

The Arbitr-Al-tor: Judgment Day for Arbitrators?

SEPTEMBER 30, 2025



Whatever your age, background or job title, Artificial Intelligence (AI) is pretty much impossible to ignore. It dominates news headlines, slips into everyday conversations both at home and at work, and often secures a spot on conference agendas. It sometimes feels that nowhere does the noise feel louder than when it comes to dispute resolution, where speculation over its potential and pitfalls continues to reverberate. From breathless predictions of robots replacing lawyers to dire warnings about machines taking over the justice system, it's sometimes been tricky to separate the reality from the rhetoric, so I've held off blogging about it until now. But the recent announcement by the American Arbitration Association (AAA) / International Centre for Dispute Resolution (ICDR) about their plans to launch, this November, an AI arbitrator capable of drafting awards for "documents-only" construction disputes caught my attention and so it felt like a good time to take a closer look.

The tool, being developed in collaboration with QuantumBlack (McKinsey's Al arm), is built on a dataset of over 1,500 construction awards, annotated and refined with human-arbitrator input and promises to deliver fast, cost-effective, and trusted dispute resolution. According to the press release, it says: "The Al arbitrator effectively tech-enables the approach to alternative dispute resolution, accelerating case management and legal reasoning and operating within a "human-in-the-loop" framework that validates every output. With step-by-step oversight by an arbitrator, the Al tool can review filings and supporting documents, break down claims into their component arguments, and generate draft awards grounded in decades of case data and experience." All sounds promising, right?

All in a day's work

I've been acting as arbitrator almost as long as I've been adjudicating and have spent more hours than I can count drafting awards relating to a broad range of pretty meaty construction and engineering disputes. An arbitral award is not just a list of findings, it's a carefully reasoned document that explains the tribunal's findings on facts and law, addresses the parties' arguments, and complies with enforceability standards. Drafting requires clarity, neutrality, and precision and, as I'm sure many of you can attest, it's a painstakingly time-consuming process. At the more complex end of the disputes spectrum, the reasoning can stretch into hundreds of paragraphs, and, as arbitrators, we can sometimes spend days fine-tuning language to get the balance of accuracy and persuasiveness just right. Let's be honest, anything that can take some of the tedium out of certain of the more mundane elements of the award writing process and boost efficiency might be a welcome addition.

So, with the first (I believe) institutional body about to formally deploy an AI tool to draft arbitral awards, are we are on the brink of 'award-winning' progress or 'award-worrying' outcomes?

The downsides

While the arrival of an Al arbitrator may promise efficiency and innovation, it also brings risks that still can't be overlooked.

Accuracy and reliability - A lot has been written about so called "hallucinations" in AI and I think accuracy and reliability are the biggest issue. AI systems can misinterpret evidence, apply legal principles incorrectly, or oversimplify complex factual disputes. There were a couple of high-profile cases earlier this year which brought AI misuse into sharp focus.

The combined cases of Ayinde v London Borough of Haringey and Al-Haroun v Qatar National Bank [2025] EWHC 1383 (Admin) involved the use or suspected use of generative Al resulting in fictitious caselaw, fake citations, and misstatements of law being relied upon by parties to litigation. In the Ayinde case, counsel put forward five non-existent cases as evidence. Counsel denied using generative Al, claiming instead that the authorities came from general internet searches, yet when pressed, she was unable to point to any online source for the fictitious cases.

In the second case, Al-Haroun v Qatar National, the claimant's solicitor submitted a witness statement containing 45 authorities, of which 18 turned out to be fictitious, while some others weren't accurately applicable to the arguments or were misquoted. The solicitor had relied on research undertaken by his client, without independently verifying it.

In response to the cases came a stern warning from Dame Victoria Sharp, President of the King's Bench Division of the High Court, who said that while AI is powerful, the likes of ChatGPT 'are not capable of conducting reliable legal research' and that there are 'serious implications for the administration of justice and public confidence in the justice system if artificial intelligence is misused'. In both cases, the Court declined to initiate contempt proceedings, instead referring the matter to the Bar Standards Board and Solicitors Regulatory Authority respectively for investigation.

Decided just a week before the hearing in these cases was *Bandla v Solicitors Regulation Authority* [2025] EWHC 1167 (Admin), where the Court encountered the same issue. The judge said at paragraph 53: "In my judgment, the Court needs to take decisive action to protect the integrity of its processes against any citation of fake authority". So, there's a clear trend emerging here, and it seems inevitable that regulators will take a dim view of any proven reliance on bogus case law especially given the concerns of the senior judiciary and the growing frequency of these incidents.

Confidentiality is another crucial factor. One of arbitration's core attractions is its privacy, with parties expecting their commercially sensitive information to stay protected. While this concern isn't unique to AI, uploading case materials into AI systems, especially those hosted on external servers, raises obvious risks of data breaches or unauthorised use.

Human skills and judgement - Arbitration is valued, not just for technical accuracy, but also for the ability of arbitrators to weigh up credibility, context, and fairness. Al could rob decision makers of some of these skills, and I think over-reliance on Al risks reducing awards to formulaic reasoning that could miss the nuances of human decision-making.

Enforceability poses another challenge. Some commentators have suggested that national courts could question whether an award partially drafted by AI satisfies the legal requirement for a reasoned decision. This could provide fertile ground for parties to resist enforcement, undermining arbitration's effectiveness as a final and binding process.

Underlying data for me is another area of risk. Because these AI systems are trained on existing

arbitral awards, they inevitably reflect the reasoning patterns and tendencies and the flaws of that material. If past awards contain certain limitations or biases, there's a real risk that AI tools will replicate or even amplify those patterns.

Transparency and legitimacy. If parties are not clearly informed about the extent of Al involvement in drafting an award, they may feel that the process lacks fairness. This perception could damage trust in arbitration as a party-driven and impartial method of dispute resolution. I've certainly come across submissions that had the hallmarks of Al assistance as they were particularly heavy on verbiage, but light on the substance that actually addressed the facts. Tools are in development (and some already exist) to detect Al-generated content, flag hallucinations, and verify authorities, which might help tribunals and opposing parties challenge suspicious filings.

Over-prioritising efficiency - Institutions and arbitrators may be tempted to lean on AI to speed up proceedings and reduce costs. However, while I'm all for working smarter and not harder, efficiency should not come at the expense of careful reasoning or case-specific deliberation. Awards that feel a bit "off the shelf" or generic could again erode confidence in the arbitral process.

How AI can assist arbitrators

Whether AI can draft truly sound and enforceable awards remains to be seen. What seems more realistic, at least for now, is its ability to take on the 'bronze-tier' work, so the tedious, repetitive sections that slow the process down. Rather than producing the full award, AI tools could be used to generate first drafts of standard parts such as background facts, procedural history, and summaries of submissions, leaving arbitrators free to focus on analysis, reasoning, and the ultimate decision. However, I often find summarising submissions helps in my understanding of a party's case, and so I would question whether this will really assist arbitrators.

Any tool which could review the draft award to detect internal inconsistencies or contradictions (which aren't uncommon risks in lengthy awards), could be enormously beneficial or which would help improve clarity and standardise terminology especially as many awards are written in a second language for at least one of the parties in international arbitration.

These types of considerations were the topic of a recent academic study entitled "Interacting with Al at Work: Perceptions and Opportunities from the UK Judiciary". It was based on a focus group of 12 UK judges from a range of civil, family and criminal courts. Whilst not necessarily representative of the judiciary as a whole, it does provide some insight into the issues and benefits that may be being looked at. I'll let you read the article for yourselves but, in high level terms, some of the potential benefits included: (i) increased consistency, efficiency, access to justice; (ii) improving information; and (iii) reducing bias, cost, and tedious work.

Obviously, such an AI model could equally apply to adjudication if adjudicators' decisions were made publicly available in a structured dataset. Adjudication is designed to be quick and cost-effective, but the same issues of consistency, efficiency, and clarity arise there too. It isn't hard to imagine an AI tool supporting adjudicators in producing draft decisions, detecting inconsistencies, or standardising terminology, all while leaving the substantive decision making firmly in human hands.

Hasta la vista, Arbitrator?

As dispute resolvers, it would be entirely remiss of any of us to ignore AI. The AI-generated horse has well and truly bolted, and we have to decide how and when we want to ride with it or rein it in (pun intended). The use of AI to write awards, in particular, is both promising and controversial. It offers very real opportunities to improve efficiency, consistency, and clarity in arbitration, but while AI can certainly draft like us, and sound like us, can it truly think or decide like us?

None of the reported cases on Al misuse to date have succeeded because someone trusted the

technology blindly and it would seem that, in every instance, human diligence was still required. The human oversight model is clearly crucial and AAA-ICDR is emphasising that human arbitrators remain ultimately responsible for the award, but whether courts will accept those generated (even partially) by AI will certainly be one to watch. Accountability for me remains an open question. If an AI-assisted award were ever found to contain flawed reasoning or factual errors, it could raise complex questions of responsibility: would it rest with the arbitrator who signed it? the institution that deployed the tool? or the developers who built the system? What will also be interesting to see is how the success of this pilot could influence whether other institutions or jurisdictions adopt similar AI-assisted tools. Though, until some of these legal, ethical and professional boundaries are clarified, I can't help but feel the use of AI in award writing will remain controversial for many.

The proposed AAA-ICDR AI arbitrator is, for now at least, confined to 'documents-only' construction disputes, so cases decided entirely on written submissions and evidence, without live hearings. These disputes are usually straightforward, relatively lower-value matters, which makes them a sensible testing ground, but they're a far cry from the complex, high-value disputes for which international arbitration is best known. Cost, of course, is another consideration. Developing and maintaining AI platforms isn't cheap, and those expenses will inevitably filter down to users. While automation might reduce arbitrators' drafting time, it remains to be seen whether those savings will outweigh the licensing fees on these types of smaller disputes, so parties may well end up paying more for the technology than they save on tribunal time.

All that said, Al undoubtedly has the potential to support arbitrators in meaningful ways and as an industry we must continue to explore and develop these ideas. Just as word processors revolutionised legal drafting decades ago, Al tools could become trusted digital assistants provided the right safeguards are in place and ensuring that efficiency gains don't come at the cost of fairness or enforceability. Robots may be able to help draft, but they can't yet replace the human judgment that gives an award its legitimacy. Not yet anyway. As for the moment, we arbitrators live to fight another day and responsibility will (and must) remain with human decision-makers to ensure awards are reasoned, impartial, and enforceable. So, the good news is, as Arnold Schwarzenegger famously put it: 'I'll be back.'



JONATHAN COPE