

ARTICLE

The Document that AI Never Finds: Rethinking Document Production in Arbitration

BY SARAH MALIK AND PAVITHRA RAJENDRAN



We all know that document production is rarely the most straightforward stage of international arbitration proceedings. A dispute arising from a long commercial relationship can easily involve years of emails, contracts, financial records, and internal correspondence held by numerous custodians across disparate systems globally.

Technology-assisted review has been a mainstay of arbitral practice for several years. What has shifted fundamentally is the capability of the tools. Current artificial intelligence (“AI”) systems do not merely match keywords; they analyse large datasets, evaluate documents by content and context, and rank material for human consideration. Used effectively, they significantly reduce both the time and cost of document disclosure.

However, a problematic issue arises when AI is relied upon to decide which documents require no human review whatsoever. A simple hypothetical scenario demonstrates the procedural risk.

A HYPOTHETICAL SCENARIO

A party collects 50,000 documents and runs them through an AI-assisted review platform. The system flags 20,000 as potentially relevant. The legal team reviews those 20,000 documents and the remaining 30,000 documents are not reviewed.

Later in the proceedings, the party serves a witness statement from a senior executive. He states that he was not involved in a key contractual decision and had no knowledge of the relevant discussions at the time.

During a later stage of the arbitration, the opposing party produces an email retrieved from its own disclosure. It is a contemporaneous exchange directly involving the witness, and it fundamentally contradicts his account. Unbeknownst to the legal team, a copy of that same email had resided in the first party’s collection all along, buried within the 30,000 documents the system screened out.

Missing an individual document does not, of itself, render a review process unreasonable; human reviewers miss critical documents with regular frequency. However, the party will now face immediate and uncomfortable scrutiny from both the opposing party and the tribunal:

- What was the review criteria?
- How was the review criteria configured?
- Was the exclusion methodology ever sampled or validated? If not, why not?
- To what extent did lawyers oversee the automated decisions?
- Can the party properly account for the role AI played in filtering the record?

It is easy to see from the above example that the consequences extend far beyond a procedural reprimand. By the time the email surfaces, the witness has committed to his statement, the party has built a cornerstone of its case around that evidence, and the tribunal must decide what weight, if any, can now be attached to his testimony.

DOCUMENT PRODUCTION AS AN EVIDENTIARY EXERCISE

Document production in international arbitration remains a targeted exercise, restricted to documents that are relevant to the issues in dispute and material to their outcome.

Article 3 of the IBA Rules on the Taking of Evidence in International Arbitration (the “**IBA Rules**”) reflects this balance. It requires a requesting party to identify specific documents or narrow, well-defined categories, and to articulate their relevance and materiality. Production may be resisted on established grounds, including legal privilege, commercial confidentiality, unreasonable burden, or broader procedural fairness.

Whatever institutional framework applies, the overarching objective is to place before the tribunal the evidence necessary to resolve the dispute fairly. A party is not required to search every conceivable source that might contain something of passing interest. That core purpose must govern how AI tools are deployed. AI allows counsel to interrogate far larger populations than would otherwise be feasible, identifying conceptual linkages that keyword searches routinely miss. This is particularly valuable where parties used inconsistent terminology, where documentation spans multiple languages, or where a crucial concession is buried within a long chain of correspondence.

Using the scenario above, a standard keyword search may well have missed the email if the witness used unconventional phrasing. Conversely, a well-configured AI tool ought to have flagged it precisely because of its context.

However, AI platforms remain entirely dependent on the parameters set by their operators. If the legal issues are mischaracterized at the outset, or if the tool is given insufficient background, an automated review will simply produce an efficient processing of the wrong material. Had counsel framed the review criteria strictly around the technical terms of the contract rather than the identity of the decision-makers, the executive’s involvement would never have registered as a relevant query.

The IBA Rules draw no qualitative distinction between manual and automated search methods. Standards of relevance, materiality, privilege, and proportionality apply with equal force whether a document is identified by a junior associate, a boolean search string, or an advanced machine-learning algorithm. AI alters the mechanism of discovery; it does not alter the evidentiary standards by which disclosure is judged.

This principle is reinforced by emerging soft-law instruments, notably the Silicon Valley Arbitration and Mediation Center (“**SVAMC**”) Guidelines on the Use of Artificial Intelligence in Arbitration. Under Guideline 4, counsel remains responsible for any uncorrected errors or inaccuracies in AI output. Guideline 3, while imposing no general duty to disclose the use of AI, contemplates that disclosure may be required where it bears on the evidence or the integrity of the proceedings. In document review, that could extend to how the tool was selected, configured and validated. Whatever the tool, responsibility for the completeness of the record rests with counsel.

AI ASSISTANCE VS. AI DETERMINATION

AI performs various functions within a document review. At one end of the spectrum, its role is purely administrative: deduplication, document classification, language identification, metadata extraction, and batch prioritization. At the other end, the technology evaluates semantic meaning, maps relational networks across a corporate structure, and makes predictive determinations on relevance.

The critical distinction in practice is between AI-assisted review, where technology aids decisions ultimately made by legal counsel, and AI-determined production, where software independently decides that a subset of documents receives no human consideration.

The hypothetical scenario illustrates AI-determined production. No lawyer concluded that those 30,000 documents were irrelevant; the platform reached that conclusion, and counsel adopted it without verification.

The closer a process moves toward automated determination, the greater the obligation on counsel to understand how the system reached its output and how those results were tested. Administrative processes rarely require granular justification. By contrast, automated determinations on relevance and privilege demand rigorous oversight, particularly regarding documents that sit on the margin or whose significance is context-dependent. Read in isolation, an email arranging a routine calendar invite appears irrelevant. However, read alongside a witness's denial of involvement, it may be the central document in the arbitration.

Crucially, human sampling feeds back into the technology. As lawyers review output, they identify gaps and refine search strings, which in turn ensures the emerging case strategy informs the software. Had the legal team reviewed the 30,000 excluded documents and discovered a single exchange involving the witness, they would have had immediate cause to recalibrate the system before the witness statement was finalized and served.

MANAGING THE OMITTED DOCUMENT

In a review of 50,000 documents, errors will inherently occur whatever methodology is utilized, what matters to a tribunal is how the parties account for those errors when they emerge.

When a relevant document surfaces post-production, the immediate obligation is to establish why it was missed. The nature of the document will also dictate the tribunal's reaction.

At that stage, the mechanics of the review become a central procedural dispute. A party that can demonstrate how its platform was configured, how statistical sampling was conducted, and where legal oversight was applied will be well placed to argue that the missed document was an isolated slip within a defensible process. Conversely, a party that cannot explain why 30,000 documents were excluded, or that applied no human auditing to that group, will struggle to convince the tribunal that its production is complete.

Contemporaneous record-keeping is therefore an essential component of the disclosure exercise itself. Legal teams should maintain a clear audit trail documenting:

- The scope and parameters of the harvested data population;
- The search, filtering, and classification methodologies applied;
- The specific role allocated to AI tools versus human reviewers; and
- The sampling methodology used to validate automated exclusions.

A party is not required to volunteer this documentation to the tribunal as a matter of course. It is, however, the precise record required should the adequacy of disclosure be challenged. Reconstructing such an audit trail after a critical document has been exposed is rarely persuasive.

Tribunals do not need to inspect source code or proprietary algorithms; they require sufficient operational detail to satisfy themselves that the review methodology was reasonable, proportionate, and fair.


THE DIFC AND ADGM FRAMEWORKS

Arbitrations seated in the Dubai International Financial Centre (DIFC) or the Abu Dhabi Global Market (ADGM) are governed by arbitration legislation based on the UNCITRAL Model Law. Both jurisdictions afford tribunals wide procedural autonomy, and neither regime contains specific statutory provisions governing AI. Instead, general principles governing the conduct of proceedings and the assessment of evidence apply.

Under Article 24 of the DIFC Arbitration Law (DIFC Law No. 1 of 2008), parties must be treated with equality and given a full opportunity to present their case. Article 25 confers broad power on the tribunal to determine the admissibility, relevance, materiality, and weight of any evidence. The ADGM Arbitration Regulations 2015 reflect the same framework: the tribunal exercises wide discretion over procedure, subject always to mandatory duties of due process and equal treatment.

Arbitrations seated in the DIFC and ADGM reflect a unified legislative philosophy derived from the UNCITRAL Model Law. Under Article 24 of the DIFC Law No. 1 of 2008 (“**DIFC Arbitration Law**”) and Section 33 of the ADGM Arbitration Regulations 2015 (“**ADGM Arbitration Law**”), tribunals operate under a mandate to ensure equal treatment and procedural fairness. This duty is paired with expansive procedural autonomy codified in Article 25 of the DIFC Arbitration Law and Section 34 of the ADGM Arbitration Law which assigns exclusive authority to the tribunal to rule on the admissibility, relevance, and probative weight of all evidence submitted.

Under both legal frameworks, ultimate responsibility for the integrity of the evidentiary record rests with the tribunal and the parties. An automated software platform bears no legal or procedural standing.



Consider the hypothetical scenario within a DIFC-seated arbitration where the tribunal has issued a procedural order for document production. The sheer volume of documents returned by a platform indicates very little about compliance with that order. The key inquiry is whether the chosen methodology was capable of identifying responsive documents within the ordered parameters, and whether the resulting record affords both sides a fair hearing.

If the tribunal concludes that a document ought to have been produced earlier, it possesses a broad spectrum of procedural remedies. It may:

- Direct further targeted searches across the previously excluded population;
- Permit the witness to submit a supplementary statement addressing the document;
- Grant the opposing party leave to adduce responsive evidence;
- Reallocate costs to reflect the procedural delay; or
- Draw adverse inferences when evaluating the witness's credibility.

The issue typically reaches the tribunal in one of two ways: either an opposing party applies for re-review on the basis that a party's filtering methodology was inherently flawed, or a party seeks leave to rely on a helpful document discovered late within its own system. In either scenario, the tribunal's focus will remain on the adequacy of the process and its impact on procedural fairness.

Advancement in AI tools does not necessitate a fundamental rewrite of arbitration law in the DIFC, the ADGM, or internationally. Existing arbitral principles remain fully capable of addressing defective disclosure.

What changes when AI enters document review is not the standard of disclosure, but the point at which a party's internal workflow opens itself to scrutiny. Allowing a platform to discard thousands of documents without human validation elevates that methodology into a substantive evidentiary issue. Should an unproduced document later prove central to the dispute, counsel will find it extraordinarily difficult to justify why no lawyer ever laid eyes on it.

The answer is not to abandon these tools but to use them with discipline. Relying on automated filtering without building clear audit trails or sampling excluded datasets is a gamble that assumes the system will never fail. When a critical document is missed, blind trust in an algorithm offers no protection before a tribunal. As these tools become increasingly autonomous, counsel must actively guard against complacency. Ensuring proper oversight, documenting methodologies as they are executed, and proactively interrogating the system's blind spots are essential safeguards for protecting the integrity of the case.

Sarah Malik FCIArb

CEO and International Disputes Counsel

[www.linkedin.com](https://www.linkedin.com/in/sarahmalik)

Sarah Malik is an award-winning lawyer who achieved 'Litigator of the Year' at the first GCC Women in Law Awards 2022 and received an honorable mention as 'Law Firm Leader of the Year' at the same awards. SOL was set up in 2018 and has become internationally recognized brand of repute in all aspects of dispute resolution. In 2022, SOL was listed as a 'Firm to Watch in Dispute Resolution: Arbitration and International Litigation' by Legal 500 EMEA. In 2023, Sarah was recognized in the Arbitration Powerlist in Private Practice in the Middle East by Legal 500 EMEA.



Sarah is a published author of two legal texts. She is also appointed as Adjunct Faculty in Law teaching LLM and LLB students in Dubai, LLB students in Riyadh and Fashion Law students in Milan. Sarah acts as Counsel in multi-jurisdictional disputes and is a sought after international public speaker on thought leadership. Sarah holds various international advisory and Board positions.

Sarah is also retained as a personal advisor to high net-worth clients, acts as external GC, and undertakes sports, fashion, and entertainment law. She runs multiple businesses in addition to SOL and is recognized for her business acumen.

Pavithra Rajendran

Trainee

[https://in.linkedin.com/in](https://in.linkedin.com/in/pavithra)

Pavithra Rajendran joined SOL International Ltd in November 2024 and is currently a Trainee Solicitor, focusing on dispute resolution, employment and commercial law. Her responsibilities include drafting court submissions and pleadings, preparing bundles, and supporting clients in litigation from inception to conclusion.

Pavithra holds an LL.M in International Business Law from Middlesex University and a B.A., LL.B (Hons) from the School of Law, SASTRA University.

Pavithra has been involved in a number of matters, including shipping arbitration administered by the Hong Kong International Arbitration Centre (HKIAC), litigation involving hospitality claims before the Dubai International Financial Centre (DIFC) Courts, and claims before the Abu Dhabi Global Market (ADGM) Courts. She was also previously involved in a multi-million-dollar claim before the DIFC Courts concerning the enforcement of a SIAC award, which raised complex cross-border issues relating to service and jurisdiction.

She has experience drafting Letters Before Action contributing to the amicable resolution of disputes. Her work also extends to corporate matters, including company re-domiciliation, and drafting and reviewing Share Purchase Agreements, Distribution Agreements, Service Agreements, Consultancy Agreements, Subcontracting Agreements. She has liaised with regulatory authorities to obtain permits and licenses required for business operations.

