



Litigation at one-tenth the cost, or ten times the output.

Rebuild litigation support as an AI-enabled operating system, run with manufacturing discipline instead of billable improvisation. Measured on real matters, with an error record that beats the leading legal-AI tools.

378K+

documents indexed into 5.5M+ searchable passages, on local hardware

0

fabricated citations or invented legal authority ever reached a filing

~8 hrs

zero to a filing-ready employment package, exhibits included

THE PRETRIAL PARADOX

Trials are rare. Cases are won or lost in the pretrial phase: building the record, meeting deadlines, drafting discovery, and framing the motions that decide whether a case survives at all. In one federal district, **56% of self-represented claims did not survive even a motion to dismiss** (study: U. Chicago Law Review). Pretrial is where cost concentrates and where the fatal errors live: a missed statute of limitations, a blown response window, an unanswered request for admission that quietly becomes an admission.

The question this paper answers: what happens when you run the pretrial phase as a machine? Measured on real matters, the same work gets done at roughly **one-tenth the cost, or ten times the output, and often both**, with a documented error record that beats both the leading legal-AI tools and human baselines.

Not a lawyer. That is the whole point.

This practice is led by an operator, not a vendor and not an attorney. In college he spent two years as a document-coding QA lead and coding manager at a litigation-support firm, running quality control on the coded record. He later ran his own litigation-support firm. Then, over the course of a year, he built the platform this paper describes, and he tested it in the hardest place in all of law.

A business cannot represent itself in court; a corporation must appear through counsel, so in commercial litigation the only pro se party is a lone individual. They lose almost always: pro se plaintiffs win about 3% of final judgments in federal court, and they are a small minority even of higher-value commercial cases who almost never litigate one to a conclusion. Across federal diversity commercial cases from 2021 to 2026 (cases that by law put more than \$75,000 at stake), a self-represented party reached a pretrial conference in only about 1 in 7 of their cases, and trial in roughly 1 in 75 (analysis of the Federal Judicial Center Integrated Database). The empirical literature treats an affirmative counterclaim as a marker a case has moved beyond pro se viability.

Against that, one non-lawyer used this platform to carry multiple complex, multi-defendant business matters through pleadings, counterclaims, and into discovery, pro se. Pro se litigants lose on procedure, missed deadlines and untimely responses, not the merits, and this platform is a procedural-compliance machine. Read that again, then picture the same machine in the hands of your firm. The point was never the litigant. It is the machine. Licensed attorneys own the legal advice and every filing; nothing here is the practice of law.

THE THESIS

AI is the operating system, not the tool.

The industry's own data says the value is not in any single tool. In a 2026 litigation-support survey, **integration depth predicted user satisfaction four times better than any individual feature** (study: US Legal Support). The market has moved past general-purpose chatbots toward platforms that run a whole case, because generic models are built to sound right, not to be verifiably right. Legal work demands the opposite: a citation you can check, a source you can trace, a record you can defend.

That is the Synaptic thesis applied to law. AI becomes the operating system underneath the lawyer: it ingests the record, maps every document to the legal elements it supports, computes every deadline with its rule citation, drafts the discovery, tracks the admissions, and makes filing errors structurally hard. **AI drafts; you file. AI suggests; you decide.**

The repeatable methodology behind every matter.

Every matter runs the same gated pipeline. It is named for what it produces:

-
- I Ingest.** Docling parses every file type on local hardware. 378,000+ documents become 5.5 million+ passages, embedded with nomic-embed-text at 768 dimensions into a self-hosted PostgreSQL and pgvector store with an HNSW index. Nothing leaves the box.
-
- M Map.** Local vector search, at zero cost per query, ties each passage to the legal element it proves, on intake.
-
- P Pressure-test.** Local search over the whole record, with a frontier model (Claude) doing the legal reasoning, runs a gap analysis that finds where the record is thin and sets the next round of discovery.
-
- A Adversarially review.** Three separate gates before anything is filed, each catching what the others miss: internal QA rules, an independent outside-model red-team, and a human read.
-
- C Cite to source.** Every assertion traces to a Bates-stamped page, a statute, or a prior filing, and each citation is checked against the source.
-
- T Track.** A rules engine computes every deadline with the rule that generated it and reconciles the docket.

One example of what that produces. Point it at a fact and the system finds the source document by meaning, pulls the original from wherever it lives, converts it to PDF and OCRs it if it is not already searchable, renames it to the court's caption and exhibit nomenclature, and files it into place, as a standalone exhibit or appended to a combined exhibit PDF, formatted to that court's rules. The pleading comes out assembled and ready to submit. Where a court's electronic filing system allows it, the filing itself is automated.

The method is enforced in code, not by memory. Each defect ever discovered became a permanent standing check: 30-plus blocking QA categories and a set of code guards, each born from a specific documented failure.

This is not a login you can buy. Ingestion, search, and the document store are self-hosted (Docling, Ollama, pgvector, PostgreSQL), tuned for a year against live dockets and run on hardware you control, so searching the whole corpus costs nothing per query and no case data leaves the building. The legal reasoning and drafting run on a frontier model (Claude), which does cost tokens, but because retrieval is local you pay only to synthesize the handful of passages that matter, not to search hundreds of

thousands of documents. That keeps token cost a fraction of a naive approach. Your firm cannot download this, and a general-purpose chatbot cannot approximate it. That is why we install it and run it for you.

THE ERROR RECORD

The honest version, built to be verified.

A fabricated citation reaching a judge is the error that draws sanctions and ends cases. Here is the record, stated the way it would have to survive cross-examination.

~34%

Westlaw AI hallucination rate
(Stanford study)

~17%

LexisNexis AI hallucination
rate (Stanford study)

1,000+

U.S. court filings caught citing
fabricated AI cases, and
climbing

Sources: Stanford RegLab and HAI, Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools (2024): Westlaw's AI hallucinated more than 34% of the time, Lexis+ AI more than 17% (hai.stanford.edu). The fabricated-citation count is Damien Charlotin's AI Hallucination Cases Database, which as of mid-2026 logged more than 1,000 U.S. court decisions (about 1,500 worldwide), growing by several a day: damiencharlotin.com/hallucinations.

THIS PIPELINE'S RECORD: FROM THE FILING LEDGER AND ROOT-CAUSE LOG

Across roughly 40 accepted filings on two active state-court dockets:

- **No fabricated case, invented statute, or made-up legal authority has ever reached a filed document. Zero.** That is the error the profession is measured on, and it is the one we hold.
- Two errors did reach the docket, and we own them: one motion filed under the wrong matter's caption, and one citation corrected after filing. **Both were fixed by formal errata and accepted by the court.** Errata are routine for every firm; the difference is ours are logged with a root cause and engineered against, so the same defect cannot recur.
- Every draft clears an **adversarial-review gate** before it can be filed: internal QA rules, an independent external-model red-team, and a human read. That gate, not a claim that the machine never errs, is why the record above holds.

How this is derived: the denominator is the filing ledger's count of accepted envelopes; the errors are counted from the root-cause log the system keeps for every defect it finds. We do not claim zero mistakes. We claim two auditable things: no fabricated authority has reached a filing, and every error that did reach the docket was corrected on the record. In the operator's own words to colleagues and opposing counsel, the residual error rate is down to typos, not wrong cases and not invented law.

Against real lawyers and paralegals, the fair comparison is not who never errs, because everyone does. It is the catastrophic error: fabricated authority reaching a judge. There the profession has produced more than a thousand documented failures in the U.S. alone, and counting. This pipeline has produced none.

THE FIVE THINGS LAWYERS SAY ABOUT LEGAL AI

And how this answers each.

1

"It hallucinates and I can't trust the citations."

Answer. The adversarial-review gate plus source-parity checking. Every citation traces back to the exact source document, Bates-stamped if you want it and tracked either way, down to the exhibit number, and the record above shows zero fabricated citations reaching a filing.

2

"It's generic, not built for how I actually practice."

Answer. It is built to your practice, not the other way around. We shape the pipeline and the workflow to how you actually litigate, encode the local rules and conventions of your courts, and add or change functionality on request, immediately. No SaaS roadmap, no feature committee, no waiting a quarter for a vendor to decide whether to build it. If it should work differently, it does, that day.

3

"It's a shelf of disconnected tools, and the gains get lost across tabs."

Answer. One command center. Ingestion, evidence mapping, discovery, deadlines, exhibits, and filing live in a single system: exactly the integration depth the research says predicts satisfaction four times better than any feature.

4

"The pricing is opaque and it nickel-and-dimes per feature."

Answer. A transparent operating-system install, not a per-click meter. The heavy lifting, embedding and searching hundreds of thousands of documents, runs locally on your hardware at no per-query cost; you pay tokens only for the final drafting a frontier model does over the few passages that matter, a fraction of the cost of pushing everything through a paid tool.

5

"There's no defensible record when I need to prove something."

Answer. Every production is Bates-stamped with a full chain-of-custody audit trail, every assertion traces to source, and the code gates leave an auditable trail of what was checked.

Zero to a filing-ready package in about 8 working hours.

Anonymized. An employment matter; no parties named. We ran our own method on it, IMPACT, end to end. We ate our own cooking.

Here is IMPACT on one real matter, start to finish, from no pleading at all. **Ingest:** more than 7,800 emails pulled in and made semantically searchable. **Map:** each relevant message tied to the legal element it proved, on intake. **Pressure-test:** gap analysis showed what the record could and could not support, which shaped a four-count petition, and flagged one exhibit as an impeachment risk, so it was pulled, the process protecting the case, not just producing volume. **Adversarially review:** the draft cleared the review gates before it was called final. **Cite to source:** seven exhibits stamped, organized, and tied to the assertions they back. **Track:** a discovery plan with a custodian matrix and every deadline computed. The result: a single filing-ready **108-page package** plus a served preservation demand, in about **eight working hours**.

The ROI, estimated conservatively. Reviewing 7,800 emails at a brisk ~100 per hour is roughly 78 hours of review alone. Add the four-count petition, the discovery plan, seven exhibits, and a preservation demand, and a firm is realistically at **100 to 150 professional hours** at a blended rate near \$250, or **\$25,000 to \$37,500**, over two to four calendar weeks. The system delivered it in about eight hours. That is a **12x to 19x time compression**, and valued at the practice's own \$300 hourly rate (about \$2,400 of operator time), a **roughly 10x to 15x cost advantage on a single matter**, before counting the value of never missing a deadline or filing a fabricated citation.

BEYOND THE COURTROOM · THE CONSUMER-PROTECTION EXTENSION

The same system runs pre-suit consumer campaigns.

The operating system is not limited to court dockets. The same pipeline can run a complete consumer-protection pre-suit campaign end to end: a **DTPA statutory pre-suit notice**, a formal **Demand for Cure**, **state Attorney General and licensing-agency complaints**, and **executive demand letters**, all generated, page-cap-validated, and mailed certified through a certified-mail API, by a single operator, in minutes.

The problem it solves: the pre-suit phase normally requires a lawyer or paralegal to draft, format, and certify-mail each piece separately, which is why most consumers never send them. Collapsing that to one operator and a few minutes is the seed of a productized consumer service, where claimants file online and the system runs the full pre-suit campaign, now in development.

THE OFFER

AI Litigation Support Consulting.

Private consulting

A **free 30-minute consultation**, then **\$300 per hour**: the rate attorneys already pay for expert coaching. It stays private: every engagement comes with a simple one-page mutual NDA, I never reveal a client's name (including for marketing), and references are written opt-in only.

The Free-Case Challenge: risk reversal

Give us one live case to ingest and build out. If the system does not beat your three best paralegals combined on speed, accuracy, strategic quality, and overall work product, **you owe nothing**.

Engagements

A Litigation Operating System Assessment, a 90-day pilot on a live matter set, or a fractional litigation-operations lead for a plaintiff or litigation-support practice.

Book the free 30-minute consultation: ivan.madrigo@synapticssystems.ai

The offer is deliberately low-friction and evidence-first. The free consult and the free-case guarantee put the risk on the system, not on you: you see the work product before you commit to anything, and if it does not beat your team, you owe nothing.

THE MATH, AND THE CLOSE

Pick the lever your organization needs.

One-tenth the cost: pretrial labor that scaled with document volume collapses to review and judgment; one operator carries what took a team. **Ten times the output:** freed from mechanical work and never losing a deadline, the same operator produces far more finished work product across more matters at once. **Or both:** the system removes the labor cost and multiplies the throughput simultaneously.

And measure the gain the way your controller would. The number that counts is the one your finance function can verify on your own books: your realized rates against your recorded hours, not an outside estimate. We install the method; your billing proves the saving.

The pretrial phase is where cases are won, where costs concentrate, and where the fatal mistakes are made. It is the right place to install an operating system. This paper is the argument, the error record is the proof, and the free-case challenge is the invitation.

So the choice is simple. Wait for the legal-AI tool you already pay for, Westlaw, LexisNexis, Harvey, to put the function you need on its roadmap and clear its committee. Or be in complete control now: more effective, higher quality, and not carrying the hallucination risk those same tools do, on the studies' own numbers. One of those is a subscription. The other is an operating system.

And a system needs an operator. What moves the numbers is not another login but a partner: someone local who installs the system, leads it, and shows your team how to run it. That is what turns the technology into the promise this paper opened with, litigation at one-tenth the cost, or ten times the output, or both.

Synaptic Systems installs AI-enabled operating systems that increase throughput, reduce cost, and enforce discipline. For a free 30-minute litigation-support consultation, contact ivan.madrigo@synapticssystems.ai.

Draft v9, 2026-07-18. Pro se progression figures are an analysis of the Federal Judicial Center Integrated Database (civil terminations 2021 to 2026, diversity-jurisdiction cases in commercial nature-of-suit categories), reproducible from the FJC's public data. The fabricated-AI-citation count is Damien Charlotin's AI Hallucination Cases Database. All case detail is aggregated and anonymized. "System" figures come from the reference system's own logs and filing registry; "study" figures are cited public research. Roadmap note: a productized consumer-protection service, where claimants file online and the system handles the full pre-suit campaign, is in development.