# Artificial Intelligence and Expert Evidence

## Opportunities, Risks and Responsibilities

**Joanna Smith J**

Good morning and thank you for inviting me to take part in the Bond Solon Expert Witness Conference. It is a real privilege to be here.

Some of you may know that the role of expert witnesses in our court system is close to my heart, no doubt in part because of my background at the Bar. As a barrister practising in professional negligence and commercial litigation, where the expert evidence is often crucial to the outcome of the case, I spent many happy hours with experts in many different fields – from accountants to engineers, patent agents to (would you believe) international gun experts (the latter giving a whole new meaning to the term "hired gun")!

As a judge sitting in the BPD I now spend a considerable amount of time listening to experts being cross examined and making findings about the credibility or otherwise of their evidence.

And the credibility of the expert, which so often turns on his or her compliance with the fundamental duties to the court, is really what lies at the heart of my speech today.

The advent and rapid development of AI has not changed, and will not change, these duties. The obligations of independence, objectivity, transparency and professional competence remain exactly the same as they were before the emergence of AI.

The question is therefore not whether experts should use AI. Nor is it whether AI should be banned from the expert process. The question is how experts can use these tools responsibly, safely and in a manner consistent with their obligations to the court.

I want to begin by running a straw poll to see how many of you at this conference have used AI in your practices in connection with the preparation of expert evidence and whether that use has generated or influenced the substantive content of your reports.

Against that background, my aim this morning is to try to offer a practical framework for your ongoing work as expert witnesses. I shall begin by discussing the recommendations of The Civil Justice Council in its report published last week on the "Use of AI for Preparing Court Documents". The Civil Justice Council, as many of you will know, is an advisory public body that reviews civil justice and makes recommendations designed to improve the system. I was co-chair of the working group that produced that report. I shall then identify sensible uses of AI within the expert process, highlight briefly some significant risks, and conclude with some suggested practical safeguards that experts may wish to adopt.

**The CJC's Recommendations**

The courts recognise that AI technologies can improve efficiency and productivity. They can help users to process large volumes of information, identify patterns and communicate complex material more effectively.

Equally, the risks are now well understood. Generative AI systems are capable of producing entirely false information presented with complete confidence. They may invent sources, misunderstand context, misstate facts and provide persuasive but incorrect analyses. There are already examples in other jurisdictions of experts including AI-produced fictitious content in their reports.

The Civil Justice Council's report attempts to reflect a balanced approach to these opportunities and risks.

The report acknowledges that increasing numbers of expert witnesses are using AI in their role as experts, something that is borne out by the results of the 2026 Bond Solon Expert Witness survey published, fortuitously, today and the straw poll we have just run.

When the Working Group went out to consultation we wanted to understand whether experts considered that this use should be transparent. Experts must already verify their reports with a statement of truth designed to support compliance with their duties, but was it necessary we asked, for them to declare AI use?

The answer we received from those responding to the consultations was overwhelmingly Yes.

This chimed with the views of the Working Group and also with guidance issued earlier this year by the Academy of Experts, with which many of you will be

familiar. Thus we concluded that the integrity of expert evidence, including the need for proper scrutiny of the expert's methodology and the maintenance of a level playing field between experts, requires a degree of transparency around the use of AI which can only be achieved by the introduction of a new AI use specific declaration. We have recommended a rule change to that effect.

Importantly, that declaration would only be required where an expert has used AI tools to generate or influence the substantive content of his or her report – this will include where the AI use formed part of an expert's methodology or substantive analysis. AI is already in widespread use for so many tasks, and embedded in so many existing tools that might be termed administrative, that we did not consider there to be any need for a declaration in such cases. We have concentrated instead on the integrity of the expert output, together with the need for human accountability and verification. Where a declaration is required we believe that it will need to provide sufficient information to facilitate effective scrutiny of the expert evidence. We envisage that this may include disclosure of the methodology used together with confirmation as to the appropriateness of AI use for that purpose and, exceptionally, but not always, the tool that has been used.

At the same time, and recognising the need for guidance in this area, which is also highlighted in the Bond Solon Expert Witness Survey, we have recommended that revisions should be made to the CJC's Guidance for the Instruction of Experts in Civil Claims (2014) to introduce good practice guidance on AI use by experts into revised Guidance.

We have indicated our view that this guidance needs to facilitate consideration and application of a three-part test: first experts should consider whether their use of AI has generated or influenced the substance of their evidence; second (and if a declaration needs to be made) they should consider what information should be disclosed to enable the opposing party and the court to scrutinise their methodology and analyse it effectively; and third they should consider whether it is necessary in the particular circumstances to disclose the AI tool used.

**Where AI May Be Used Safely in the Expert Process without the need for a declaration (assuming our recommendations are put into practice)**

AI may assist with a whole host of administrative tasks, information management, and drafting support which I doubt will engage the need to provide a declaration.

Many of you will already be using AI for organisational tasks which help to reduce the burden of dealing with large volumes of documentation and complex data. Some of you may not even know that you are doing so – some data analytics tools are already using AI in ways which are not immediately apparent.

You will be using it to organise large document collections, create chronologies, classify documents. You may also be using it for summarising lengthy documents and identifying themes, for carrying out research, for structuring your reports and for checking your final reports for errors or inconsistencies, for spelling and grammar checks. It can be a brilliant "second pair of eyes".

While, as I say, these things may very well not require a declaration because they do not affect the substance, you do need to exercise considerable caution.

An AI summary of a document may not pick up a crucial feature that would have jumped out at you; it may miss important matters of context or nuance. In this way, if used in a report without checking the underlying document thoroughly, it may influence the substantive content. An administrative tool that helps to structure user content may influence (albeit perhaps only in a subtle way) the substantive nature of the content. AI transcription may create substantive content that does not accurately reflect what was said. This is why at all times your focus as an expert should be on understanding what the AI tool that you are using is capable of.

What I am really saying is that, where appropriate, you can of course make use of AI as a tool for research and as a tool to achieve greater efficiency and speed, but if it is used in any way that may directly or indirectly generate or influence the content of your evidence its use will have to be disclosed (assuming of course that the CJC's recommendations are accepted by the rules committee) and you will have to provide sufficient information to the other side and to the court for them to be satisfied that your report is nevertheless genuinely a product of your own opinions.

If proper disclosure is not provided you can expect your report to be the subject of searching questions in cross-examination during which you will have to justify your use of the AI tool and explain why it has not in any way influenced your expert opinion.

**Pitfalls and Risks: What Can Go Wrong?**

I am going to touch on this now only briefly. Most of you will, I am sure be aware that the key risks of AI use include hallucinations, confidentiality breaches and technical misunderstandings. At a more philosophical level, the key risk from the use of AI by experts is in the erosion of independent judgment, the outsourcing of your opinions to AI and thus the loss of expert integrity.

Returning to a more practical level, significant risks, highlighted by some of the reported cases of which many of you will be aware, are the dangers of a lack of training in the use of the relevant AI tool, a lack of understanding as to what it is capable of and a failure to ensure adequate levels of oversight. If you work in organisations which use teams to analyse material and carry out research, you will need to ensure that you understand what AI tools these teams are using and how they have been used (even if you have not used them yourself). You will also need to ensure that there is a strict protocol in place to ensure that everyone in the team understands the importance of you, as the expert, being clear about what use they have made of AI.

At risk of repeating myself, the way to guard against these risks of course has its foundation in the fundamental duties of an expert. The need to verify all factual assertions, references and conclusions. To take nothing at face value. To provide your own independent opinion in your own words

**Practical Safeguards and Good Practice**

I do not wish to be prescriptive and nor will the things I am going to suggest be relevant in every case, but I can identify a number of important questions that every expert should bear in mind upon receipt of instructions (and I am afraid there are rather a lot!):

1. What are my instructions on the subject of AI? May I use AI or is it prohibited for some reason (such as intellectual property concerns, data protection or regulatory compliance). Even if your instructions do not raise the issue, I think it makes good practical sense for you to disclose your intended use in the first instance and to discuss that use with your instructing solicitors.
2. Does the technology that I intend to use incorporate AI? Do I need to carry out due diligence to investigate that question?

3. Is the AI tool I intend to use fit for the purpose for which I intend to use it? Why has it been chosen over any other similar tool? Am I using the most up to date tool for this task? It might be necessary to carry out research into other similar tools that might be more suited to the purpose. One can readily imagine cross-examination around the question of whether you have used the best, most up to date tool for the task and, if not, why not.
4. Is the tool I intend to use secure such that confidential or privileged information is not going to find its way into the public domain – a topic I suspect it will always be worth discussing with your instructing solicitors.
5. Where you are working in a team – do all members of the team understand the importance of disclosing their use of AI and of keeping a written record of the key steps in the process, including how the output was used? Have they all had appropriate training?
6. Whether you are working in a team or not: Do I understand how the AI system operates, including its limitations and the steps that have been taken by others in my team to validate its outputs? Could I explain these things in court if I had to do so?
7. Do I have a proper record of how the AI tool has been used and the extent to which, if at all, it has generated or influenced substantive content?
8. If it has generated or influenced substantive content, then I think there are 10 things you need to be asking yourselves:
    a. Do I need to disclose the AI tool that has been used and can I justify its use?
    b. Is the output robust and can it be reproduced?
    c. Might it be subject to any known biases in the AI tool or underlying data set on which the tool was trained and, if so, how might those impact the output?
    d. Might it include hidden assumptions?
    e. Have I done all that I can to verify the output?
    f. Have I applied sufficient human oversight to ensure that any hallucinations or errors have not found their way into the evidence?
    g. Are the opinions expressed genuinely my own?
    h. Can I explain every conclusion that I have arrived at?
    i. Would I be comfortable describing exactly how I used AI to a judge under oath?
    j. Coming full circle, have I complied with my professional duties?

**Conclusion**

AI is a tool. Much as we may be prone to anthropomorphise it, AI is not an expert and should not be a substitute for expert opinion. Responsibility will always remain with the expert witness. Ultimately, as I hope I have shown, transparency, verification and professional judgment are essential safeguards and should be your guiding principles.

As it seems to me, the future belongs to experts who understand both the capabilities and the limitations of AI, who keep abreast of new technologies and who use AI thoughtfully, appropriately and (where necessary) transparently in support of their overriding duty to the court.

Thank you.

JS

5.11.26