

Decision-making: is it as easy as AI, B, C?

APRIL 29, 2026



Six months is a long time in AI. What feels like a breakthrough one month can feel standard or outdated the next, as new models, technologies and capabilities roll out at speed. From drafting contracts, case management to predicting case outcomes, AI is steadily expanding into areas that have historically always been considered firmly and encouragingly human.

There is, of course, no universally agreed legal or technical definition of AI. The UK Jurisdiction Taskforce, in its draft Legal Statement on Liability for AI Harms (January 2026), adopts a technology-agnostic definition: AI is autonomous technology. In that sense, autonomy brings three familiar features: an unpredictable relationship between input and output, opacity in the reasoning that produces the output; and limited user control over the output once the system is deployed.

So, for all the buzz around AI entering the legal profession, the everyday uptake still feels understandably quieter, slower and more selective, well, for now anyway. Six months ago, the American Arbitration Association rolled out its AI Arbitrator which, as Jonathan [blogged](#) about at the time, is for documents-only construction disputes. Though, as of early 2026, reports suggest only a single case has actually gone into the system which I think alone indicates it's still finding its feet and not reshaping arbitration overnight. Even the term "AI arbitrator" can perhaps be taken a little too far, because the system doesn't operate independently; a human arbitrator remains firmly in control, reviewing and verifying the output and issuing the final Award. So, while admittedly not as catchy, it's perhaps more a case, in reality, of augmenting legal judgment rather than replacing the role itself.

Lawyers and arbitrators have been actively refining the system since its launch, and the AAA has begun rolling out related early evaluation tools like a ["Resolution Simulator"](#) to help legal teams test strategy and inform negotiations by predicting outcomes before proceedings even begin. This feels to me like the system is growing in support and experimental contexts, not in widespread use just yet.

This relatively slow uptake perhaps isn't entirely surprising. Legal systems are built on trust, precedent and careful judgment and introducing AI into that mix raises immediate questions about fairness, confidentiality, verification and accountability. These were just some of the themes touched on at last night's Worshipful Company of Arbitrators event, "Law, Ethics and AI: The Risks and Opportunities of AI Arbitrators," kindly hosted at Watson Farley & Williams in London, along with sponsors TrialView and BlueLight Management.

As I said in my opening address, we are at a period of considerable transition: for our profession as well as the wider legal and commercial world. The City of London has always stood at the intersection

of law, commerce and innovation and we now face a new challenge: the increasing presence of artificial intelligence in the processes by which disputes are managed, analysed and perhaps, in time, decided. On that very theme, a particular highlight was the keynote address from Bridget McCormack, President and CEO of the American Arbitration Association. Even delivered remotely, her perspective on the development of AI-enabled ADR, including the AAA's AI Arbitrator initiative I mentioned above, brought a valuable institutional lens to the discussion.

From A to C: Three levels of AI in Arbitration

As I also suggested last night, it is often helpful in this context to think about AI in three roles:

A - AI as Assistant

At this level, AI is already embedded in our day-to-day work, often in ways we barely notice. It reviews documents, organises evidence, and helps support analysis far more quickly than most individuals or even teams could manage alone. In many cases, it's quietly handling the heavy lifting in the background by sorting, summarising, and streamlining tasks that used to take hours, even days. To be honest, the efficiencies are hard to ignore and with a growing number of senior judiciaries both using and advocating for it, it's becoming harder still to overlook. It's become part of the expected toolkit with stakeholders looking for faster turnaround times, lower costs and the ability to focus more on the parts of the job that actually require human judgment. Senior judiciary or otherwise, I can't think of many who would object to that!

B - AI as Co-pilot

This is where things start to get more interesting. AI isn't just helping with the admin anymore; it's starting to influence how a case is looked at in the first place. It can pull out patterns, point you towards certain cases, and even suggest how an argument might be structured. That might sound helpful, but it also means it's quietly steering the direction of thinking and often, that influence might not be obvious to everyone. You might be tempted to follow a line of reasoning because it's there, or because it seems well put together, without always stopping to ask why that route was suggested over another. In that sense, AI becomes less of a background tool and more of a sounding board which can shape how a dispute is framed from the outset. Used well, this can improve consistency and depth of analysis, but it also requires awareness. The more influence AI has over framing disputes, the more important it becomes to question, test and challenge the direction and output it suggests.

C - AI as Decision-Maker

As came through clearly in last night's discussions, the issue is no longer just whether AI can assist arbitration, but whether it might begin to play a more meaningful role in decision-making and where things become even trickier. It's one thing for AI to help organise material or suggest arguments, but it's another for it to move closer to actually deciding outcomes. As we know, arbitration isn't just about getting to an answer; it's about trust in the process and confidence that someone has properly weighed the case. Once AI starts edging into that space, the questions and considerations need to change – so, not just what outcome did it produce, but how did it get there and whether that can be properly understood, tested and challenged. The jury, as they say, is still out ...

The tensions we can't ignore

We obviously can't ignore the pressures building from the other direction. Cost, speed and accessibility are real drivers, and if AI can deliver outcomes more quickly, significantly more cost effectively and with greater consistency, it's easy to see why expectations will shift. In time, the question may no longer be whether AI should be used, but whether it can realistically be resisted.

That, in turn, raises a deeper and perhaps uncomfortable question more for a rainy day, but should there be a right to human decision-making? Parties obviously come to arbitration for an answer, but also for a process that is independent, fair and accountable. That process depends on judgment, discretion, an ability to understand nuance and context, and, at times, to move beyond rigid logic in a way that reflects commercial reality. Will AI ever be able to do that? This is where the tension emerges: what is economically attractive and what is institutionally acceptable. Time will tell.

For now, and for me, human oversight and accountability remain critical. AI is not a legal person. The UKJT draft Legal Statement records that AI does not have legal personality under the law of England and Wales and, accordingly, cannot itself be held legally responsible for harm; liability for harm arising from the use of AI must be attributed to legal persons under ordinary legal principle. That responsibility must remain with arbitrators, institutions and those deploying the technology, and not something to be outsourced.

So, is decision making as easy as AI, B, C?

Not quite. AI can already do a huge amount in dispute resolution, and, over time, it may even adequately resolve disputes in ways we're only beginning to understand, but there's a behavioural shift to keep in mind in the meantime. For newer practitioners especially, there's a real risk that answers get taken at face value simply because they look convincing.

That said, arbitration has always been adaptable, and with that comes responsibility. We shouldn't embrace AI blindly, nor push back against it for the sake of it. What's needed is a bit of balance: a thoughtful, deliberate approach grounded in the principles of independence, fairness and accountability.

That's why my Master's theme of "*Uniting Generations: Honouring our roots, growing our future*" feels particularly relevant. AI should enhance what we do, not replace how we learn. Roles like tribunal secretaries aren't just admin; they're where people learn how decisions are actually made, tested and challenged. If we lose that layer, we risk ending up with a generation relying on outputs they're not fully equipped to question, and we lose an important layer of human sense-checking along the way.

So, the real task is making sure AI strengthens arbitration, rather than undermines it. I think it will, and should, play a bigger role going forward, but ultimately, I don't think disputes will be decided by what the technology can do; but what we decide it should do. Those decisions, ironically and thankfully, are still human.



MATTHEW MOLLOY