



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 framework of a solicitors' ethical and professional obligations. For solicitors to be able to gauge and assess whether AI is being used within this framework, it is recommended they have a general understanding of how the technology works.

This guideline focusses on generative AI, and sets out potential issues and considerations 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:

1. Understand the tool you are using and know how it can assist your legal practice.
2. 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 commercial 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 client or matter intake process;
- summarising cases, reports and other documents;



- drafting correspondence;
- undertaking legal research;
- recording and transcribing minutes of meetings;
- drafting and reviewing contracts;
- preparing articles or presentations; and
- analysing case law to provide 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 fewer 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 might 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. This is known as hallucination, where the algorithm simply makes things up. While using a generative AI tool 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 does not understand who the client is. It cannot weigh up 'best interests,' optimise outcomes or provide explanations or reasoning. It lacks the skills of critical thinking and relevance. Humans are accountable, but generative AI is not. For this reason, it must undergo conscious critiquing, while contemplating both your 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 as 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. Many generative AI products aim to further train and improve their AI by learning from users' prompts. This could mean potentially exposing confidential client data to the service provider, which could be accessed by any of the service provider's customers. Organisations are also finding it harder to respond as quickly to potential AI vulnerabilities to hacking and data breaches, given the lack of reporting and detection capabilities that exist to monitor AI tools.

Good accountability and governance are key to maintaining your 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. Furthermore, the AI models currently in use cannot properly identify when the output they are returning is factual. 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.

In *Handa & Mallick*², a practitioner was afforded an opportunity make submissions to the Federal Circuit and Family Court as to why conduct relating to tendering a list consisting of fictitious authorities generated by an AI tool should not be referred to the Victorian Legal Services Board and Commissioner for investigation.

Furthermore, solicitors should not rely on generative AI to verify sources produced by AI. This has been known to fail. In *Mata v. Avianca, Inc.*³, a lawyer who appeared before the United States District Court for the Southern District of New York 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. Again, a number of the cases cited were fictitious.

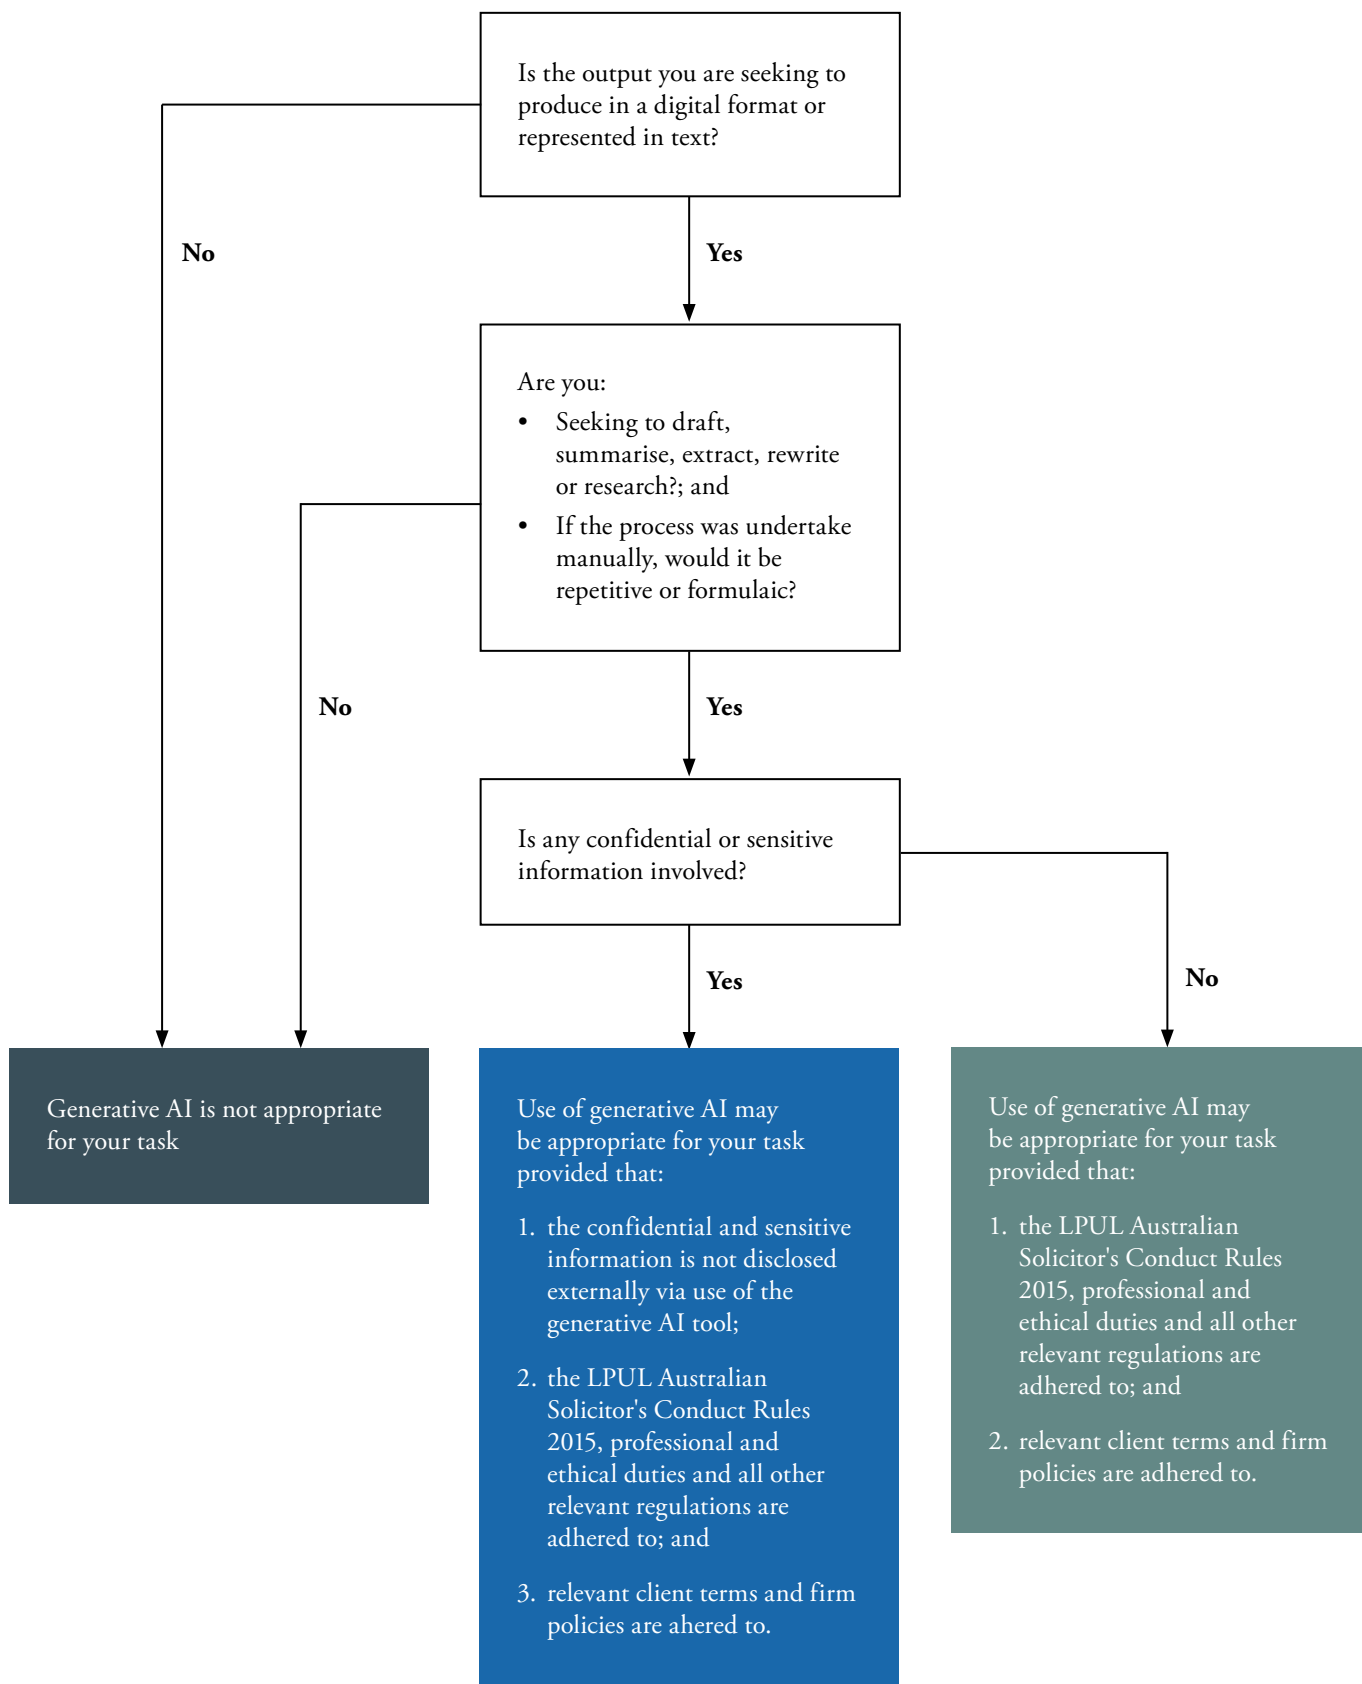
Rule 37 – Supervision of legal services

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.

Key considerations for procurement and use of generative AI in legal practice

Used appropriately, generative AI may be an effective tool to assist you to provide legal services. The below high-level infographic has been prepared in order to give an example of the kinds of questions that should be considered when determining whether the use of a generative AI product may be appropriate in the provision of legal services. A more detailed list of key considerations is provided following this.

Used appropriately, generative AI may be an effective tool to assist you to provide legal services.



The following considerations may help you decide whether your firm's use of generative AI in the delivery of legal services is appropriate and consistent with the Australian Solicitor Conduct Rules (ASCR) discussed above.

Pre-Adoption

- Consider establishing the following:
 - An appropriate generative AI use case, or specific scenario in which generative AI could be used to solve a particular problem;
 - a desired output for the generative AI tool; and
 - benchmarks to measure the generative AI tool's effectiveness.
- Evaluate your organisation's existing internal policies and risk management frameworks governing the use of generative AI and consider forming new ones.
- Consider undertaking robust research on vendors and tools available for the generative AI use case and evaluate integration with existing IT systems, tools and apps. It may be prudent to give preferential consideration to generative AI tools which present their output in a manner which enables the user to interrogate the sources or authorities upon which the output is generated.
- Prior to purchasing a generative AI tool, endeavour to undertake a full contract review to understand the key issues relating to copyright, intellectual property, data ownership, data sharing, and access to data provided to the tool. This includes details on how prompts, queries and data are stored to ensure that use of the tool would not breach confidentiality obligations.
- If the generative AI tool will have access to internal documents and systems, the principle of least privilege should ideally be followed, meaning the tool is only provided with the data it needs to operate and nothing more.
- Check the source(s) of data upon which the tool has been trained, noting that databases may contain misinformation or incomplete information.
- Check that use of the tool is compliant with relevant legislation, including in relation to legislation regarding the use of AI, cyber-security, data usage, and privacy.
- Consider the establishment of access control mechanisms to ensure that only authorised users can access and utilise the generative AI tool.
- Explore the provision of 'prompt engineering' training to the users of the generative AI tool.
- Consider undertaking thorough testing of the generative AI tool prior to its adoption. Outputs should ideally be expected, determinable, impartial, and fit for purpose.
- Ensure your client is aware and approves of the use of generative AI in the provision of their legal services and consider the implications using a generative AI tool would have on a client's potential requirement for onshore data hosting.
- Institute a process to ensure that practitioners with the appropriate skills and knowledge level are supervising and reviewing the generative AI outputs and their accuracy.

Post-Adoption

- Consider continual testing of the generative AI tool against benchmarks to ensure that its outputs remain useful and fit for purpose.
- If applicable, consider maintaining records of each generative AI output for discovery and retention purposes.
- If any generative AI output is intended for use in litigation, consider any relevant court protocols.
- Consider ongoing professional development opportunities to ensure that all practitioners have the skills and knowledge to:
 - Appropriately prompt the generative AI tool;
 - Evaluate the common limitations of the tool, including its susceptibility to AI bias; and
 - Supervise the outputs of tool and their accuracy.

Summary

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 with any legal assistant or paralegal. Solicitors must exercise independent forensic judgement, based on their own training, experience and research, and review, check, edit and correct any output to be confident it is reliable and correct. Solicitors need to demonstrate understanding, and may need to explain this understanding in verbal or written forms or provide additional supporting material not obtained via generative AI.

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 constant human oversight, at all levels of your business. Gaps in content need to be identified. In terms of privacy and data security, sensitive and client information should not be inputted into a public generative AI tool, particularly client information. Once in, it is hard to remove. Educating your employees on the risks and benefits of generative AI is essential.

-
- 1 Legal Profession Uniform Law Australian Solicitors' Conduct Rules 2015, rule 9.
 - 2 *Handa & Mallick* [2024] FedCFamC2F 957 (19 July 2024)
 - 3 *Mata v. Avianca, Inc.*, 1:22-cv-01461, (S.D.N.Y.)