they have a general understanding of how the technology works.

This guideline focusses on generative AI and sets out potential issues that arise in legal practice with the use of generative AI systems. Importantly, this guideline seeks to remind legal practitioners of their relevant legal and ethical obligations that may arise when engaging AI tools to perform legal tasks.

While these guidelines are mainly concerned with the use of generative AI systems, in essence, it conveys a simple principle that can be applied to any new technology in legal practice:

- Understand the tool you are using and know how it can assist your legal practice.
- Know what your professional and ethical obligations are and how they are to be applied when using any given tool.

Knowing your tools

The appropriate use of any new technology in legal practice first requires a basic understanding of how the technology works. This ensures solicitors avoid breaching their professional and ethical obligations.

For example, take the obligation not to disclose confidential client information.¹ Before electronic documents and communications, standard practice for safeguarding confidential client information involved filing documents into a lockable filing cabinet or safe. With the advent of electronic communications, and more recently with cloud computing, antivirus software and data encryption has become part and parcel of securing confidential client information.

How a piece of technology works determines the measures we need to put in place for its appropriate and ethical use. This is no different when it comes to generative AI systems, whether it's a publicly available system like ChatGPT or a proprietary system commissioned for in-house use or for sale.

What is Generative AI?

Generative AI is a learning model that produces outputs (this can be text or images or sound) when given a prompt. The outputs are based on probabilistic modelling applied to a set of data (the underlying data). The AI learns from the aggregate underlying data and generates content based on the data it is fed. It is generally trained over a large corpus of data, by recognising relationships between syntax. Generative AI is considered a new iteration of AI intelligence.

How can generative AI assist in legal practice?

Although in its relative infancy as a tool of legal practice, generative AI has already been known to assist lawyers with many tasks, including streamlining the intake process, undertaking legal research, recording minutes of meetings, drafting and reviewing contracts and summarising cases. Generative AI has also been used to systematically analyse case law to provide solicitors with helpful insights into the potential outcomes of a particular case.



Given that it has a range of applications, generative-AI can potentially improve your client engagement and add substantial benefits to your practice if used to, say, assist in the explanation and nature of the processes a client is about to embark on. It can add a personalised experience to regular online engagement, by gathering information about what kind of referral services your client/s may need.

The abilities of generative AI can make legal services more affordable, lowering cost barriers to particular clientele. Discrete pieces of legal work may take less time; with less costs being passed onto the client. High-volume, more repeatable tasks are the most likely to be positively impacted, for example, when large volumes of documents need to be reviewed quickly.

The remarkable capacity of AI is transforming the legal landscape and, no doubt, will become increasingly common in legal practice across Australia and the world. Best practice would therefore be to get across the technology early, so you know where you can apply it and learn how you can use it to unbundle your legal services.

As solicitors welcome and embrace this new and efficient way of working, it is imperative that the issues arising from the use of AI remain front of mind with a clear understanding of its limitations.

Issues arising from the use of generative AI

Accuracy

At a glance, generative AI systems may resemble search engines (like the Google Search Engine) but it is not, and knowing the distinction is important because it changes how you use the technology to assist your legal practice. Of the distinctions between a search engine and generative AI, the most relevant for solicitors to consider is the way in which these different systems generate content, which in turn affects the accuracy of information generated.

While both systems produce information based on prompts given, a search engine is essentially a very powerful cataloging and indexing system, which uses complex algorithms to retrieve relevant information from billions of web pages based on a user's query.

On the other hand, generative AI is designed to comprehend a user's prompt using AI and produce a response in conversational, everyday language. While search engine systems generally draw on web-based sources and links its users to those sources, generative AI tools are concerned with interpreting information and producing appropriate conversational responses, having drawn on web-based sources.

The accuracy of information produced by a generative AI system such as ChatGPT can, therefore, be compromised in the system's pursuit of facilitating human-like conversations in its delivery of information. For example, it has been known to produce fictitious case law.

Other aspects of generative AI that can affect accuracy include hallucinations, where the algorithm simply makes things up.

While it may be quick and easy, if you are using it in an area of law that is not your area of specialty, you must thoroughly check its outputs to avoid compromising your integrity.

Generative AI is binary, so it does not understand who the client is. It cannot weigh up 'best interests,' optimise outcomes or provide explanations or reasoning. It lacks the skill of critical thinking and cannot determine relevance. Additionally, outputs can become resistant to change, so it is best to use generated output to 'fact-check' only. Humans are accountable, but generative AI is not. For this reason, it must undergo conscious critiquing while contemplating both our ethical and professional obligations.

Bias

Another issue lawyers need to be cognisant of when using generative AI is 'AI bias'. In computer science terminology, AI bias refers to the tendency for an AI system to produce skewed results that are systematically prejudiced against individuals or groups. AI bias can be caused by a number of factors, including unrepresentative, limited or incomplete training data sets. Put simply, biased data that is used to train AI systems may lead to biased information being generated.

Intellectual property

In addition to the quality of information, solicitors also need to be mindful of ownership rights attached to information produced by the various generative AI models. Material produced by generative AI is usually based on existing data or creations, which means that the use of such material may be infringing on someone else's copyright.

Privacy and data security

More generally, lawyers should consider privacy and data security frameworks to support any generative AI system they use. Given the relative infancy of publicly available generative AI systems, such systems may have inadequate checks and balances regarding information sharing and privacy.

For example, some generative AI systems may share user data with third parties without explicit consent, or for purposes beyond what was initially communicated. There may also be AI systems with inadequate data anonymisation techniques. Like any other tool publicly available over the Internet, generative AI systems are at risk of data breaches.

Practitioners should not necessarily rule out the use of AI systems because of these issues. Rather, these issues highlight the need for appropriate checks and balances in legal practice management systems and processes to ensure human oversight is everywhere, at all levels of an organisation, so to avoid malpractice and harm to clients. Practitioners should always oversee and review AI-generated work to ensure accuracy and ethical compliance. Having appropriate cyber security measures in place is also critical.

Good accountability and governance are key to maintaining our professional responsibilities as well as protect clients' privacy and data security.

Relevant rules to consider under the Legal Profession Uniform Law Australian Solicitors' Conduct Rules 2015 (Conduct Rules)

When using generative AI systems solicitors should be particularly mindful of their obligations under the following Conduct Rules:

Rule 4 – Competence, integrity and honesty

Solicitors are responsible for the accuracy of any advice they give, irrespective of where they may have done their research. As with any other research, for full transparency, it may be appropriate to disclose to clients the fact that a generative AI program was used.

As previously mentioned, generative AI can create fabricated material, including fictitious case law. The materials it generates can therefore be completely wrong, inaccurate or outdated. There is also a risk that generative AI may contain plagiarised material or that it breaches intellectual property rights. Be aware that the sources of its information are often obscure.

Rule 9 – Confidentiality

Generative AI takes information inputted and uses the information to learn and then discloses that information to other users. Therefore, any placing of client confidential information into a publicly available generative AI system is akin to putting it in the public domain. This is likely to be a breach of confidentiality and clients may lose privilege.

Rule 17 – Independence (avoidance of personal bias)

Solicitors are not a mere mouthpiece for their clients. Solicitors have an obligation to exercise their best judgment independently, irrespective of the views of any generative AI system.

Rule 19 – Duty to the court

Solicitors must ensure they do not mislead or deceive the Court, even if inadvertently. The validity of any material presented to the Court needs to be tested by solicitors, whether or not that material has been produced by generative AI. When using AI, solicitors should be particularly cautious, given the limitations discussed above.

Solicitors should not rely on generative AI to verify sources produced by AI. This has been known to fail.²

Rule 37 – Supervision of legal services

A solicitor who is charged with supervising legal practice by others, and the provision of legal services generally, needs to be particularly cognisant of the risks of generative AI in legal practice. Being able to exercise reasonable supervision where AI is employed requires critical evaluation of the accuracy and completeness of AI's outputs. A more junior practitioner may not have the experience or knowledge to make that evaluation

Practice management tips

When a solicitor uses generative AI to assist in their legal practice, they should employ the same level of care and caution as they would to any legal assistant or paralegal. Solicitors must exercise independent forensic judgement, based on their own training, experience and research, and review and edit any 'product' to be confident it is reliable and correct. Solicitors need to understand it and may need to explain it or provide additional supporting materials.

Clients are entitled to expect that any work done by a solicitor is the solicitor's own work, reflecting the solicitor's experience, knowledge, application and judgment. AI must, therefore, be used responsibly to supplement (rather than substitute) the legal services on offer.

Organisations can benefit from AI while limiting their risks. However, the only way of managing these risks is to know of them. Incorporating and/or adopting AI tools needs human oversight everywhere, at all levels of your business. Gaps in content need to be identified. In terms of privacy and data security, no sensitive information should be inputted into the tool, particularly client information. Once in, it is hard to remove.

Educating your employees on the risks and benefits of generative AI is essential, given the emerging responsibilities on lawyers to develop guidance on AI tools. Some thought may need to go into different human resource considerations at entry level. For example, basic 'prompt engineering' may be on your skills list when hiring graduates because, while generative AI offers a neutral position, whether it is capable of arguing two sides of a case is yet to be seen. A young graduate who knows what prompts to input to produce another perspective on earlier outputs may prove invaluable.

Simplistic prompt engineering methods can include the accurate use of prompt wording. Being succinct with the words you input will improve specificity. For example, by explicitly stating that you want the tool to act, say, as a forensic expert for a particular criminal matter will provide narrower and more targeted outputs. Similarly, using (both) positive and negative prompting will (either) encourage the model to include specific outputs and generate certain responses or discourage them. These prompt varying techniques can influence the direction and quality of information generated.

When a solicitor uses generative AI to assist in their legal practice, they should employ the same level of care and caution as they would to any legal assistant or paralegal.

How to use Generative AI efficiently and safely

Used appropriately, generative AI can be an effective tool to assist you to provide legal services. Here are some examples of how generative AI can assist your legal practice, as well as regulatory and ethical considerations that may arise.

Confirmatory communications

Generative AI can generate conversational responses to queries. It can therefore be helpful for more common writing tasks. Consider whether your client should know you plan to use generative AI for their advice, and also consider whether having access to such a system has a cyber-risk element.

Initial legal research

As a preliminary research tool, generative AI may help to give a broad picture, so long as users are:

- very careful not to disclose any confidential information to it.
- aware of its limitations. Generative AI should not be blindly relied upon, even at this preliminary stage, noting it has been known to fabricate case references.
- careful not to rely on it to move into areas of law where the user does not have sufficient expertise.

Obtaining further particulars from our client

Using generative AI for this is not a complete solution, as it cannot read humans like humans can. However, generative AI tools can offer a more personalised and tailored experience, by seeking further particulars in a more interactive way than simply filling in an online form.

In-depth legal research

Caution must be exercised. Refer to the above discussion on how generative AI works and the impact it has on the responses it may generate when conducting legal research. Issues requiring considerations include:

- Confidentiality (an applicable consideration whenever using generative AI).
- The source/s of the generative AI's information (noting that, at least for public generative IT, it would likely be impossible to trace).
- Bias inherent in generative AI. Such bias can affect the response generated in many ways, and produce information that is:
 - Inaccurate
 - out of date
 - geographically limited
 - culturally limited, or
 - not sufficiently specific to legal issues
- Consideration should also be given to whether there could be a breach of legal rights given you do not know the source of the information, for example, are you breaching:
 - intellectual property rights
 - privacy considerations
 - human rights
 - rights to procedural fairness

Analysing legal research

Practitioners are responsible for the accuracy of the advice they give. That responsibility extends to (both) client/s and to the administration of justice. Practitioners must therefore check/analyse any research generated by AI, just as you would any other information provided by another.

Crafting the advice

Generative AI may help with setting out a basic framework within which to communicate any carefully crafted advice that you have prepared. It can also help with providing the contra argument – again using its ability to generate words from patterns.

Draft article or presentation

Generative AI can be very good at pulling together research notes into a coherent form and translating them to slides. Again, though, practitioners must be very careful to ensure accuracy and confidentiality.

Supervision

Generative AI is probably best compared to a paralegal or an administrative assistant, depending on the tasks it assists with. Practitioners will need to ensure they supervise the work it produces, thoroughly review it and make any necessary changes and/or amendments, depending on the particular circumstances.

Summary

Generative AI is likely here to stay. Practitioners should neither avoid generative AI completely nor embrace it without first understanding its limitations and giving critical thought to maintaining their professional obligations while using it. Generative AI presents a host of new and positive opportunities for practitioners and their legal practice. Key to capitalising on these opportunities is understanding the inherent limitations of the everchanging and evolving nature of generative AI and steadfast adherence to solicitor professional and ethical obligations.

-
- 1 Legal Profession Uniform Law Australian Solicitors' Conduct Rules 2015, rule 9.
 - 2 The lawyer who appeared before the United States District Court for the Southern District of New York in the matter of *Mata v. Avianca, Inc.*, 1:22-cv-01461, (S.D.N.Y.) gave affidavit evidence to the Court to the effect that he had asked ChatGPT whether the cases it cited were real and ChatGPT answered the cases existed and could be found on legal research databases such as Westlaw and LexisNexis; also refer to Benjamin Weiser (The New York Times), *Here's What Happens You're your Lawyer Uses ChatGPT* (27 May 2023) <https://www.nytimes.com/2023/05/27/nyregion/avianca-airline-lawsuit-chatgpt.html>.