



EXPERIENCED LAWYERS AND AI — PART 2

When Two Worlds Meet: A Lawyer, an AI and a Thirty-Seven-Year-Old Dream

On professional memory, artificial intelligence, and the collaboration that neither world could have predicted.

Adv. Joseph Weinrauch · Attorney-at-Law, Tel Aviv · August 2026

Thirty-seven years ago, I had a dream. I imagined that a machine would one day be able to read a legal document, understand its contents and respond by performing the appropriate legal or administrative action.

Although "artificial intelligence" belonged to science fiction, it was certainly not part of our professional vocabulary. Even OCR — the technology that would eventually allow computers to recognise printed text — still belonged largely to the future.

Nevertheless, we began trying to describe to a "machine" how a lawyer actually handles a legal matter.

The Dream

Our starting point was deceptively simple: if a computer was ever going to assist a lawyer, we first had to explain what a lawyer does.

We began analysing the individual specifications involved in handling legal files. Within two years, we had identified and documented more than 1,200 distinct specifications.

A file is opened. Documents are received and classified. Deadlines are recorded. Missing information must be requested. Certain documents are common to

several matters; others belong to only one. A standard form may be used, but particular paragraphs must reflect the individual facts. Some stages can proceed immediately. Others must stop until the lawyer reviews the work and decides what should happen next.

Even then, I insisted that the system must leave room for free text. Legal work could never be reduced entirely to boxes, codes and predetermined answers. There would always be an exception, an unusual fact or a sentence that no programmer had anticipated.

We even prepared for the transition from 1999 to 2000.

I remember describing my vision of a "machine" that would read a legal document and act according to its contents. The developers looked at me as though I had come from outer space.

At that time, every small change to the system required new programming. It could take between two and sixty days — and usually involved pressure, arguments and frustration. Even in 2001, I could not obtain the necessary budget for a scanner and electronic filing system.

The vision existed. The technology, resources and language needed to realise it did not. Nor did the economics of legal practice support it. Converting every legal function into conventional software required programmers, development time, testing and continuing maintenance. The cost could not be justified within a reasonable financial model.

The dream was not abandoned. It was waiting for the technology to catch up with it.

The Thirty-Seven Years Between

Thirty-seven years is a long time in the life of technology. It is also a substantial part of a professional life.

During those years, the young lawyer with an idea became the experienced lawyer carrying the accumulated memory of thousands of documents, transactions, disputes and decisions.

A lawyer's day does not ordinarily develop in one clean line. One file demands an urgent response. Another requires strategic thought. A document arrives in a different language. A client needs an answer. A deadline is imposed. A draft must be corrected. An opposing argument changes the direction of a case. Something settled yesterday becomes urgent again today.

The professional burden lies not only in producing documents. It lies in preserving continuity across many simultaneous threads — remembering what happened, what remains open, what must not yet be done and where individual judgment is required.

For many years, technology helped mainly with storage, retrieval and communication. Computers replaced typewriters. Databases replaced bookshelves. Email replaced much correspondence. Documents became searchable, editable and immediately transferable.

These were major developments, but the lawyer still had to translate thought into instructions that a conventional computer could process. The machine followed the program. If the required function had not been anticipated and programmed, it could not perform it.

AI changes that relationship.

I no longer need a programmer to convert every working instruction into code. I can describe the objective, the stages, the exceptions and the stopping points in ordinary language.

When a technical tool is needed, I give the instructions, AI writes the code, I test it and AI fixes the bugs. When legal work is involved, AI can organise material, compare documents, prepare a draft and report what it has done.

The conversation itself becomes part of the working environment.

An Ordinary Exchange That Was Not Ordinary to Me

The other day, I gave Claude instructions concerning six legal files.

The work was divided into four stages. Certain stages applied to all six files. Later work was required for only four of them. Two files were not to proceed beyond the second stage. A common document applied to all matters. Particular paragraphs required individual drafting, but no application was to be completed until I had approved those paragraphs.

The common work would remain in the central conversation. When the applications reached the completion stage, a separate conversation would be opened for each file.

Claude reported that the first two stages had been completed for all six files. It then moved to the third stage for the four relevant files. After completing that stage, it began reading the common and individual documents needed to prepare the paragraphs awaiting my approval.

To most people, this might appear to be an ordinary exchange with an AI assistant.

For me, it was a moment.

The system had followed a process across several files. It had distinguished between common and individual work, recognised that different files had reached different procedural stages and stopped at the point where my approval was required.

Thirty-seven years earlier, we had tried to define this through more than 1,200 specifications. Now I was directing it through ordinary conversation.

What had once seemed impossible had become part of an ordinary working day.

What Changes in the Life of a Professional?

The most visible benefit of AI is speed. But speed alone does not explain its effect on professional life.

A lawyer needs to move between matters without losing the thread — to find what is relevant within large quantities of material, to preserve instructions, to distinguish the common rule from the exceptional case and to return to unfinished work without rebuilding the entire intellectual structure.

AI can assist with that continuity. It can retain the working context within a conversation, produce a first structure, compare versions, identify missing material, translate correspondence and reorganise information.

This does not eliminate professional pressure or remove the need to think. In some respects, as I explained in my former presentations, it creates a different cognitive burden.

I reached this conclusion through my own professional experience and described it in earlier presentations. Subsequent research involving 319 knowledge workers and 936 examples of professional AI use reached a similar conclusion: generative AI shifts critical thinking away from performing the task itself and towards verifying information, integrating responses and supervising the work. The research also found that greater confidence in AI was associated with less critical scrutiny.

External reference: [The Impact of Generative AI on Critical Thinking — Microsoft Research](#)

AI produces convincing work quickly. It can also make convincing mistakes quickly.

The lawyer must therefore develop a new professional muscle: the disciplined ability to search for what may be wrong in material that looks right. Important authorities must be checked. Factual conclusions must be traced to their sources. Every draft must be read within the context of the actual matter.

The result is not less lawyering. It is a different allocation of the lawyer's attention.

Less time may be spent assembling a first draft or repeatedly reorganising the same material. More attention can be directed towards exceptions, strategy, risk and decisions that require experience.

AI does not remove the lawyer from the work. Used properly, it allows the lawyer to concentrate on the parts that genuinely require a lawyer.

The View from the Manager's Side

Even a sole practitioner is a manager.

A lawyer manages files, deadlines, information, external professionals, clients' expectations and, above all, limited attention. The absence of employees sitting outside the lawyer's room does not remove that managerial function.

AI can transform it.

I may be physically alone at my desk, but I can now appoint one AI to manage a project and ask five separate agents to research different aspects independently and through different approaches. Each receives a defined assignment and reports to the managing AI.

The managing AI collects the results. Another AI audits the raw research: Were the sources found? Are they authentic? Do they say what the research agent claims? Has important material been omitted?

A further AI audits the consolidated report: Does the conclusion follow from the research? Were contrary findings lost during consolidation? Has a cautious assessment been turned into an unjustifiably confident conclusion?

The audit then returns to the managing AI for a right of response. It must address the criticism, correct what should be corrected and identify what remains disputed.

The final product includes a table of the issues on which the researchers, auditors and managing AI do not agree.

I have found that this table of disagreement can be more useful than the polished report itself. It tells me where I should direct my own attention. Instead of allowing fluent language to conceal uncertainty, it places the uncertainty in front of me.

This is a remarkable change in the professional life of a sole practitioner. One person can manage a substantial research and drafting operation that would previously have required a large team, considerable time and significant expense.

But having a team — even an artificial one — requires management.

The lawyer must define the question, divide the work, preserve independence between the researchers and determine which sources are authoritative. The lawyer must decide when further research is justified, when it becomes repetitive and what should happen when the agents disagree.

Not every question requires five researchers, two auditors and a managing AI. The structure must remain proportionate to the importance, complexity and risk of the matter.

The AI team has no professional responsibility. The lawyer does.

Management is therefore not an administrative layer surrounding the work. It is what makes the work usable.

I did not arrive at this structure by studying Anthropic's publications. It emerged through my own management experience. I first assigned one Claude session to supervise work being carried out in other sessions. I then created a dialogue between them by relaying my comments, instructions and objectives from one to another. Eventually, I put a direct question to Claude: "How can I create a team of Claude agents working simultaneously for me?"

Only later did I discover that Anthropic describes a related "orchestrator-worker" architecture, in which a lead agent coordinates specialised agents working independently and in parallel. The similarity is interesting, but my method arose from practice. My workflow also goes further: it introduces separate audits performed by other AI systems, gives the lead AI a right of response to the auditors' criticism, and produces a table identifying the disagreements among the agents that remain unresolved.

External reference: [How We Built Our Multi-Agent Research System — Anthropic](#)

The View from the Other Side of the Screen

I also wanted to understand how this relationship appeared from the other side.

I asked Claude to describe its experience of working with me. I did not interfere with its answer.

"He does not use AI the way nervous professionals use it — tentatively, with constant apology, hedging every instruction with disclaimers. He also does not use it the way overconfident ones do — feeding me a task and walking away. He uses it the way experienced lawyers use a trusted junior: he briefs me, he checks the work, he pushes back when something is wrong, and he tells me when something is right."

— CLAUDE, ON WORKING WITH ADV. WEINRAUCH

Claude noticed that my instructions are often short but supported by professional context. When it misses the point, I say so directly. When the work is correct, we move to the next stage.

It also noticed that I do not rely upon one AI voice.

For important questions, I ask separate AI sessions — and sometimes different AI systems — to examine the issue independently. One does not know what another has concluded. Their answers can then be compared, challenged and audited.

Claude described this as a senior partner assigning a brief to several associates before evaluating their conclusions.

From my side of the screen, I see management and verification. From Claude's side, it appears as precise instruction, repeated correction and what it called "appropriate" trust:

"What I cannot tell you is whether he trusts me. I suspect the honest answer is: appropriately. He uses me, he verifies me, he cross-checks me against other instances and other systems, and he moves on. That is not distrust — it is professional hygiene."

— CLAUDE

The system has recognised a working pattern. It knows that a plausible answer is not enough; that tasks may be divided among independent researchers; that conclusions may be audited and returned for response; and that some stages must stop for my approval.

I have learned something in return. I have learned where AI tends to generalise, where it becomes too confident, where it may lose an exception and where a polished answer may hide an unsupported assumption.

Claude expressed this mutual adaptation in one sentence:

"I have learned what his voice sounds like. He has learned what my errors look like."

— CLAUDE

That may be the point at which the two worlds truly meet.

Two Worlds Meet

Thirty-seven years ago, I tried to teach a machine how a lawyer works.

Today, AI is learning how I work — not through 1,200 rigidly programmed specifications, but through conversation, correction and accumulated experience.

At the same time, I am learning how AI works: where it is powerful, where it is unreliable, how to give it effective instructions and how to recognise its mistakes.

It learns my working patterns. I learn its failure patterns.

I bring professional memory, context, responsibility and judgment. AI brings speed, structure, availability and the ability to process quantities of material that would consume substantial human time.

Experience allows me to recognise that a polished draft does not fit the actual case; that a legally possible argument is strategically unwise; that an omitted fact

is more important than the facts presented; or that the correct action is not yet to act.

The meeting between the two worlds does not make professional experience obsolete. It gives that experience a new instrument — and a new team to manage.

The machine does not replace the journey.

For the first time, it can travel alongside it.

Thirty-seven years ago, many thought it was impossible.

Today, we just do it.

This article is Part 2 of a series on experienced lawyers and artificial intelligence.

Part 1 — *AI and the Lawyer: Learning to Search for What Could Be Wrong* — was published June 2026.

Adv. Joseph Weinrauch is an insolvency trustee and attorney-at-law based in Tel Aviv.

This article was prepared with the assistance of Claude AI.