

Episode 36 - Transcript

AI unplugged

Intro

The introduction of Artificial Intelligence tools like ChatGPT have generated new conversations and questions about their use in legal practice. Host, Julian Morrow chats with SJ—Schellie-Jayne Price, leader of AI Practice and partner at Stirling & Rose, about how these tools are used in legal practice, the risks of doing so, and how you can protect yourself and your practice against misuse.

Julian: Welcome to Risk on Air. I'm Julian Morrow, and today, we're chatting about ChatGPT and AI in the law. And to do that, I'm joined by a living, breathing human being and lawyer who's built a career around the dynamic intersection of law and emerging technology. SJ Price is a partner at Stirling & Rose, which I see from its website, is a law firm and emerging technology policy institute. SJ, welcome.

SJ: Thank you, Julian.

Julian: And thanks for joining us to talk. Everyone is chatting about ChatGPT and using the letters AI, but what really is AI? And especially, what is it for lawyers to think about?

SJ: So that's a great question because even the experts don't agree on what AI is. So, we all have our own little version of what artificial intelligence is. For me, it's when you have a machine that is doing something that, typically or previously, required human intelligence. So, for example, in the olden days, last millennium, I don't know if you remember this, but we used to use maps.

Julian: That's right.

SJ: Physical maps to actually try and find where to go. Now, we just use our phones. So that, for me, is artificial intelligence. ChatGPT is a very interesting type of artificial intelligence, and it is what is called machine learning. And what is really fascinating about machine learning is that you don't have a human writing the computer code that creates the system. What you have is you have what's called a learning algorithm. That learning algorithm identifies very deep patterns in that information and uses those patterns or statistical models to create the AI system. And that's exactly what ChatGPT is.

In fact, ChatGPT has been trained on so much data that Julian, if you and I were reading together, we were both reading for 24 hours a day, it would take us a thousand years, both of us, to actually read all the training data.

- Julian:** So if I type in a legal-ish question into ChatGPT, what's actually happening when I see that answer generated in front of me?
- SJ:** First of all, what's not happening is ChatGPT is not going off and searching for the answer. If you use Bing, that's a little bit different, but ChatGPT is actually predicting the next most likely word in response to your query. So, it's actually creating or generating new content in response to your query. It's just a statistical model that has learnt so much about the patterns in our language that it's able to predict really well.
- Julian:** So, in terms of legal practice, under what types of circumstances could something like ChatGPT be used, and I suppose in what ways is AI already being used in legal practice?
- SJ:** It's actually been used for a really long time, particularly in discovery. Artificial intelligence is great at going through lots of information and just pulling out things that are relevant. Now, I don't know if you remember this, Julian, but last millennium, when I was a junior lawyer, I remember doing discoveries where you would go into enormous rooms with boxes packed to the ceiling, and you would take a box down, take the lid off and physically have to go through the documents and read them. That's not the case now, it's all done digitally.
- So, lawyers have been using AI for a really long time. What is new, though, is the ability to use something like ChatGPT to actually generate what could look very much like legal advice rather than legal services.
- Julian:** Hmm—
- SJ:** And that creates a whole new category of risk.
- Julian:** What's at the top of your list for the risks that come with using generative AI in the context of legal practice?
- SJ:** Julian, I've got a top-four list.
- Julian:** Fantastic. Well, let's start with number one, shall we?
- SJ:** Okay, so this is the number one top risk: hallucinations. This is where ChatGPT creates very authentic and authoritative-sounding content that is subtly wrong, and sometimes it's not so subtly wrong.
- Julian:** We've all heard the case of the American lawyer, who relied in court, on a list of precedents that had been generated by ChatGPT that sounded very impressive; the only problem was that none of the cases exist.
- SJ:** Exactly, everyone's talking about the lawyer that used ChatGPT.
- Julian:** And we all live in fear that there's a little bit of Mr. Schwartz in all of us.
- SJ:** And I think there is, and that's why it's my number one risk because you have very clever people like lawyers who are seduced by automation bias. Sometimes, it's not subtly wrong, it's just wrong. For example, someone used it to ask about the record for walking across the English Channel and it proceeded to tell you who held the record. And someone else got it to generate reasoning as to why mayonnaise is racist. So, it is capable of generating data that's not even subtly wrong; it's just wrong. But it's the subtle nuances of our language that are very dangerous.

Lawyers really need to check very carefully the output that comes from ChatGPT in two ways; number one, check the content with reliable sources. But number two, really read the language carefully, because at a first reading, it can sound really good. But if you start reading it with a better and more detailed eye, you can actually identify some of the subtleties that are just simply not correct.

To give you an example, some of my friends are doing medical research, and they've asked me to use ChatGPT to generate abstracts. Now, the abstracts to me sound very good, right? But they actually say to me, "Well, it is useful, but there are subtle points I need to tweak to make it correct." That's why hallucinations is number one on my top four list.

Julian: Okay, so, hallucinations, we'll lock that in at number one. It's interesting, though, that the sort of strategies you described to minimise the risks of a hallucination getting through without being picked, they sounded quite old-fashioned and traditional. Checking sources and reading carefully, they're two things that lawyers should be pretty familiar with in other forms, and it's just a skill that we've got to apply in a new technological context.

SJ: Exactly, and you know what? We are familiar with it because for hundreds of years, lawyers have had juniors, and those juniors have done the first draft. And then their supervisors have gone through those drafts; that's a really good way to look at ChatGPT. Treat it absolutely with the scepticism, the healthy scepticism, that you review a graduate lawyer's work.

Julian: Maybe even treat it with a little bit more respect than the way you treat a graduate lawyer, dare I say. All right, so we got through number one, but there are three more big items on the risk list. What else is there?

SJ: So, my number two risk is ensuring that lawyers maintain the duty of confidentiality. And here, perhaps I should say we're talking here about the consumer version of ChatGPT. We're talking about the one where you've got it on your phone or your computer; you go to the web browser, and you log in. We're not talking about a situation where your law firm or your organisation has API-ed or whitelisted ChatGPT or another large language model within your organisation.

Confidentiality is really important because, of course, what we do when we use ChatGPT is we write a prompt, and in that prompt might tell you a lot about what's going on in your law firm and that's something we need to be very careful of. We need to ensure that we never upload in our prompts anything that's confidential, or sensitive, or proprietary, or any personal identifying information or any client confidential information.

Julian: Can typing information into a non-human piece of technology like ChatGPT, can you breach confidentiality to a computer?

SJ: That is a whole other podcast.

Julian: Yeah, yeah.

Episode 36 - Transcript

AI unplugged

- SJ:** But what we do know is we do know that OpenAI, who is the maker and producer of ChatGPT, actually has humans who look at these prompts. And one of the reasons, which is a completely rational reason, they look at the prompts to ascertain whether the users are complying with the service conditions and their guard rails and things that it shouldn't be used for. So, we do know that humans are looking at it. So, as long as you have a human looking at your prompt, there's a risk of a breach of confidentiality.
- Julian:** So regardless of what you might think of the way technology is working, it's always prudent to work on the assumption that because it's a third-party service provider, you are opening up information to a human somewhere.
- SJ:** That's exactly right. And even if you try to anonymise it, be very mindful that even anonymised data can actually pinpoint people. For example, if I upload something about a ninety-nine-year-old man living in Perth, that really narrows the field as to who that person is.
- Julian:** Yes, say hello to him for me as well, by the way.
- SJ:** I will, indeed.
- Julian:** So we've talked about confidentiality generally; what about the question of client-legal privilege in the context of ChatGPT or other generative AIs?
- SJ:** Julian, I'm so glad you asked that.
- SJ:** One of the things that I think is really important is to remind clients that if they upload your legal advice into ChatGPT, they may well be taken to have waived privilege. And some law firms, including ours, have actually updated our engagement letters just to highlight that to our clients, who might inadvertently not really think about the implications of uploading their legal advice into a tool such as this.
- Julian:** Really interesting because, obviously, that's an action that a client could take inadvertently. But then, there could be a professional obligation to, in the best interest of the client, warn them about that possibility.
- SJ:** Exactly. And to build on that, in litigation, and I'm not in litigation, but if I was, I would actually be asking clients, have you talked about your problem with ChatGPT? Because it occurs to me that uploading prompts, to the extent that they're discoverable or can be subpoenaed, might actually be evidence of admissions.
- Julian:** Gasps of breath taken in at the prospect, but better to be aware of it now. All right, what's next on the list?
- SJ:** So, number three is IP—Intellectual Property, and there's three aspects of this that are pretty interesting, and there's cases going on, on a lot of these. So, number one, you've probably seen the case against OpenAI, who is the maker of ChatGPT, where people are alleging that the training data that was used for ChatGPT, the training data that would take us thousands of years to read, was actually uploaded in breach of copyright. So that's your first issue, and there is an ongoing case with regard to that. The second one is, it is statistically possible but very unlikely that ChatGPT might create something that breaches copyright, that actually replicates the prior works of human authors. Very unlikely, but certainly, it's possible.

And then the third aspect is because you do very little in the production of the content from ChatGPT, arguably, you may not be able to claim copyright over those outputs. So, they're the three issues relating to IP.

Julian: All of which make sense, and again, very sensible to be aware of those when we've got those fingers hovering over the keyboard of ChatGPT.

SJ: And Julian, to just build on that about maybe we don't own the output of ChatGPT, I wonder if that's why the New South Wales Bar Association, in their guidelines for ChatGPT, have recommended that barristers keep records of their prompts and the outputs. Essentially, like a provenance log, and I wonder if that's in case they need to demonstrate that they have put enough human content or creativity or input to claim copyright.

Julian: That's really interesting. So I suppose that aspect is also keeping records as a lawyer of your use of assistive technology, which could come in handy for all sorts of reasons.

SJ: Exactly.

Julian: Okay, so we're three-quarters of the way through SJ's list of top risks to be aware of in AI. What's coming in at number four, SJ?

SJ: So, number four is a little bit of an unusual but super important one, and that is judgment. AI is not appropriate to be used in all circumstances. And an example of this was, there was a very tragic shooting in the United States at a university and a leader of an adjacent university created, using ChatGPT, what I would say was an incredibly empathetic note of sympathy and call to embrace diversity and so on, but people were outraged when they realised it had been written, not by a human.

But it's a very tricky one, but always think about, "Is it appropriate for me to be using AI in these circumstances?" And to build on the good judgment that's one aspect, but the other aspect is as lawyers, we have a duty of diligence and to do our work in a way that is most efficient for our clients. So, I think in the future, there is going to be a question asked about whether we should, in fact, be using these tools responsibly as part of discharging our obligation of due diligence and efficiency for our clients.

Julian: Because if it can be done more efficiently using the assistive technology, there's going to be a lower bill, and provided the appropriate judgment and oversight has been applied, then maybe that is a better result for the client.

SJ: Exactly, and I think it's important because what I'm seeing in legal practice is clients are using ChatGPT to generate term sheets, which they bring to you. They're using it to generate contracts, which they bring to you. They're using it to generate instructions to you. So, our clients are already using it, and I've even spoken to general counsel of organisations, asking them, well you know, if a lawyer is using this tool, do you want to know? Do you feel that they should be upfront and tell you?

SJ: And he laughed and said, "Well, actually, I want to know if they're not using it." So I thought that was a really interesting juxtaposition in his view that there is an expectation that lawyers use all appropriate tools in a responsible way to deliver their legal advice to the clients.

Episode 36 - Transcript

AI unplugged

Julian: And that brings us, I think, to the next question. We've talked about the things that individual lawyers should be aware of when engaging with the technology. What about firms? What measures should be put in place by a legal organisation to ensure that the individual users are employing SJ's top four tips, and I suppose to ensure that there are records and processes in place to be able to refer to and rely on if things don't quite go as well as we might hope?

SJ: That's exactly right, Julian. Two top tips: Education and Governance. And this is incredibly important because research has shown that 30-43% of people in organisations are using ChatGPT to do their work, to enhance their productivity. And of those, 68% are not telling anyone, so this is shadow use. So, my first tip would be: address it; tell your staff how you want them to use it or if you don't want to use it. I think that's really important. A lot of organisations have called a pause and have said we don't want anyone using this for work until we figure out how we're going to embrace it, how we're going to manage it and govern it.

The next thing is, there's tremendous value in educating people. Educating and consulting; so educating people in your law firms about what artificial intelligence is, what machine learning is, how it's used, how it can be used, and what it does and engaging in the discussion. There's lots of stakeholders; there's the lawyers in your firm, but there's also your clients and other staff that also have opportunities to improve productivity in a safe and responsible way.

The second thing is, once we have this education then it enables sensible decisions to be made about how we govern this technology. For example, coming out with some sort of philosophy about how you will use it or principles about how you will use it and a strategy, because there's tremendous opportunities for organisations that are prepared to embrace this, to be competitive, to deliver a more personalised service for their client. And then what we're seeing is a lot of organisations are actually having a policy about how their staff can use ChatGPT. And that policy is essentially, we encourage you to experiment and try this, but don't upload any confidential, personal, proprietary information. Use good judgment, and let us know if you're using this because we want to make sure that it is being used in an appropriate way.

Julian: I suppose if you ask ChatGPT to generate you an AI policy in your law firm, you should probably get that looked over, be wary of hallucination and all the other tips as well.

SJ: Exactly, Julian.

Julian: ChatGPT is one aspect of technology-related concerns that are really the focus of lawyers and others' attention at the moment. The other, of course, is cyber security; how do AI and cyber security fit into each other? Are they separate buckets, or do they need to be looked at together?

SJ: They need to be looked at in a consolidative fashion, and the reason I say that is if you're using artificial intelligence and you're uploading data, you want to make sure that data is safe. The second thing, and I don't know if you've noticed this, Julian, but I am getting really sophisticated phishing emails. It's really changing, and my view is that it's from scammers using ChatGPT. Because now I'm getting things that, if I don't look carefully, it looks like it comes from someone at my law firm, Stirling & Rose, and it's asking me to do very fun things like get my units for my employee share scheme and all sorts of really enticing opportunities.

Episode 36 - Transcript

AI unplugged

Julian: So there are many risks to attend to, but I suppose we should also stay tuned to the fact that there's immense potential here as well. Could you give us a couple of specific examples that, in your mind, point to some of the efficient, creative, or just interesting ways that AI technology can be applied in the field of law?

SJ: Absolutely! So, let me start with two really fun facts, Julian. The first fun fact is that GPT-4, which is the next iteration from ChatGPT, is able to pass the US Uniform Bar Exam in the ninetieth percentile.

Julian: Wow.

SJ: It can also pass the LSAT, which is the examination to get into law school in the eighty-eighth percentile, and this is important because you only need to pass it in the ninetieth percentile to be considered at Harvard Law School, so, GPT-4 almost going to Harvard Law School.

Now, more fun is some of the things we're using it for in real life, in the wild. Here's some examples: to help you get really cool and fun titles for articles or presentations.

Julian: An area where lawyers could potentially improve.

SJ: I'll give an example. We had to do a presentation, I and some of my legal colleagues, to an international audience. So, we wanted to give it an Australian flavour, right, to lure people in with the whole Aussie vibe. I used GPT, and what it came up with was "Sizzle or Fizzle: Cooking Up Innovation Down Under."

Julian: Not bad!

SJ: So that was great, right? And we had this whole barbecue thing happening, so that's one example. Another one is style transfer, if you have to write an email or a message to someone and you want it to have a particular tone, for example, if you have someone in your law firm maybe who isn't keeping up with their time sheets and what you really want to say is, "Look dude, you need to keep up with your timesheets! Right? Well, ChatGPT can actually make that sound very polite and very nice, and you know, "We understand you're very busy, but this is really important and please can you..." and makes it a very nice style.

For example, you can learn things. So if I had a client whose domain specialty was quantum computing, I could ask ChatGPT to explain to me well what is quantum computing. And it might say something, and I say, oh, that's too complicated for me; make it simpler for me. So, you can use it as a personal tutor to enable you to understand new domain areas. Julian, one of the other things that ChatGPT is incredibly useful for is to check your spelling and your grammar.

Julian: Oh, good!

SJ: So there are no excuses anymore for having bad spelling and bad grammar. And this is particularly important if English is your second language; you can use this as a tremendous English graduate to check your work.

ChatGPT is also great to summarise notes. For example, when we get Copilot, Microsoft Copilot, it will actually take a transcript of a meeting you're having, say, on Teams, and then it can summarise the meeting and identify the action items, who has to do them, and by what time. This is an incredible productivity boost.

Episode 36 - Transcript

AI unplugged

Now of course, the one I hear lots of lawyers using it for is what we've discussed, which is to do your first draft, your first draft of a particular clause, your first draft of a contract. But, the uses are only limited by the imagination of lawyers who are incredibly clever people.

Julian: So watch this space. I suppose it's worth bearing in mind also that ChatGPT is one large language model and generative AI. There are others out there as well, and I imagine that we'll see AI integrating into all sorts of apps and technologies that we're already using. It might become part of our basic word processing software, it already has, I suppose. So there really are going to be lots of both general work and law-specific options becoming available and worthy of exploration.

SJ: That's exactly right, Julian, and you know what? It's closer than we think because those of us who are in the Microsoft O365 environment...if you're in the United States, it's there now.

But shortly, watch this space; it's coming to Australia. It's called Microsoft O365 Copilot, and what it is, it's GPT-4 within your desktop environment. And what this enables you to do is to generate content by pointing it to the documents within your organisation that you have access to. That means that you can say, "Look at this Word document and create me a PowerPoint. Or look at this Word document and put that information in a tabular form for me."

Julian: Interesting stuff! It sounds like the sort of thing that, in the past, you might have been getting junior lawyers to do. And I'm sure that's not going to lead to junior lawyer redundancy. It just means that the humans are going to be doing more high-level and more interesting stuff, and the computers are going to take the grunt out of it.

SJ: Exactly. In fact, it's a great time to be a lawyer because you're going to bypass all those boring, boring tasks that lawyers like us had to do when we first joined the workforce.

Julian: Well, SJ, it's been great chatting ChatGPT with you, and lots to think about, but also lots of benefits along the way as well. So, I suppose it's good luck in the wild west of new technology.

SJ: Thank you, Julian.

Outro

Thanks for listening to Risk on Air by Lawcover. Join us for the next episode on current risks in legal practice to stay up to date.