

AI and the Lawyer: Learning to Search for What Could Be Wrong

Adv. Joseph Weinrauch, Attorney-at-Law, Tel Aviv

June 2026

Artificial Intelligence has quietly made itself at home in legal practice. For most lawyers, entry was gradual — an email draft here, a document summary there. Then contracts. Then research memos. Then pleadings. Before we registered the change, AI had become part of the workflow.

In many ways, it resembles an exceptionally efficient junior lawyer: it drafts without fatigue, summarises without complaint, suggests case law at speed, and is available at three in the morning. So, what is the problem with this new assistant? Quite similar to the assistant of the past: accuracy — but with far greater confidence. When I was a young legal assistant, preparing documents for my principal, I was afraid to fabricate. I might cut corners, yes — but only at the margins. The core had to be real. Fear was the safeguard. Fear of being caught, fear of embarrassment, fear of professional consequences. That fear was not a weakness — it was a professional filter. AI has no such fear. No shame. No hesitation. It does not know what it does not know, and it does not care. It will invent a Supreme Court decision — give it a name, a date, a docket number, a holding — and present it to you in clean, confident prose, indistinguishable in tone from a citation that actually exists. It will do this without a tremor. And then it will wait for your approval.

The Problem Has a Name

AI hallucinations occur when a system generates information that appears entirely credible — well-written, structured, plausible — but is factually wrong, fabricated, or unsupported. In legal practice, this is not a theoretical concern. It is a disciplinary one. In June 2024, when my colleague Pim Betist and I presented together in Amsterdam on the use of AI in cross-border legal work, we posed this paradox as the central challenge:

"If you use it, there is trouble. If you don't, it will be double trouble." — Pim

Betist, Amsterdam Accord, June 2024

At that time, a Stanford University study had just found that AI hallucinated in **51% of legal research responses** — more than half. That figure shocked audiences. It was the number that changed the conversation. And yet the cases that made the headlines were still ahead of us.

The Israeli Reckoning

On 20 February 2025, Israel's Supreme Court issued a landmark ruling that legal practitioners across the country could not ignore. Justice Gila Canfy-Steinitz addressed a petition in a divorce matter under Sharia law in which the petitioner's counsel had submitted AI-generated pleadings containing:

ISRAELI SUPREME COURT — FEBRUARY 2025

- **36 cited Supreme Court decisions** — internally inconsistent
- References to **non-existent judgments**
- Quotes that appeared in no legal database

The Court found that the lawyer had relied blindly on an unnamed website "recommended by colleagues" without any independent verification. Justice Canfy-Steinitz wrote that the AI hallucination output had appeared so credible that the lawyer did not bother to check its veracity. The Court did not sanction the lawyer — noting this was the first such case before it — but warned explicitly that future cases would face the full range of judicial tools. That warning proved prophetic almost immediately. **Days later**, another petition reached the Supreme Court — this time on a law facilitating stray dog euthanasia — again containing hallucinated cases. The Court rejected it outright. **Within weeks**, a debtor filed an application citing non-existent cases with the Enforcement and Collection Authority. His apologetic follow-up pleadings, remarkably, included further AI-hallucinated citations. He was fined 1,000 NIS. In May 2025, the Tel Aviv District Court struck down a class action request against health funds for the same reason.

Israel was not alone. A researcher tracking AI hallucination cases in court filings globally has now identified over **1,600 documented cases** since 2023, with the overwhelming majority occurring in 2025 alone. In the United States, lawyers have been sanctioned, suspended, and fined. An Alabama attorney was sanctioned for filing briefs with fabricated citations — and then cited a non-existent case again in the very next sentence after being corrected.

Sources: Library of Congress, June 2025; Lexology, March–April 2025; [Damien Charlotin, AI Hallucination Cases Database](#), 2025; Scientific American, 2025.

What the Research Now Shows

The data between 2024 and 2026 tells a story of meaningful improvement that still falls far short of professional reliability.

The 2024 Stanford study — "Large Legal Fictions" — tested general-purpose models on over 800,000 verifiable legal questions. Hallucination rates ranged from **58% for GPT-4** to **88% for older models**. A second Stanford study, published in 2025, then examined the premium specialised legal tools — Lexis+ AI, Westlaw AI-Assisted Research, and Thomson Reuters's Ask Practical Law AI — the tools that lawyers actually pay professional rates for:

Stanford / Journal of Empirical Legal Studies, 2025:

Lexis+ AI: **17% hallucination rate**

Westlaw AI-Assisted Research: **33% hallucination rate**

GPT-4 (baseline): **43% hallucination rate**

Conclusion: *"Legal RAG can reduce hallucinations compared to general-purpose AI, but hallucinations remain substantial, wide-ranging, and potentially insidious."*

Source: Magesh et al., "Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools," Journal of Empirical Legal Studies, 2025.

To place these figures in context: across all domains, the overall hallucination rate for AI even among top models stands at approximately **6.4% for legal information** — compared to just **0.8% for general knowledge questions**. Law is not a domain where AI underperforms slightly. It underperforms dramatically and specifically, because legal reasoning requires contextual synthesis of authority that language models cannot reliably perform.

Hallucination rates by field — top-performing models (2024–2025):

General knowledge: **0.8%**

Financial data: **2.1%**

Medical / healthcare: **4.3%**

Legal information: **6.4%**

Source: AI Hallucination Report 2025–2026, AllAboutAI / drainpipe.io compilation.

Improvement is real. The best current models hallucinate less than models from two years ago — some by as much as 64% less. But an 85% accuracy rate in legal citation is not a passing grade. It is a disciplinary risk.

A Different Set of Muscles

The deeper professional challenge is not technical. It is cognitive — and it is more tiring than it first appears.

For most of my career, legal work followed a familiar rhythm. I accumulated expertise through thousands of documents, hundreds of transactions. When a new matter arrived, I recalled a precedent, retrieved a similar document, adapted it. The effort was creative but grounded: I was working with material I had already internalised, adjusting for what was new.

AI inverts this entirely. It delivers a well-organised draft almost instantly. The document looks good. The citations appear correct. The reasoning flows. And this is precisely where the danger begins.

The skill the lawyer now needs is not retrieval — it is detection. Instead of searching memory for the right precedent, the lawyer must interrogate the AI's output systematically:

- Do these cited cases actually exist in the database?
- Do the holdings say what the AI claims they say?
- Is the statutory reference accurate to the current version of the law?
- Has the AI presented a selective picture — omitting contrary authority?
- Does the reasoning connect to the actual facts of this case, or is it generic?

This is not the same mental effort as drafting. It is closer to auditing. And it has a particular quality that makes it genuinely demanding: you are searching for errors in output that **looks correct**. The AI does not flag its mistakes. It presents a fabricated citation with the same confident prose as a verified one. The lawyer must supply the skepticism that the tool entirely lacks.

This shift in cognitive effort is real, and it can be tiring precisely because it cuts against professional instinct. Experienced lawyers trust their reading of a document. They have earned that trust over decades. With AI output, that instinct must be consciously suspended and replaced with systematic verification — every time, on every document, regardless of how credible the draft looks.

One can identify errors only if one possesses a strong foundation of legal knowledge. AI does not make legal expertise less necessary. It makes continuous learning more necessary than ever.

Should Lawyers Use AI?

Yes. Used with appropriate discipline, AI saves genuine time. A faster first draft means more time for analysis. A summarised set of documents means more time for strategy. The efficiency benefits are real, and in many cases they translate into lower costs for clients.

But the word *discipline* is carrying significant weight in that sentence. Responsible use means:

- Treating every AI draft as a starting point, never a finished product
- Verifying every cited authority independently against official databases
- Never submitting AI-generated pleadings without reading them as carefully as if you had drafted them yourself — because before the court, you did
- Investing time in prompt quality: precise, contextualised instructions produce significantly better output and reduce the verification burden

Perhaps the clearest analogy is this: AI is an extraordinarily efficient law librarian — one that retrieves information in seconds. The responsibility for understanding, verifying, and applying that information remains entirely with the lawyer. The law librarian cannot be sanctioned. You can.

The Present Picture

Two years after Amsterdam, the credibility of AI in legal practice has improved — measurably, meaningfully — but not enough to work unsupervised. Hallucination rates in specialised legal tools have dropped from the mid-fifties to the mid-to-upper teens. That is progress. It is not yet safety.

The courts have drawn the line with clarity. Bar associations from California to New York to Tel Aviv have issued the same message: AI does not reduce professional responsibility. It creates a new layer of it.

The lawyer of the coming decade will need to master two distinct competencies that have rarely needed to coexist: deep legal knowledge, and rigorous skepticism toward the tools that claim to help apply it. That combination — expertise and structured doubt — is what responsible AI-assisted practice looks like in June 2026.

During one of our working sessions on this article, I shared a draft direction with Claude. Its response was: *"Love this direction — personal, authentic, and sharp. Let me complete it in your voice."*

Coming from an AI — that made me smile.

Sources: Stanford RegLab / Institute for Human-Centered AI (January 2024); Magesh et al., "Hallucination-Free?", *Journal of Empirical Legal Studies* (2025); Library of Congress Global Legal Monitor (June 2025); Lexology (March–April 2025); Damien Charlotin, AI Hallucination Cases Database (2025); *AI Hallucination Report 2025–2026* (AllAboutAI); *Scientific American* (October 2025). Amsterdam Accord presentation: Adv. Joseph Weinrauch & Pim Betist, June 7, 2024.

This article was prepared with the human assistance of Ms. Preeti Desi, and with Claude AI.