How using AI won't save you money

The accessibility and ease of using AI has exploded in ways that most people could not have predicted only a short time ago. For many people, AI has become the tool in your pocket that turns large amounts of complex information into plain English. Rather than read a whole document, people can pass it to AI and in no time at all the document becomes a list of key points to focus on. It saves time, and because time is money, its saves money. Or does it?

When you feed information (like a legal contract) into an AI, and then ask the AI to tell you what risks or issues it finds in the contract, the AI will generally respond with a number of key points that appear to be relevant. When raising these points, the AI might provide you with references to different legislation that should be considered and on face value it will make a very convincing argument that there is definitely something that needs to be looked at further. To the average person it will seem to make sense, it will pass the pub test. The next step then is asking an actual lawyer to explain and address all these concerns that that have been raised by the AI.

Fast forward a little and a several pages long email is now in their lawyer's inbox. It contains a number of questions that have been generated by the AI that relate to the contract. All the questions appear important. Knowing that their client doesn't want to incur legal fees on trivial matters, the lawyer naturally assumes that every question that has been raised deserves deep consideration and attention to detail. Even though most of the questions don't actually make sense to the lawyer.

Why is it that they don't make sense to the lawyer, when they made sense to the client?

Because the lawyer is not only trying to answer the questions, they are also trying to figure out why the client:

- is referring to the section of an Act that is in NSW, when the business is in Queensland;
- is asking about the implications of specific licensing requirements, for services the client doesn't provide;
- is asking for a recommendation on what action to take, when the lawyer's recommendation was already provided in the previous letter of advice; and so on….

Identifying the difference between a genuine question and a question that simply isn't relevant can be time consuming. The lawyer assumes the client has a legitimate question that needs answering and therefore the lawyer needs to take the time to either answer the question or explain why it doesn't need to be answered. The more complex and granular the question, the longer it generally takes to explain why it isn't relevant.

The issue here is that the AI that generated the questions doesn't have all the context that the client and the lawyer take for granted. The AI's questions are based on the limited amount of information in the document the client has fed to the AI, mixed up with the extremely broad scope of information in the AI's data base drawn from all over the internet. Ultimately, this often leads to intelligent questions being asked that actually have very little relevance to the specific circumstances the client is in.

What could have been a 15 minute conversation with the lawyer becomes a several page letter of legal advice and a much larger bill for legal fees. Did the AI save money in this scenario? No. Could money have been saved if the client read the document their lawyer asked them to read? Yes.

Liability limited by a scheme approved under Professional Standards Legislation
Disclaimer – This article is provided for information purposes only and should not be regarded as legal advice.