

AI in adjudication: a benefit or a burden?

AUGUST 11, 2026



I recently had the pleasure of being one of the trainers at the [TECBAR Adjudication Training Day](#), organised by Mercy Milgo and chaired by Calum Lamont KC. Alongside, Mr Justice Constable, Alexander Nissen KC and Martin Bowdery KC, we covered “*Adjudication from Cradle to (Almost) Grave*” and, for my sessions, I was very generously trusted with training delegates on “*How to Run the Perfect Adjudication*”. I have to admit that “perfect” isn’t a word many would instinctively associate with adjudication. In fact, I doubt there is even such a thing as a “perfect adjudication”, let alone a single, one-size-fits-all way of conducting one. Ask ten experienced adjudicators how they approach the task and you will almost certainly receive ten slightly different answers. Nevertheless, having spent the best part of two decades acting as an adjudicator, I was very happy to share some of the practical lessons I have learned along the way, together with some of the issues and frustrations that regularly arise throughout the lifecycle of an adjudication, from appointment, jurisdictional challenges, submissions, requests for more time, etc., to, ultimately, producing a robust, “bulletproof” decision.

These days, it’s impossible to discuss any form of dispute resolution without mentioning AI and much of the conversation tends to focus on familiar concerns. “Hallucinations” seem to have become something of a catch-all expression, and certainly make for good headlines, so I won’t focus on those too much as we’ve all heard the horror stories, and perhaps encountered for ourselves case citations that simply don’t exist. While those risks are real, I suspect (and certainly hope) they are now reasonably well understood across much of the legal profession.

One of the issues we discussed during the training session was how adjudicators should deal with submissions that have been prepared, or at least substantially assisted, by AI. Closely related to that is perhaps the more difficult question: to what extent, if at all, should adjudicators themselves be using AI? With senior members of the judiciary continuing to advocate its responsible use, those questions certainly keep giving us all plenty of food for thought.

The AI ship has sailed

AI-assisted submissions are increasingly part of everyday practice, whether we like it or not. I’m sure many readers will have heard or read Dame Victoria Sharp’s recent lecture, [Without Fear or Favour: Judicial Independence, Past, Present and Future](#). During the lecture, she acknowledged the considerable benefits that AI can bring to access to justice, including helping people to understand

their rights, draft documents and navigate unfamiliar procedures. The same is surely true in construction disputes. If AI enables a small contractor, an unrepresented party or a business without access to specialist construction lawyers to present its case more clearly, we should not dismiss that benefit too quickly.

There is, however, another side to this. Dame Victoria referred to the *“boom in lengthy applications”*, observing that each requires careful judicial appraisal and that many are less than legally sound. She also highlighted the now familiar problem of false legal material, including invented authorities and quotations. Although her subject was judicial independence, a number of the issues she identified have obvious parallels for those of us adjudicating and arbitrating.

AI is extremely good at producing language, and most adjudicators reading this will already have encountered submissions running to many pages, carefully rehearsing and often repeating every aspect of the factual background multiple times. Reading those submissions can be time consuming, but parties have always been free to present their cases as they see fit and it has always been part of the adjudicator’s job to separate what matters from what doesn’t. What AI is far less good at is exercising judgment: distinguishing the central issues from the peripheral ones and identifying the points that will actually determine the dispute. I recently dealt with an adjudication where the Response narrative ran to circa 20 pages and the Rejoinder to over 250 pages. Much of that material was repetitive and some really important issues received little or no attention. There is perhaps some irony in technology intended to improve efficiency instead creating considerably more work for the adjudicator.

What about adjudicators using AI?

So far, I’ve been talking about parties using AI, but what about adjudicators? I have already mentioned Dame Victoria Sharp’s recent lecture and a couple of interesting points she made particularly resonated with me where she also expressed some concerns about the dangers of AI to judicial independence. She said: *“A judge who treats AI as a source of “judicial reasoning” rather than a useful tool risks outsourcing part of the judicial function. Summarising material, translating text, locating authorities already known to exist, or managing large quantities of documents may be legitimate uses, provided confidentiality, accuracy and responsibility are preserved. But there is a qualitative difference between assistance and influence. The danger is not merely that an AI system will make a crude mistake. It is that it will produce an attractive formulation which subtly narrows the judge’s own analysis.”*

Again, I think that this distinction applies just as much to adjudicators and arbitrators. Personally, I would not use AI to summarise the parties’ submissions, not to mention to any extent decide the issues. One of the most valuable parts of writing a decision for me is producing my own summary of the parties’ cases and the exercise forces me to identify the issues that really matter, to understand where the parties agree and disagree, and quite often to realise that what initially appeared to be the central issue is not one that will ultimately determine the dispute. Every adjudicator understands that the decision-making process involves far more than producing the final written decision – you have to read submissions, identify what matters, weigh competing evidence and work through difficult legal and factual questions as part of reaching an independent conclusion. If somebody else prepared that summary for me (be it a colleague, a pupil or an AI system), I can’t help but feel that I’d lose something valuable. Others may take a different view, but for me, that exercise is all part and parcel of the decision-making process itself and not an administrative task.

Dame Victoria also highlighted another reason for caution: “automation bias”, that is our natural tendency to place confidence in systems that *“appear technical, neutral and authoritative”*. So the danger is not necessarily that AI produces an obviously wrong answer. It is that it produces a plausible one, expressed with such confidence that it subtly influences the decision-maker’s own thinking before they have fully analysed the evidence for themselves. That strikes me as perhaps a

far greater concern than the occasional imaginary case.

None of this is to say that AI has no use for an adjudicator. Back in April this year, the City of London Law Society hosted a panel on legal professional privilege in the age of AI. In his keynote speech, Sir Colin Birss, Chancellor of the High Court, said that he was using AI to identify internal inconsistencies in his own work. He said:

"Once I have finished writing a judgment, I give it to the secure Copilot system on my computer and ask it to identify any internal inconsistencies. It is remarkably effective. What I choose to do with the proposals is up to me. I don't always agree with the AI but it has been helpful and I have clarified wording in draft judgments as a result."

Recently I used Microsoft Co-pilot with one of my own decisions. By that stage the decision had already been drafted, checked by me, amended, proofread by my colleague, Helena, and amended again. I wasn't asking AI to check whether my findings were right or wrong, just to look for things that are easy for any human to miss, so we're talking internal inconsistencies, grammatical slips and arithmetical errors etc. Interestingly, even after two rounds of proofreading, it identified genuine internal inconsistencies that neither Helena nor I had spotted, for example incorrect references to years and so on. Equally, it made plenty of suggestions that I simply ignored. The important point is that it is acting as nothing more than a final proof-reader and not influencing the reasoning or the conclusions.

Knowing where to draw the line

Adjudication has many strengths, and there is plenty of best practice to draw on (not sure if you've read mine and Matt's [book](#) ) , but by its very nature it is never going to be perfect. Modern dispute resolution routinely involves thousands of pages of electronic documents, witness statements, expert reports, contracts, payment notices and authorities, all of which, certainly in adjudication, have to be considered within an exceptionally short timetable.

I don't think any of us can deny that AI offers some opportunities to improve efficiency across the whole lifecycle of an adjudication, particularly when it comes to genuinely administrative tasks. For parties, AI should be a drafting assistant rather than a substitute for legal analysis. For adjudicators, and dispute resolvers more generally, it should remain an administrative tool rather than a decision-making partner. The challenge is knowing where assistance ends and influence begins, but perhaps getting that balance right will take us as close as we are ever likely to come to the perfect adjudication!



JONATHAN COPE