GAR Global Arbitration Review

Lessons on AI from the Hackathon

9 min read | 30 April 2026

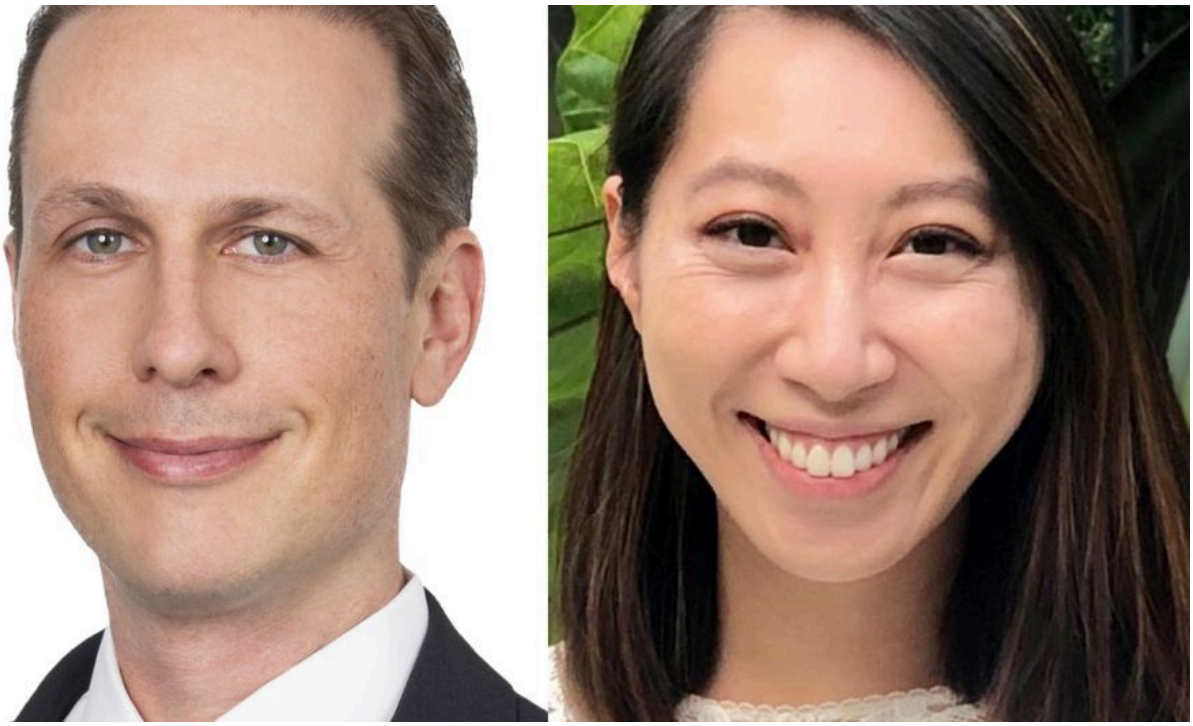

*Three Crowns CEO **Hugh Carlson** and **Megan Ma** of Stanford University reflect on lessons learned from judging the inaugural GAR-LCIA Hackathon, including the economics restraining AI innovation in international arbitration, and the future of work for all those in the space – from students to seasoned partners.*

AI has given renewed currency to author William Gibson's aphorism, "The future is already here – it's just not evenly distributed."

Sectors like programming and customer service are in the midst of a transformation but others, including seemingly vulnerable white-collar industries, have experienced only modest change. The extent of AI transformation grows more uneven daily.

And the law? The same sawtooth terrain in miniature – albeit one with far fewer dramatic peaks and valleys than the breathless commentary might suggest. On the transactional side, some practices are redirecting first-instance junior work to AI tools, while others are shrinking junior associate class sizes or attributing business services layoffs to AI.

But for most practitioners, the daily practice of law remains largely as it was a decade ago – and where practice has held steady, the business of law has been under little

pressure to evolve.

International arbitration's economic moat

International arbitration practitioners can be forgiven for questioning the loftier and gloomier predictions surrounding AI. While certain domain-specific AI tools are shifting how junior associates bill and spend their time, the practice remains substantially intact, especially at the senior level. More than most practices, international arbitration benefits from a modest but meaningful moat separating it from AI disruption.

The reasons will be familiar to GAR readers. The responsibilities of junior associates are sufficiently varied and judgment-intensive that current AI systems are a supplement more than a looming replacement. Public pools of international arbitration data are comparatively scarce, rendering domain-specific AI tools less predictive than their equivalents in, say, US commercial litigation or M&A due diligence.

Those same practices dwarf international arbitration in volume, attracting outsized investment in bespoke AI tooling that the arbitration market, for all its growth, cannot yet command.

The economics of the billable hour reinforce this. Senior litigators at top US firms now routinely bill more than $2,000 per hour; some exceed $3,000 and even third-year associates have crossed $1,000. At these rates, even marginal efficiency gains from AI tools are financially attractive. Capital follows returns, and international arbitration does not sit atop that hierarchy.

But change may be afoot. For the first time, two connected trends are introducing a credible challenge to arbitration's protective moat.

The first is the accelerating capability of AI. Eighteen months ago, cutting-edge AI systems could draft passable legal memoranda but struggled with the sustained reasoning demanded by complex arbitration. Today, "frontier models" like Claude and Gemini can synthesise large documentary records, identify inconsistencies across witness statements and reason persuasively about contested points of law.

Context windows – the amount of information that large language models can retain at a time – have grown dramatically, meaning that AI systems can now hold and reason across large swaths of case files. Agentic workflows – in which AI systems autonomously plan and execute multi-step tasks – are moving from research demonstrations to products in areas directly relevant to international arbitration: document review pipelines, procedural compliance monitoring and the assembly of chronologies across multilingual evidentiary records.

The advancements are already being applied. Practitioners report using AI as a first-instance drafter of written submissions – work that until recently would have been entrusted without question to junior associates.

We raise these not only as proof that arbitration's moat is narrowing, but as an illustration of something more consequential: an innovation curve that shows no sign of flattening.

The underlying driver is a set of empirical regularities known as scaling laws, describing the relationship between the computational resources invested in training an AI system and the performance of the resulting model. What makes them remarkable – and what distinguishes AI from most technologies – is that the relationship is not linear but follows a power law: as computation increases, performance improves in smooth and, to date, predictable curves.

The practical significance is difficult to overstate. Labs can invest hundreds of billions of dollars in next-generation training runs with a degree of confidence that, for now, the resulting models will be materially more capable than their predecessors.

The second trend is "vibe-coding" tools that allow those without programming experience to build bespoke, sophisticated and reliable applications by describing their needs in natural language. This challenges the long-standing assumption that practice-specific software requires significant capital investment, specialised engineering talent and years of development. A resourceful lawyer or law student can now prototype tools tailored to international arbitration's demands at a fraction of the cost and time previously required.

Of course, there are limitations to vibe-coded tools, at least when built by those unfamiliar with application development. Code generated through natural language prompting can contain subtle errors that a non-technical user may not detect, and the resulting applications often lack the robustness, testing and audit trails that enterprise environments demand.

But the pace of improvement here, too, is instructive. Claude Code, perhaps the exemplary vibe-coding tool, is barely a year old but its creator says it writes itself. Software companies can now design, build and ship products at a recently unimaginable pace. Even if an associate's vibe-coded prototype is not yet fit for enterprise deployment, the gap between prototype and polished product is closing.

The same trajectory applies to concerns over the security of client data. Legal tools built upon frontier AI models operate within secure, sandboxed environments no less robust than the cloud platforms on which firms already rely. The more pressing risk lies in tools built by amateurs, where data handling, access controls and logging may fall short of enterprise standards. But even here, efforts are being made to close the gap.

Consider the implications for international arbitration. The practice is not short of technology but is yet to experience the kind of disruption beginning to reshape other areas of law.

But now, pressure is emerging from several directions. The arbitration community, armed with vibe-coding tools, is beginning to build for itself. Domain-specific legal AI

companies are building tools with increasing relevance to arbitration, and a smaller number are now building specifically for it.

And the frontier AI labs themselves are producing new models whose capabilities extend naturally into the kind of complex, judgment-intensive reasoning that arbitration demands. Anthropic's recent launch of a legal plugin for Claude Cowork, which **wiped billions from legal tech stocks overnight** in what Wall Street analysts dubbed the "SaaSpocalypse", indicates that these labs are advancing beyond only supplying the models that legal technology companies build upon; they are beginning to blur the line between model provider and legal technology vendor.

Four takeaways from the GAR-LCIA Hackathon

It was an awareness of these converging trends that led to the first GAR-LCIA Hackathon.

Its premise was straightforward: the international arbitration community is now in a position to preview what AI-driven innovation in the space will look like. Earlier, such an exercise would have been largely speculative, not because the ideas were lacking but because execution remained too difficult for lay people. Today, that gap has narrowed sufficiently that the innovations produced by lawyers and students can move beyond the conceptual.

So what did we learn from judging this year’s competition?

First, in several leading international arbitration practices, lawyers are building sophisticated and useful tools. Practitioners are engaging with AI development iteratively, rapidly and in direct response to frictions in their daily work. We were shown live demonstrations of functioning software, not speculative proposals or slide decks.

One professional team built a tool to map the logical dependencies between issues in a case at the outset of proceedings, enabling tribunals to identify dispositive questions and sequence issues more efficiently from the first case management conference. Another tackled a more foundational problem: the inability of existing AI tools to operate across an entire case record – thousands of documents spanning PDFs, email archives, transcripts, and client systems – rather than the isolated slices most tools currently handle. These are practical instruments built by practitioners who understand the problems first-hand.

What set the strongest professional entries apart was not technical sophistication alone (several student entries were comparably ambitious) but an understanding, sharpened by years of practice, for what warrants a tool.

The second lesson is that many in the cohort entering the profession will regard AI-assisted development as more baseline competency than a novel skill. Candidly, several of the student proposals were among the most imaginative submissions across

the entire hackathon, reflecting a generation that has come of age with these tools and is less inclined to take their limitations for granted.

What also distinguished the strongest student entries, beyond their technical ambition, was their willingness to think laterally about problems to which the professional community has grown inured. One team built a passive background system that continuously processes incoming case correspondence and incrementally assembles a structured understanding of each dispute without requiring any change to existing workflows. The winning student team developed a tool that maps fragmented consent trails across emails, term sheets, and contract redlines to produce an evidence-linked record of what parties agreed to when negotiating an arbitration clause, and then stress-tests the resulting clause.

Third, there was little convergence on issues or solutions. Across more than 40 participating teams, no two entries addressed substantially the same problem and no two solutions were similar, let alone redundant.

This exceeded our expectations; we anticipated breadth and diversity but more overlap. What we received suggests that international arbitration contains a wide and largely untapped landscape of inefficiencies that the community can identify and that AI tools can address.

Finally, each solution targeted a specific procedural or substantive issue. Notably absent were sweeping ideas that might cut across multiple phases of arbitration or re-imagine key parts of the process. Every submission was supplementary in character, designed to enhance existing workflows rather than to replace them. This restraint is itself a form of judgment; knowing what not to automate is as important as knowing what to build.

Remember that the tools available to these participants will be materially more capable next year. If the inaugural GAR-LCIA Hackathon established that the community can produce functional tools for specific issues, future iterations may demonstrate something more, particularly if the insights of practitioner-builders and the infrastructure of dedicated legal technology platforms begin to converge.

The depth of one's understanding of international arbitration is becoming a direct predictor of the quality of the technology one can build for it. As AI grows more powerful, those practitioners whose command of the field has been forged over years of complex, high-stakes practice will be best positioned to build tools of consequence.

Hugh Carlson *is the CEO of Three Crowns, a lecturer at Harvard Law School and co-founder of legal AI conference Generative Legal.*

Dr Megan Ma *is the executive director of the Stanford Legal Innovation through Frontier Technology Lab (liftlab) and a lecturer in computational law. She sits on the*

board of directors of the International Institute for Conflict Prevention & Revolution (CPR) and co-chairs the ICC Task Force on AI.

The authors are grateful to friends and colleagues for their input into this piece, including ***Christopher Bogart*** *of Burford Capital and several members of the Three Crowns team.*

The judging panel for the 2026 GAR-LCIA Hackathon also included GAR's publisher ***David Samuels****, LCIA director general* ***Kevin Nash****, and independent arbitrator* ***Lucy Greenwood****.*