SLTA | SWISS LEGALTECH ASSOCIATION

LAW&TECH INSIGHTS MAGAZINE

2026-2

CONSTRUCTION ARBITRATION:

THE AI AUGMENTATION EXPERIMENT UNDER ICDR PROCEDURE

ARINZE EZEAGU, NICOLAS TORRENT, ZIYU LIU,
NOMUNGEREL JAMSRANJAV

Editorial Team :

Nicolas Torrent, Ziyu Liu, Nomungerel Jamsranjav

www.swisslegaltech.ch

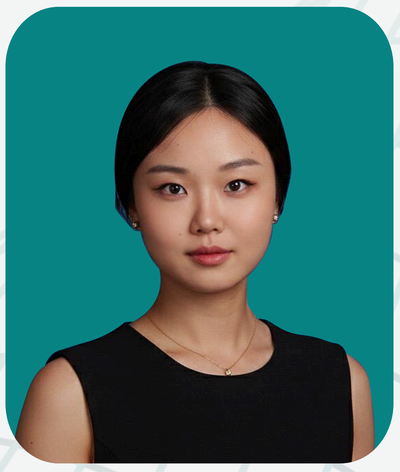
INTRODUCTION

On 3 November 2025, the American Arbitration Association's international division, the International Centre for Dispute Resolution (AAA-ICDR), launched what it calls an "AI arbitrator" for two-party, documents-only construction disputes[1].

The initiative is deliberately narrow in scope. It applies only to lower-value construction cases administered under AAA Construction Rules, requires the consent of both parties, and operates within a human-in-the-loop framework[2] where a trained human arbitrator reviews, revises if necessary, and signs every award. Yet its significance extends well beyond its current perimeter. This is the first time a major arbitral institution has deployed AI to generate draft awards in real disputes[3].

This article examines the design of the AAA-ICDR experiment, assesses the legal questions it raises (particularly around the enforceability of resulting awards under the New York Convention), and considers what the initiative signals for the broader arbitration landscape.






I. DESIGN OF THE AI ARBITRATOR

The AI arbitrator was developed in collaboration with QuantumBlack, AI by McKinsey[4], and trained on over 1,500 actual construction arbitration awards from AAA-ICDR's case repository. Construction cases were selected as the initial use case because they produce reasoned awards with clear chains of legal reasoning, which allowed the system to learn decision-making patterns from structured documentary records.

The process follows a three-tier structure:

- First, the parties upload their submissions and supporting evidence. The AI analyses the documents and produces summaries highlighting the factual and legal arguments of each side[5].
- Second, the parties review these summaries and provide feedback on whether the AI has correctly understood their positions. This step is a notable design choice: it gives parties a degree of procedural control that traditional arbitration does not always afford, since parties rarely get to confirm that the tribunal has grasped their arguments before deliberation begins[6].
- Third, the AI generates a draft award, which is then reviewed by a human arbitrator selected from AAA-ICDR's Construction Panel[7]. That arbitrator has full access to the case materials, can revise the draft, and ultimately issues andsigns the final award[8].

The commercial rationale is straightforward. Construction disputes are cash-flow sensitive. Project timelines are tight[9]. Traditional arbitration can be slow and expensive, with representation costs accounting for roughly 83% of total arbitration costs according to ICC data.

The AAA-ICDR projects that the AI arbitrator could reduce costs by 30 to 50% and resolution time by 20 to 35%[10]. These are the institution's own estimates, and they have yet to be independently verified, but the order of magnitude is plausible for document-only disputes where the main bottleneck is the time required for human review of evidence and drafting of the award.

II. ENFORCEABILITY AND PROCEDURAL LEGITIMACY

a. The “natural person” requirement

The central legal question is whether an award produced through this process can withstand scrutiny under the New York Convention on the Recognition and Enforcement of Foreign Arbitral Awards (the New-York Convention)[11]. The New-York Convention does not explicitly require a human arbitrator, but it was drafted in 1958 on the assumption that arbitrators would be natural persons.

Many national laws make this explicit. Article 1450 of the French Code of Civil Procedure, for instance, provides that only a natural person enjoying full legal capacity may serve as an arbitrator[12]. Swiss law, under Article 180 of the Private International Law Act (PILA), similarly assumes a human arbitrator when it addresses grounds for challenge.

The AAA-ICDR’s human-in-the-loop design is crafted precisely to satisfy this requirement[13]. The human arbitrator is intended to be more than a rubber-stamp; the institution emphasises that the arbitrator reviews the full case record, exercises independent judgment, and issues the award under their own name. In formal and practical terms, the AI is positioned as a tool that assists the arbitrator – as should be the case with any AI in our opinion.

Whether this framing holds in practice will depend on how actively human arbitrators actually engage with the AI’s output. If the review becomes performative, the distinction between an AI-assisted award and an AI-generated award could fail to meet the legal criterion.

The question, as Herbert Smith Freehills[14] has noted, is whether arbitrators can effectively scrutinise AI outputs while retaining the time-saving advantages the tool is designed to deliver. Representative experiment data will not surface before one or two years, considering the time necessary for sufficient awards to be challenged and the appeal courts to issue their decisions.

b. Due process and Article V(1)(d) of the Convention

Article V(1)(d) of the New-York Convention allows enforcement to be refused if the arbitral procedure was not in accordance with the agreement of the parties or, failing such agreement, with the law of the seat[15].

The opt-in mechanism addresses the first limb: both parties must consent to the use of the AI arbitrator, and parties that decline remain under AAA-ICDR's traditional process. Consent, however, must be informed. Parties need to understand what they are consenting to, including the role of the AI, the nature and limits of the human review, and the implications for any subsequent challenge.

The second limb raises questions that are harder to resolve in advance. If the seat of arbitration is in a jurisdiction whose arbitration law requires the arbitrator to personally conduct the proceedings (as many do), the use of AI to generate the substantive analysis and draft award could be characterised as an improper delegation of the arbitral mandate. For instance, the Chartered Institute of Arbitrators' 2025 Guideline on the Use of AI in Arbitration[16] states clearly that arbitrators should not delegate their decision-making powers to AI, though they may use AI to support the processing of submitted information.

c. The EU AI Act dimension

The EU AI Act[17] classifies AI systems used in the administration of justice and democratic processes as high-risk under Annex III, Category 8. An AI system involved in generating arbitral awards would, in all likelihood, fall within this classification if deployed within the EU or in respect of EU-seated arbitrations.

High-risk classification triggers obligations around transparency, human oversight, data governance, and risk management, which significantly increases the cost of development for the corresponding AI systems.

The AAA-ICDR's human-in-the-loop framework[18] is broadly consistent with the AI Act's oversight requirements, but compliance would require formal documentation of the system's design, training data provenance, and performance monitoring than is not currently publicly available to our knowledge.

d. Lessons from LaPaglia v. Valve Corp.

An instructive counterpoint to the AAA-ICDR experiment is the case of LaPaglia v. Valve Corp. (S.D. Cal., No. 3:25-cv-00833).

In April 2025, a party filed a petition to vacate an AAA arbitration award, alleging that the arbitrator had used ChatGPT to draft substantial portions of a 29-page award issued just 15 days after final briefing, over the holiday period. The petitioner argued that this constituted an improper delegation of the arbitral function under the Federal Arbitration Act.

TIn December 2025, the court declined to hear the challenge, but the case crystallised a question that the AAA-ICDR experiment must also confront: at what point does the use of AI in award-drafting cross the line from legitimate assistance to impermissible delegation?

In our opinion, the critical distinction lies in consent and institutional design. In LaPaglia[19] , the arbitrator allegedly used AI unilaterally and without disclosure. The AAA-ICDR experiment, by contrast, builds AI into the process by design, with full transparency and explicit party buy-in and consent.

An arbitral framework that openly integrates AI, with defined safeguards and informed consent, is on far stronger procedural ground than an individual arbitrator who outsources drafting to a chatbot without the corresponding transparency.

III. FURTHER IMPACT CONSIDERATIONS

The AAA-ICDR initiative is best understood as a controlled experiment. Its narrow scope (two-party, documents-only, construction) is deliberate. The institution has announced plans to expand to other industries (insurance is mentioned as the next candidate), other case types, and higher-value claims this year (2026)[20].

Several observations are worth making at this stage.

First, the experiment will generate real data on whether AI-assisted arbitration delivers on its efficiency promises. The projected cost and time savings expectations are significant, but they remain projections. Actual outcomes will nonetheless help determine the impact of AI on court proceedings – and these learnings may have far-reaching impacts, notably on state courts.

Second, the enforceability of the resulting awards will be tested only when a losing party challenges one. Until that happens, the legal questions discussed above remain theoretical. The human-in-the-loop model provides a strong foundation, but it has not been stress-tested in adversarial proceedings before a national court.

Third, the initiative puts pressure on other arbitral institutions. The ICC, LCIA, SIAC, and Swiss institutions such as SCAI will need to decide whether and how to engage with AI in their own proceedings. Inaction is itself a position, and as market expectations shift, it may become an increasingly difficult one to maintain.

Fourth, the experiment raises accountability questions that exceed procedural legitimacy:

If an AI-assisted award contains an error, who bears responsibility? The human arbitrator who signed it? The institution that deployed the system? The developer that trained the model (although most decline responsibility)?

The CIArb Guideline takes a clear position: arbitrators remain responsible for the award regardless of the AI system used[21]. But the practical allocation of liability in cases of serious irregularity remains untested.

CONCLUSION

The AAA-ICDR AI arbitrator is a narrowly scoped experiment that marks a first in institutional arbitration – and in the administration of justice in a broader sense. Its human-in-the-loop design addresses the most immediate legal objections, particularly the "natural person" requirement and the need for party consent. But the harder questions, around the quality and extent of human oversight, the enforceability of awards across jurisdictions, and the accountability framework for

For arbitration practitioners, the practical takeaway is that the AI arbitrator has just appeared as a new colleague to work with.

Whether other institutions follow the AAA-ICDR's lead, and how national courts respond to the awards this system produces, will determine whether this experiment remains an outlier or becomes the template for a broader transformation of the arbitral process – and perhaps of the administration of justice in general.

REFERENCES

1. See Dominic Sinnott, The AAA-ICDR AI Arbitrator: a new era in construction arbitration? (Gowling WLG, 6 November 2025) *https://gowlingwlg.com/en-ae/insights-resources/articles/2025/the-aaa-icdr-ai-arbitrator.* Accessed 5 February 2025.
2. For an understanding of Human-in-the-loop System, please see Swiss Legal Tech Association Tech & Law insight 2025-03: Arbitrating Tomorrow's Disputes in a World of Artificial Intelligence. Available at https://*www.swisslegaltech.ch/docs/en/documentation/Law___Tech_insights_2025-3__Arbitrating_Tomorrow_s_disputes_in_a_world_of_Artificial_Intelligence.pdf Accessed 5 February 2026.*
3. American Arbitration Association- international Centre for Dispute Resolution, "AAA-ICDR introduces AI introduces AI arbitrator for Construction Disputes" *https://www.adr.org/* accessed 5 February 2026.
4. The AAA-ICDR announced on September 17, 2025, in an official release that it collaborated with Quantum Black AI by McKinsey to develop a digital dispute resolution platform and AI arbitrator. *https://www.adr.org/* accessed 5 February 2026.
5. Ibid.
6. Ibid.
7. Sinnott, The AAA-ICDR AI Arbitrator, n 1
8. Ibid.
9. See Ahmed Ismail, `Harnessing AI to improve construction dispute resolution, https://www.consultancy-me.com/news/7098/harnessing-ai-to-improve-construction-dispute-resolution accessed 6 February 2026
10. Sinnott, The AAA-ICDR AI Arbitrator, n 1
11. See Articles V (1) (b), V (1) (d), or V (2) (b) of the New York Convention for Recognition and Enforcement of Arbitration Award of 1958
12. See also, Article 10(1) of the UAEs Federal Law No 6 of May 3,3018).

13. American Arbitration Association-International Centre for Dispute Resolution, ‘AAA-ICDR Introduces AI Arbitrator for Construction Disputes’ (n 2)
14. Nicolas?
15. Ibid(n11)
16. See Charted Institute of Arbitrators, Guideline on the Use of AI in Arbitration (CIArb 2025) https://www.ciarb.org/media/bpndtcgu/guideline-on-the-use-of-ai-in-arbitration Accessed 6 February 2026.
17. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) art 6.
18. Ibid(n2)
19. John LaPaglia v Valve Corporation, Petition to Vacate Arbitration Award, Case No 3:25-cv-00833-RBM- DDL (US District Court for the Southern District of California, Petition filed 8 Apr 2025) https://www.acerislaw.com/wp-content/uploads/2025/04/LaPaglia-v.-Valve-Corp.pdf
20. Sinnott, The AAA-ICDR AI Arbitrator, n 1
21. Chartered Institute of Arbitrators, Guideline on the Use of AI in Arbitration (n16)