

Worldly Perspectives

(continued from previous page)

The question this column posed in earlier coverage of Croatia—whether legislative architecture alone would be enough—has effectively been answered by events: Croatia has now built the architecture. A confirmed pre-litigation framework, real financial incentives paired with real cost sanctions, a Level 7 qualification floor, a national ethics and disciplinary system, express recognition of online mediation, and a well-funded National Center with its own data-collection mandate together amount to a considerably more complete package than this column described as merely forthcoming just months ago.

Judge Šimac's own framing captures the shift: Croatia has moved from promoting mediation to actively managing and strategically developing it.

What the reform cannot yet answer is whether structure translates into use. Its architects seem to recognize as much—the mandatory-exposure categories are narrow by design, framed as a testing ground, and the National Center's new statistical mandate exists precisely so Croatia will finally have longitudinal data on whether information meetings move settlement behavior or merely add a checkpoint.

One domestic gap remains squarely relevant here: the reform now rests heavily on a single centralized institution whose capacity has not been tested at scale. A separate, cross-border gap is also worth noting in passing—Croatia still has not acceded to the [Singapore](#)

[Convention on Mediation](#)—though that bears on international enforcement rather than on why domestic mediation volumes remain low.

For now, the demand for mediation training in Croatia still appears to outpace the demand for mediation itself—a mismatch that, paradoxically, may be part of the solution, particularly once a planned five-year university mediation program begins enrolling students in 2027. A growing corps of trained mediators, judges, and advocates could well become the ambassadors both datasets suggest matter most.

But for the first time, that practitioner-driven push operates alongside a legal architecture built specifically to support it—which means the coming reporting cycles, not the next legislative debate, are where Croatia's mediation paradox will actually be tested. ■


SPECIAL SERIES: AI IN ARBITRATION AND LITIGATION—THE JAGGED FRONTIER

ADR Technology/Part 3 of 6

The AI Toolkit: Use Cases, Tools, and Prompting Techniques for Dispute Resolution

BY DOAK BISHOP & KARL HENNESSEE

Series Overview: This is the third of six articles examining artificial intelligence in dispute resolution—a series called, “AI in Arbitration and Litigation—The Jagged Frontier.” The kickoff article is available at “The Nature of AI in Dispute Resolution: Competence



Without Comprehension,” 44 Alternatives 129 (July/August 2026); Part 2 appears at, “The Surprising Things Artificial Intelligence Can Do: Capabilities and Properties in Dispute Resolution,” 44 Alternatives 159 (September 2026) (both articles are available in full text on Westlaw).

This month addresses the moves from theory to practice with practical use cases, tools, and prompting techniques; subsequent articles will cover advanced AI architectures for reasoning, causation, and analogical inference; the efficiency paradox—whether AI will make dispute resolution cheaper or more expensive, and the irreplaceable role of human judgment.

tion today, organized along a spectrum from fully automated tasks to judgment-intensive augmentation. It closes with 10 prompting techniques—five basic, five advanced—to assist practitioners in achieving the best results from AI.

The taxonomy that follows—automated tasks, augmented analysis, and simulation—is illustrative, not exhaustive. It does not cover every emerging use case, but it addresses the applications most directly relevant to arbitration and litigation practice.

Automated Tasks: Where AI Operates Near Full Autonomy

These are tasks where AI works with minimal human judgment, requiring quality control review rather than substantive analysis.

Bishop is a Dallas-based independent arbitrator and a retired King & Spalding partner who founded the firm's international arbitration practice 25 years ago. He was chair of the Institute of Transnational Arbitration in 2012-2015; past member of ICCA, LCIA, SIAC, AAA, and ASIL governance boards; past member of the U.S. delegation to the NAFTA 2022 Committee on Private Commercial Disputes, and an adviser to the U.S. Restatement (3rd) of the Law of International Commercial and Investor-State Arbitration. See <https://doakbishop.com>. Hennessee is a Senior Vice-President, Litigation, and Investigations at Airbus in France, where he leads the company's global activities in these areas. He holds or has held board seats at the ICC (Chair), HKIAC, LCIA, ICCA (Treasurer), AFA, and the Queen Mary School of International Arbitration. He is a frequent author and occasionally sits as arbitrator.

The first two articles in this series concerned the technical foundation: what large language artificial intelligence models are and what they can do. This article surveys how AI is being used in dispute resolu-

Document Review and Predictive Coding: Document review and predictive coding is the most established AI application in litigation, with more than a decade of deployment and judicial approval. On technology-assisted review, or TAR, and predictive coding, see *Da Silva Moore v. Publicis Groupe*, 287 F.R.D. 182 (S.D.N.Y. 2012); *Rio Tinto v. Vale*, 306 F.R.D. 125 (S.D.N.Y. 2015) (Judge Peck).

These remain the foundational judicial endorsements of AI-assisted review. Current platforms include Relativity aiR, Reveal Brain-space, Disco, and Everlaw.

Technology-assisted review and continuous active learning systems classify documents for relevance, privilege, and issue-coding across massive document sets. The current generation adds generative AI capabilities: document summarization, privilege log drafting, key-document identification, and automatic retrieval of facts and dates into timelines and chronologies.

AI contract review benchmarks show processing speeds up to 80 times faster than manual review. LawGeex, “AI vs. Lawyers: Comparing the Accuracy and Efficiency of Contract Review” (2018) (AI reviewed NDAs in 26 seconds vs. 92 minutes for human review). See also Deloitte, “Objections Overruled: The Case for AI-Enabled Legal Technology” (2020) (reporting AI performing document analysis and data extraction up to 80 times faster than lawyers).

The most recent generation of e-discovery tools adds a generative AI feature that can answer natural-language questions about the document set. Rather than constructing Boolean queries or running keyword searches, a practitioner can now ask, “Identify all documents in which the respondent’s project manager acknowledges schedule delays attributable to design changes,” and receive a ranked list with relevant excerpts highlighted.

For multi-jurisdictional disputes involving tens of thousands of documents in multiple languages, this capability can compress months of review into days.

As well as AI works for document review, there are well-understood limitations. Privilege determinations require closer human oversight than relevance coding. More important, AI cannot assess a document’s strategic significance—the judgment that a seemingly routine internal email is actually a pivotal piece of evidence. Pattern recognition identifies what

appears frequently, but it cannot identify what matters most.

Translation and Routine Drafting: AI can produce working translations across major languages, making it immediately useful for international arbitration’s multilingual demands. It can also prepare procedural orders and standard correspondence, generate Tables of Authorities and exhibit lists, proofread documents, and transcribe conversations.

Harvey’s Workflow Builder, launched in

AI Support For ADR

The process boosters: Damage modeling, case assessments and outcome prediction, and simulated witnesses examinations are among the numerous tasks detailed in this article that artificial intelligence can accomplish for arbitration.

Sample stunner: AI can systematically compare an expert’s current opinions with positions the expert has taken in prior cases—an exercise that one firm performed across 63 prior deposition transcripts.

Caveats? You know there are plenty. ‘Impressive breadth, but little judgment.’ ‘AI identifies what is statistically prominent, not what is strategically important.’ Read on.

June 2025, reports that it has already generated more than 18,000 custom automated workflows for tasks like these. Harvey AI supports contract analysis, due diligence, compliance, litigation, and custom workflows. Harvey reported as of December 2025 that over 18,000 custom workflows had been created since its Workflow Builder launch in June 2025.

The Workflow Builder—now known as Agent Builder—is a no-code platform that allows legal teams to design and deploy customizable, multi-step AI-powered processes

tailored to their firm’s specific needs rather than relying on generic AI prompts.

Clearbrief, which the American Arbitration Association-International Centre for Dispute Resolution now provides to its panel of arbitrators and mediators, is a Microsoft Word add-in that can generate hyperlinked timelines, Tables of Authorities, and formatted briefs directly within Microsoft Word.

The AAA/ICDR provides the AI-powered legal writing platform to its panel of 5,500 arbitrators and mediators. It was launched as a six-month pilot in January 2025. Capabilities include cite-checking (both legal and factual citations), timeline generation, deposition summaries, and Table of Authorities. See “Exclusive: American Arbitration Association Partners with Clearbrief,” *LawSites* (Jan. 14, 2025) (available at <https://bit.ly/4w9lmNF>).

Even routine drafting can be error prone. The plausibility-correctness distinction discussed in the second article is relevant here: language that reads correctly nevertheless may contain technical errors that can be overlooked on casual review.

Translation quality can degrade with specialized legal terminology and less common languages. Human review remains essential, especially for dispositive documents and sworn translations.

Augmented Analysis: Where AI Assists And Humans Decide

In many instances, AI can produce useful first drafts and analytical frameworks, but substantive human judgment is still required.

Legal Research: AI-powered legal research now spans general and specialized platforms. For general legal research, Westlaw’s AI-Assisted Research and Lexis+ AI have transformed the traditional research workflow. Traditional Lexis and Westlaw research requires practitioners to construct Boolean queries using precise search operators and connector logic.

By contrast, the Lexis+ and Westlaw AI-enhanced research products allow natural-language queries (“What is the standard for piercing the corporate veil in Delaware?”), automatically retrieve and rank relevant authorities,

(continued on next page)

ADR Technology

(continued from previous page)

generate summaries of holdings, and integrate with Shepard's Citations (Lexis) and KeyCite (Westlaw) to verify that cited cases remain good law. The result is that research tasks that previously required an associate to construct multiple targeted searches, review dozens of results, and manually Shepardize each authority can now be compressed into a single conversational query with integrated citation validation.

For international arbitration specifically, **Jus AI** enables natural-language queries across a database of more than 1.5 million documents, including arbitral awards, treaties, institutional rules, court decisions, and secondary sources from publishers including Oxford University Press, Juris [which co-publishes this newsletter with the CPR Institute], and Brill.

Jus AI 2, an international arbitration platform launched in September 2025, features an agentic AI planning agent that creates multi-step research strategies tailored to each query. One firm reported that its associates are 50% faster in research using Jus AI, and that the platform has become "their first port of call" for both quick questions and involved queries.

It claims it can analyze up to 75,000 documents per minute. The platform reports a 125% increase in retrieval relevancy compared to traditional legal research tools. Jus Mundi serves more than 650 arbitration teams at firms including Freshfields, A&O Shearman, White & Case, Quinn Emanuel, and Three Crowns. See "Jus Mundi Launches Jus AI 2: 'Breakthrough' Legal AI Combines Agentic Reasoning with Research Control," *LawSites* (Sept. 10, 2025) (available at <https://bit.ly/4fyCreO>); "Smarter AI for Sharper Minds," *Daily Jus* (June 2025) (available at <https://bit.ly/4wppJ7r>).

Harvey offers a custom-trained case law model built in partnership with OpenAI, designed to find argumentative support rather than just topical matches by injecting domain-specific legal knowledge and reasoning into the foundation model.

Winston Weinberg, Harvey's co-founder and CEO, quoted in OpenAI, "Customizing Models for Legal Professionals" (2024), noted, "With case law research, you're finding ammo for your argument, and that's much more difficult to do."

Legora, an AI-powered legal workspace now deployed at major arbitration-focused law firms, offering document review, comparison, and inconsistency detection at scale. It is deployed firmwide at White & Case, Scandinavian law firm BAHN, London-based international firm Bird & Bird, and Italian law firm Portolano Cavallo, among others.

Legora's Tabular Review feature transforms large document sets into an interactive grid where each document is a row and AI-generated queries are columns—enabling clause comparison and inconsistency flagging across hundreds of documents simultaneously.

Legora's Workflows feature uses agentic AI to complete multistep legal tasks from plain-language instructions. The platform is ISO 27001, SOC 2, GDPR, and ISO 42001 (AI governance) certified. See "Arbitration at a Turning Point: How AI Is Redefining the Practice," Legora Blog (October 2025 (available at <https://bit.ly/4xbxbnc>); "White & Case to Roll Out Legora's Legal AI Software Across Office Network," *Global Legal Post* (Dec. 22, 2025) (available at <https://bit.ly/4xiFkXc>).

For arbitration specifically, Legora can organize tens of thousands of case files into a clear chronology within minutes, with every event linked to its source and missing periods in the record flagged automatically.

The hallucination problem remains the central constraint. The Stanford empirical study published in 2025—the first rigorous, "preregistered" evaluation of AI legal research tools (the researchers publicly committed to their methodology and hypotheses before collecting data in order to prevent cherry-picking results)—found that even the best performing tool, Lexis+ AI, hallucinated 17% of the time. Westlaw AI was reported to have hallucinated at a 33% rate. Viraj Magesh, Faiz Surani, Matthew Dahl, Mirac Suzgun, Christopher D. Manning & Daniel E. Ho, "Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools," 22 *J. Empirical Legal Stud.* 216–242 (2025) (preprint May 2024; published April 23, 2025). (Tool-by-tool results: Lexis+ AI: 65% accuracy, 17% hallucination rate; Westlaw AI-Assisted Research: 42% accuracy, 33% hallucination rate; GPT-4 (no RAG): 43% hallucination rate).

For legal questions about non-U.S. jurisdictions, the numbers are dramatically worse: when researchers tested queries about the law

of specific cities and countries, hallucination rates reached 45% for Los Angeles, 55% for London, 61% for Sydney, and 100% for some specialized Australian statutes. This reflects the fact that LLM training data is heavily weighted toward U.S. federal courts; state, local, and foreign law is far less reliably represented.

On jurisdiction-specific hallucination rates, see Damian Curran et al., "Place Matters: Comparing LLM Hallucination Rates for Place-Based Legal Queries," AI4A2J-ICAIL25 (2025), arXiv:2511.06700. The study tested AI legal research tools on queries about the law of specific jurisdictions and found hallucination rates of 45% for Los Angeles, 55% for London, 61% for Sydney, and 100% for some specialized Australian statutes such as the Residential Tenancies Act 2010 (NSW). See also Vals AI Legal Research Report (October 2025): ChatGPT with web search achieved 80% accuracy; dedicated legal AI tools scored 78–81%.

Without exception, a lawyer should never cite an AI-generated authority without independent verification.

Arbitrator Selection: One of the most consequential decisions in any arbitration—the selection of the tribunal—is now being informed by AI-driven data analytics. Jus Connect, developed by Jus Mundi, is a "smart directory" that uses extractive AI to pull data from available awards and enforcement proceedings, mapping arbitrators' experience, cases, industries, languages, and professional connections.

Jus Connect is Jus Mundi's "smart directory" for arbitration professionals. It uses extractive AI to pull data from available awards and enforcement proceedings, supplemented by institutional partnerships, to map arbitrators' counsel's, and experts' experience, cases, industries, languages, and professional connections. The platform includes a built-in conflicts checker.

Jus Mundi is also developing a more advanced "professional search" AI tool that would go beyond data consolidation to recommendation—identifying potential arbitrators for a case—though this raises ethical questions still under consideration. See Camille Esnou, "Selecting an Arbitrator in the Age of AI," *Daily Jus* (December 2025) (available at <https://bit.ly/4xfV1OF>) (reporting on a New York Arbitration Week panel at White & Case). See also Charles Russell Speechlys case study reporting 80% faster arbitrator appointment

checks using Jus Connect. “Charles Russell Speechlys Accelerates Arbitration Research by 50% and Arbitrator Appointments 80% with Jus Mundi,” *Daily Jus* (Dec. 4, 2025)(available at <https://bit.ly/4irlCEr>.)

The platform includes a built-in conflicts checker that can identify relationships between prospective arbitrators and parties, counsel, or witnesses.

Arbitrator Intelligence provides a complementary approach, offering data-driven arbitrator research and analytics based on party feedback—essentially a structured database of how arbitrators have actually performed, as assessed by the parties who appeared before them. It offers a complementary approach to Jus Connect’s document-based profiles.

Institutional selection processes are also evolving: the ICDR now uses data-driven tools that scan resumes for expertise, monitor an arbitrator’s current caseload to assess bandwidth, and run automated conflict checks.

The limits are important. AI can aggregate data about an arbitrator’s prior cases, co-arbitrators, and published positions. It cannot assess temperament, hearing-room demeanor, deliberative style, or the interpersonal dynamics that experienced practitioners consider when constituting a tribunal.

The selection of arbitrators remains one of the most judgment-intensive tasks in arbitration practice—but AI now provides a richer factual foundation for that judgment.

Brief and Submission Drafting: AI can produce first drafts of legal memoranda, arbitration submissions, motions, and briefs that follow standard legal practice and cite relevant authorities. The output is structured, well-organized, and often reads as if written by a capable associate. Harvey and Clearbrief provide this capability.

The limitations discussed in the second article in this series are still true: impressive breadth, but little judgment. AI drafts may contain incorrect legal conclusions, misapplied precedent, or fabricated citations. The analysis that sounds most like what a good lawyer would write may be substantively wrong in ways that require legal expertise to detect.

Clearbrief’s Cite Check Report, launched in late 2025, attempts to address this problem by providing an automated audit trail confirming that every citation—both legal and factual—has been systematically verified against primary sources.

Contract Analysis: AI can review and compare contract provisions across portfolios of hundreds of agreements, identifying deviations from standard language, flagging unusual provisions, and retrieving key terms.

Harvey’s Vault feature enables bulk analysis of large document collections, allowing lawyers to pose a single question that is applied across an entire deal room or contract portfolio simultaneously—rather than reviewing each document individually. See Harvey, “Design Customizable Workflows With Workflow Builder” (June 2025); Harvey, “Top 5 Product Releases of 2025” (December 2025).

Luminance, Kira, LegalOn, Spellbook, and Legora serve similar functions with different strengths. Luminance focuses on M&A due diligence and is widely used for cross-border transactions. Kira Systems (now part of Litera) specializes in contract data extraction. LegalOn provides contract review and negotiation support. Spellbook integrates with Microsoft Word for contract drafting. Legora offers AI-powered document review, comparison, and inconsistency detection at scale through its Tabular Review feature, and has been deployed at White & Case offices world-wide.

The critical limit for contract analysis is that AI can identify what differs between contracts, but cannot determine why a deviation matters in the commercial context. A non-standard limitation-of-liability clause may be intentional and carefully negotiated, or it may be a drafting error. AI cannot distinguish between them.

Evidence Organization and Issue Identification: AI can now assist with one of the most time-consuming aspects of case preparation, organizing evidence and identifying the legal issues it raises.

LLMs perform well at spotting emerging legal issues from a set of facts, and several platforms can automatically categorize documents by issue, party, date, and relevance. LLMs perform well when identifying emerging legal issues from a set of facts. See LegalBench (2023) (comparing LLM performance on legal reasoning tasks including issue spotting, interpretation, and rhetorical tasks). See also “Artificial Intelligence in Arbitration: Evidentiary Issues and Prospects,” *Global Arbitration Review* (2025) (discussing AI’s role in evidence identification, synthesis, and organization).

For arbitrations involving voluminous documentary records, AI can create structured

issue matrices that map each document to the claims and defenses it supports—a task that would otherwise require weeks of manual review. Jus AI can extract timelines from case documents, and Clearbrief can generate hyperlinked chronologies with citations to source material.

The limitation is the same one that applies to summarization: AI identifies what is statistically prominent, not what is strategically important. A fact mentioned once in a 200-page deposition transcript may be the most important fact in a case, but if it appears only once, AI may not flag it as significant unless specifically prompted.

Inconsistency Analysis: One of AI’s most powerful current capabilities for advocates—and one that is often underappreciated—is systematic inconsistency analysis across a case record. General-purpose LLMs can perform these analyses today when given the relevant documents in a single conversation.

Specialized platforms like Legora’s Tabular Review and Casefleet’s fact-extraction tools can perform them at scale across large document sets, and Casefleet’s AI-powered fact extraction links each factual assertion to source documents, enabling practitioners to cross-reference assertions across witnesses and submissions. See also Freshfields, “The virtual witness stand: AI’s new role in arbitration advocacy” (November 2025) (available at <https://bit.ly/4w7VjWS>).

Comparing a witness statement, for example, a practitioner can upload a witness statement and a deposition transcript and ask: “Identify every factual assertion in the witness statement that is contradicted by, unsupported by, or inconsistent with the deposition testimony.”

The AI will produce a structured comparison with specific citations to both documents. The same analysis can be performed across multiple witnesses, comparing Witness A’s account of a meeting with Witness B’s account, flagging discrepancies in dates, attendees, what was said, and what was agreed.

The applications extend beyond individual witnesses. AI can compare a party’s witness testimony against its own documentary record—checking whether witness accounts are consistent with contemporaneous emails, meeting minutes, board resolutions, and

(continued on next page)

ADR Technology

(continued from previous page)

project reports. It can track shifts in a party's factual narrative across successive submissions, identifying where the position taken in a Statement of Claim diverges from the position in a Reply. And it can systematically compare an expert's current opinions with positions the expert has taken in prior cases—an exercise that one firm performed across 63 prior deposition transcripts (see below).

The limitation is that AI identifies textual inconsistencies—factual assertions that conflict with one another—but cannot assess whether an inconsistency is material. A witness who remembers a meeting occurring on Tuesday rather than Wednesday may have made an insignificant error; a witness whose account of what was agreed at the meeting contradicts the contemporaneous minutes may have a credibility problem. That distinction remains the advocate's judgment.

Damage Modeling: For damages, AI can now build first-pass financial models, process structured datasets, perform statistical analyses, and construct lost-profits calculations.

In international arbitration, in which quantum assessments can involve complex discounted cash flow models, country risk adjustments, and multi-currency calculations, AI can construct a first-pass valuation from financial statements, sensitivity-test key assumptions across ranges, and flag anomalies in an expert's calculations—work that would take a financial analyst days to perform manually.

This is one of the higher-confidence use cases because code and computational output can be tested and verified against known results—unlike prose, where errors may be difficult to detect. Selecting the valuation methodology and assessing the reasonableness of assumptions remain a matter of expert human judgment.

Case Assessment and Outcome Prediction: For litigation, Lex Machina provides predictive analytics offering data-driven insights on case outcomes, judicial behavior, opposing counsel patterns, and historical success rates. The platform is focused on U.S. federal and state litigation rather than international arbitration.

These tools allow practitioners to benchmark a case against historical data—for

example, how a particular judge has ruled on summary judgment motions in patent cases, or what the median damage award has been in a specific claim type. For arbitration, outcome prediction remains more limited because most awards are not publicly available, although the growing databases at Jus Mundi and similar platforms are beginning to enable comparable analytics for international disputes.

Strategic Augmentation And Simulations

Beyond its more automated features, there exist advanced use cases for AI.

AI-Simulated Cross-Examination, Witness Preparation, and Outline Building: Freshfields published a detailed assessment in November 2025 describing how legal teams now create AI “witness bots” in preparation for arbitration hearings. See Freshfields, *The Virtual Witness Stand*, linked above. It describes creating AI witness bots loaded with case materials and programmed to respond as a witness would enabling 24/7 practice sessions.

The process involves loading case materials into the AI system—witness statements, deposition transcripts, documentary evidence, and background information—and instructing it to respond as the witness would likely answer during cross-examination, including evasions, qualifications, and attempts to redirect. The examining lawyer can run multiple simulations, practicing against the AI bot, refining lines of questioning.

Some simulations deliberately include inaccurate responses to test whether the examiner will catch them. The applications extend to expert witnesses. A January 2026 report documented how one law firm reported using AI to review 63 prior deposition transcripts of a particular expert across multiple cases, uncovering past statements inconsistent with the expert's current report—an exercise that might have taken associates weeks of manual review. “AI's Promising Role With Expert Testimony,” Esquire Deposition Solutions (January 2026) (available at <https://bit.ly/4hO8Jnt>).

AI can also assist with building cross-examination outlines. After analyzing the documentary record and witness statements, AI can identify the lines of attack—internal

contradictions, inconsistencies with contemporaneous documents, prior inconsistent statements, and admissions against interest.

A practitioner can prompt the system to generate a structured cross-examination outline organized by topic, with each line of questioning linked to specific documentary support. The outline is a starting point, not a finished product—the judgment about sequence, emphasis, tone, and strategic pacing remains entirely with the advocate—but the factual foundation AI provides can be comprehensive in ways that manual preparation may miss.

AI cannot replicate witness demeanor, emotional responses, cultural communication patterns, or the psychological dynamics of live examination. It simulates the content of testimony, but not the experience of testifying. These tools supplement traditional preparation, but they should not be expected to replace it.

Simulated Tribunal Decision-Making and Outcome Prediction: Practitioners are also creating AI bots that can simulate how a specific arbitrator or tribunal might analyze a dispute, or the questions they might ask counsel. Practitioners create AI bots assigned the role of a specific arbitrator, provide submissions and evidence, and ask for analysis of each argument's likely reception. Prior awards and arbitrator data from Jus Mundi and Arbitrator Intelligence can inform these simulations. The practice is emerging but not yet empirically validated for predictive accuracy.

The approach involves assigning the AI a role—“You are an experienced international arbitrator applying the ICC Rules and Swiss substantive law”—providing the parties' submissions and evidence, and asking for an analysis of each argument's likely reception. When prior awards by the specific arbitrator are available through platforms like Jus Mundi or Arbitrator Intelligence, these can be fed into the AI system to inform how the simulation reflects that arbitrator's known tendencies.

AI-simulated jury analysis has also emerged, though with heavier caveats. See Hayley Stillwell & Sean A. Harrington, “Michael Scott Is Not a Juror: The Limits of AI in Simulating Human Judgment,” (Aug. 11, 2025) (available at <https://bit.ly/4cjinLOO>) (anchoring analysis for the limitations of AI-simulated jury deliberation).

These outputs should be treated as directional indicators. Actual jury behavior depends

on demeanor assessment, group dynamics, and contextual factors that AI models cannot yet capture.

Across all simulation use cases, the Harvard/BCG study's warning from our second article applies with special force: when AI operates outside the frontier—and strategic prediction is outside it—users who rely on AI perform worse than those who work without it. Fabrizio Dell'Acqu, et al., "Navigating the Jagged Technological Frontier," Harvard Business School Working Paper No. 24-013 (2023).

So counsel should tread carefully in this area, remaining vigilant.

AI Decision-Making: The AAA-ICDR has gone further, launching in November 2025 an AI-native arbitrator program for documents-only construction disputes. The AI is trained on reasoning frameworks derived from more than 1,500 construction cases, which appear to have been annotated for issue categories, outcomes, and reasoning patterns, though AAA-ICDR has not publicly detailed the annotation methodology.

The AI analyzes submissions, deconstructs claims, generates timelines, and produces a draft award. But at least for now, the AI Arbitrator is subject to review by a human arbitrator who validates the reasoning and finalizes the decision.

The AAA estimates cost reduction in eligible cases of 30% to 50%. See "AAA's AI Arbitrator: A Glimpse of What's Next?" *Global Arbitration News* (November 2025).

AI Use by Arbitrators: AI is not only a tool for counsel. Arbitrators themselves are increasingly using AI to manage their caseloads more efficiently. The 2025 Queen Mary University/White & Case Survey found that 52% of respondents believe arbitrators will increasingly rely on AI over the next five years.

Law firm Addleshaw Goddard's AI analysis (September 2025) (available at <https://bit.ly/4wGLkc7>) identifies four categories of AI support for arbitrators: (i) administrative tasks such as transcription and compliance tracking; (ii) sorting, categorizing, or summarizing evidence and pleadings; (iii) conducting legal research or analyzing legal arguments; and (iv) assessing and evaluating cases, including drafting parts of an award. See also CI Arb Guidelines on the Use of AI in Arbitration (March 2025) (available at <https://bit.ly/3RU5mAN>) (requiring AI be used only as a "supportive

tool" and that arbitrators verify all AI-generated information).

Potential arbitrator AI applications include transcribing hearings, tracking compliance with procedural orders, sorting and summarizing large volumes of evidence and pleadings, conducting legal research, analyzing the parties' legal arguments, and—most controversially—assisting with drafting parts of an award.

The institutional guidelines converge on one point: AI may be used as a "supportive tool" but must not be permitted to influence or substitute for the arbitrator's independent decision-making. The arbitrator remains personally responsible for every aspect of the award, regardless of whether AI contributed to the drafting process.

Arbitrators should also be aware of evolving guidelines considering the extent to which they should disclose their arbitration AI use.

The Importance of Prompting Techniques

The same AI model can produce output that is excellent or of limited usefulness depending on how it is prompted.

This is not a minor efficiency point; it is a fundamental feature of how LLMs work. A well-prompted model is a genuinely useful tool. Jason Wei et al., "Chain-of-Thought Prompting Elicits Reasoning in Large Language Models," NeurIPS 2022; Tom Brown et al., "Language Models Are Few-Shot Learners," NeurIPS 2020. While prompting skill remains important today, the trajectory of AI development—toward agentic systems and reasoning models—suggests that the need for detailed prompt engineering may diminish over time.

What follows are five of the most important basic techniques and five advanced techniques for legal practitioners.

Five Essential Basic Techniques

1. **OBJECTIVE, SPECIFICITY AND CONTEXT.** Clearly define your objective. What precisely do you want AI to provide? Be specific and provide AI with the necessary context to understand the background of the problem

and what you want. Tell the AI system the applicable law, the jurisdiction, the procedural framework, and the factual context.

Example: "Under Swiss substantive law and the ICC Rules, draft the jurisdictional section of a Statement of Defense arguing that the arbitration clause in Section 14.2 of the Share Purchase Agreement does not cover tort claims arising from pre-contractual misrepresentation."

2. **FORMAT SPECIFICATION.** Specify the structure, length, and format you want. AI will match whatever format you request—bullet points, numbered analysis, IRAC (Issue, Rule, Application, Conclusion—the standard legal analysis framework) structure, executive summary, or narrative prose.

Example: "Provide your analysis in IRAC format. Limit your response to 500 words."

3. **CONSTRAINT SPECIFICATION.** Tell the AI what you do not want it to do or consider.

Example: "Only cite cases decided by the U.S. Court of Appeals for the Second Circuit after 2018. Do not include any authorities you cannot verify."

4. **ITERATIVE REFINEMENT.** Treat the first response as a draft, not a final product, and follow up with further questions to sharpen the analysis.

Example: "The analysis of the third element is too conclusory—expand it with specific factual support."

AI improves substantially with iterative feedback within a single conversation.

5. **ROLE ASSIGNMENT.** Tell the AI what role to play. A prompt that begins "You are an experienced international arbitrator with expertise in construction disputes" will produce substantially different—and usually better—output than the same question asked without assigning a role.

Example: "You are a senior associate at an international law firm specializing in investment treaty arbitration. Analyze whether the respondent state's regulatory measures constitute indirect expropriation under the applicable bilateral investment treaty."

Five Advanced Techniques

1. **CHAIN-OF-THOUGHT PROMPTING.** Instruct the AI system to reason through the problem step by step before reaching a conclusion.

(continued on next page)

ADR Technology

(continued from previous page)

Example: “Analyze this force majeure claim step by step. First, identify the contractual language. Second, determine the applicable legal standard under English law. Third, assess whether the pandemic satisfies each element. Fourth, consider counterarguments. Finally, state your conclusion and include your confidence assessment.”

2. FEW-SHOT PROMPTING (IN-CONTEXT LEARNING). Few-shot prompting means showing AI a few examples of the output you want before asking it to produce new work. This is a powerful technique for imposing consistency across large-volume tasks.

Example: “Here are three examples of how I want witness statement summaries formatted: [Examples]. Now summarize the following witness statement using the same format.”

3. ADVERSARIAL/SELF-CRITIQUE PROMPTING. After AI produces an analysis, instruct it to critique its own work.

Example: “Now argue the opposing position with equal rigor. Identify every weakness in the analysis you just produced.”

Unlike people, AI has no emotions and does not “feel” defensive about its work; it can critique its own work with the same rigor with which it produced it.

4. MULTI-PERSPECTIVE ANALYSIS. Ask AI to analyze the same problem from multiple viewpoints sequentially.

Example: “Analyze this dispute from four perspectives: (1) as counsel for the claimant; (2) as counsel for the respondent; (3) as the tribunal; (4) as a mediator.”

5. LAYERED DECOMPOSITION. Break complex legal questions into sequential sub questions, with each answer feeding into the next. This prevents the “early convergence” problem, in which AI reaches a conclusion prematurely.

Example: “Stage 1: Is there a valid arbitration agreement? Analyze only this question now.” “Stage 2: Does it cover the claims asserted?” “Stage 3: Are there procedural bars to jurisdiction?”

The Verification Imperative

Every use case described in this article—from automated document review to strategic

simulation—is subject to the same structural constraint. AI produces output with identical confidence whether the underlying analysis is right or wrong.

The verification continuum from the second article in this Special Series, AI in Arbitration and Litigation—The Jagged Frontier, provides a means of assessing the proper level of review: light review for automated formatting tasks, substantive review for augmented analysis, and full independent analysis for anything involving judgment, strategy, or causal reasoning.

Counsel must also assess whether uploading client documents to AI platforms is consistent with their confidentiality obligations under Model Rules of Professional Conduct 1.6 and any applicable protective orders.

Every tool described in this article involves transmitting client information to a third-party platform, and the ethical obligation to protect that information does not diminish because the recipient is a machine. The distinction matters in practice: general-purpose AI platforms such as ChatGPT and Google Gemini, in their standard consumer versions, do not guarantee that user inputs will not be used for model training or accessed by the provider’s personnel.

By contrast, specialized legal AI platforms such as Harvey and Clearbrief contractually guarantee that customer data will not be used for model training, implement zero-data-retention policies with underlying model providers, enforce strict access controls and workspace isolation, and maintain SOC 2 Type II and ISO 27001 certifications.

Counsel should, at a minimum, review the platform’s data processing terms, confirm that client data will not be used for training, verify data encryption and retention policies, and ensure that the platform’s security posture is consistent with the sensitivity of the materials being uploaded.

Institutional guidelines are converging on certain principles. The Silicon Valley Arbitration & Mediation Center, Chartered Institute of Arbitrators, the SCC Arbitration Institute, the AAA-ICDR, and the Vienna International Arbitral Centre have all published frameworks emphasizing transparency, human oversight, non-delegation, and verification. See SVAMC Guidelines on the Use of Artificial Intelligence in Arbitration (April 2024) (available at <https://bit.ly/4wI5v9S>); CI Arb Guidelines on the Use

of AI in Arbitration (March 2025) (available at <https://bit.ly/4hVDUx4>); SCC Guide (October 2024) (available at <https://bit.ly/4554jkl>); AAA-ICDR Guidance (March 2025) (see www.adr.org/ai-governance); VIAC Note on the Use of Artificial Intelligence in Arbitration Proceedings (April 2025) (available at <https://bit.ly/4z4oWLC>). (Article 6 of this series will include a detailed treatment of ethical frameworks).

See also ABA Formal Opinion 512 (July 29, 2024) (available at <https://bit.ly/4fQAJEt>): lawyers have a duty of competence to have a “reasonable understanding” of AI technology—a competence requirement that this series of articles is designed to help practitioners meet.

* * *

AI tools for dispute resolution are no longer experimental.

They are in daily use across the profession, and the lawyers who understand both the tools and the techniques for using them effectively will gain a significant advantage. The taxonomy in this article—automated tasks requiring light review, augmented analysis requiring substantive review, and strategic work requiring full human judgment—provides a framework for determining the appropriate level of review of AI’s output.

But the most powerful tool in the AI toolkit is the lawyer’s own skill in prompting—the ability to tell AI precisely what to do, in what format, under what constraints, and from which analytical perspective. That skill is learnable and improvable. The 10 techniques described above are not the only useful prompting techniques, but they provide a strong starting point.

* * *

The next article examines the advanced AI architectures being developed to address the limitations most visible in the use cases described here: reasoning, causation, and analogical inference when using AI, including a “reasonable understanding” of the technology’s capabilities and limitations, addressing obligations under Model Rules of Professional Conduct 1.1 (competence), 1.6 (confidentiality), and 3.3 (candor to the tribunal).

