



Cross-Examine Your AI: The Lawyer's Cure for Hallucinations

“AI walks into your office like a consulting expert who works fast, inexpensively, and speaks with knowing confidence. And, like any untested expert, is capable of being spectacularly wrong.”

- Ralph Losey

Lawyers Are Discovering AI Hallucinations the Hard Way

Courts are sanctioning attorneys who accept AI's answers at face value.

Case Example 1: *Mata v. Avianca, Inc.*

A lawyer submitted a brief filled with invented cases. The judge blamed the lawyer, not the machine.

Case Example 2: *Park v. Kim*

The Second Circuit confronted AI-generated citations that dissolved under scrutiny, leading to the case being dismissed.

700



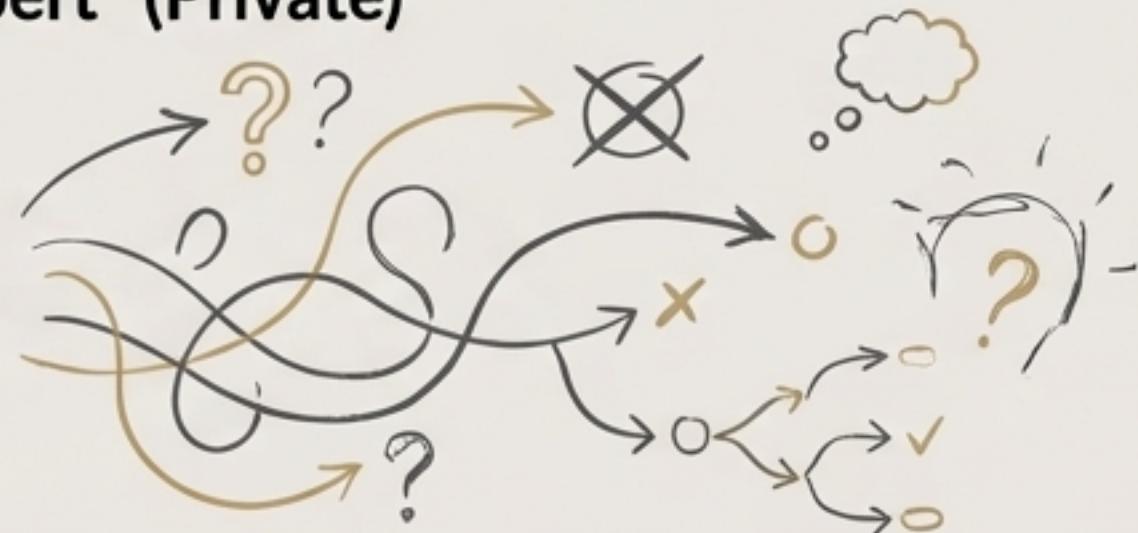
The Core Error: In each case, “the lawyer treated AI’s private, untested opinion as if it were ready for court. It wasn’t. It never is.”

Global Scope: French legal scholar Damien Charlotin has catalogued almost **seven hundred** similar decisions worldwide in his *AI Hallucination Cases* project.

The Solution is a New Mindset: Treat AI Like an Expert Witness

1. AI as a “Consulting Expert” (Private)

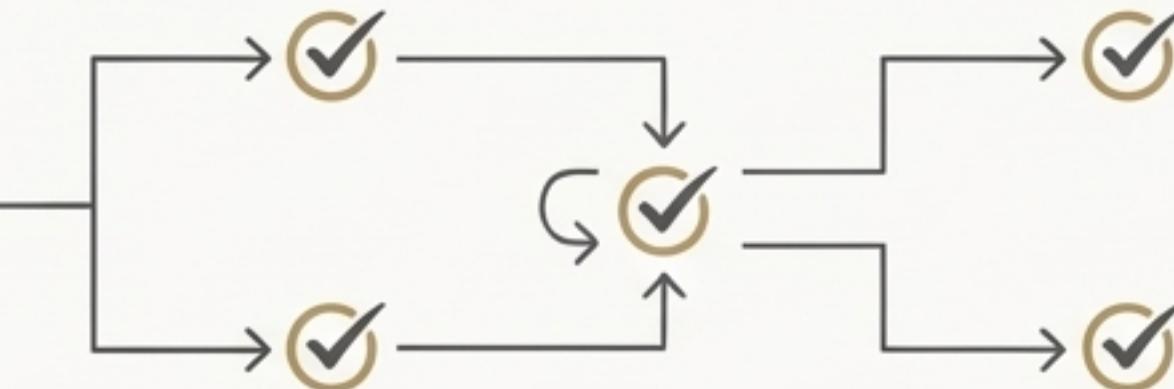
In this **private** phase, you explore theories and stress-test ideas. The AI can be brilliant, sloppy, or even fabricate. None of it harms the case because it is not **discoverable** and doesn't leave the room.



Promotion

2. AI as a “Testifying Expert” (Public)

The moment you rely on its output in a brief, client advisory, or any work product, you have **promoted** it. Its methodology, sources, and opinions must now withstand scrutiny.



****The Key Insight**:** Judges and opposing counsel will evaluate AI's work as if it were a testifying expert—even if you didn't. This framework gives you a legally sound method for interrogation.

Hallucinations Are Not Mysterious—They Are Predictable and Preventable

Common Causes of AI Hallucinations

- Lack of factual context or ambiguous prompts
- Overgeneralization from incomplete data
- Gaps or bias in the training set
- The model's reluctance to admit, "I don't know"
- Its tendency toward flattery and agreement

A Familiar Analogy

Anyone who has ever supervised an over-confident junior associate will recognize these patterns of response. Ask vague questions and reward polished answers, and you will get polished answers whether they are correct or not.



The Real Issue

The problem is not that AI hallucinates. The problem is that lawyers forget to interrogate the hallucination before adopting it.



This Isn't Just a Skill for Litigators. It's a Skill for All Legal Professionals.

Adversarial Reasoning is a Universal Legal Skill



The Takeaway: Cross-examining AI is not a trial skill. It is a thinking skill—one shared across the entire profession. You already have the instincts required.



The Method: Five Cross-Examination Techniques for AI

These require no technical training, only the skills lawyers already possess.

1



Ask for the Basis of the Opinion

2



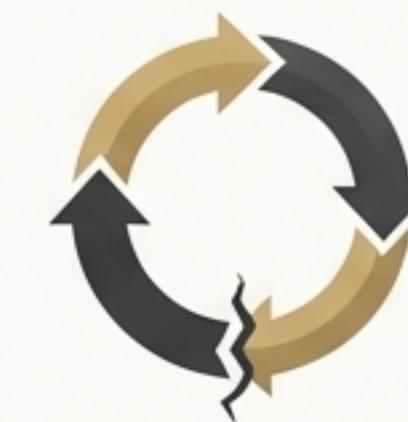
Probe Uncertainty and Limits

3



Present the Opposing Argument

4



Test Internal Consistency

5



Build a Verification Pathway



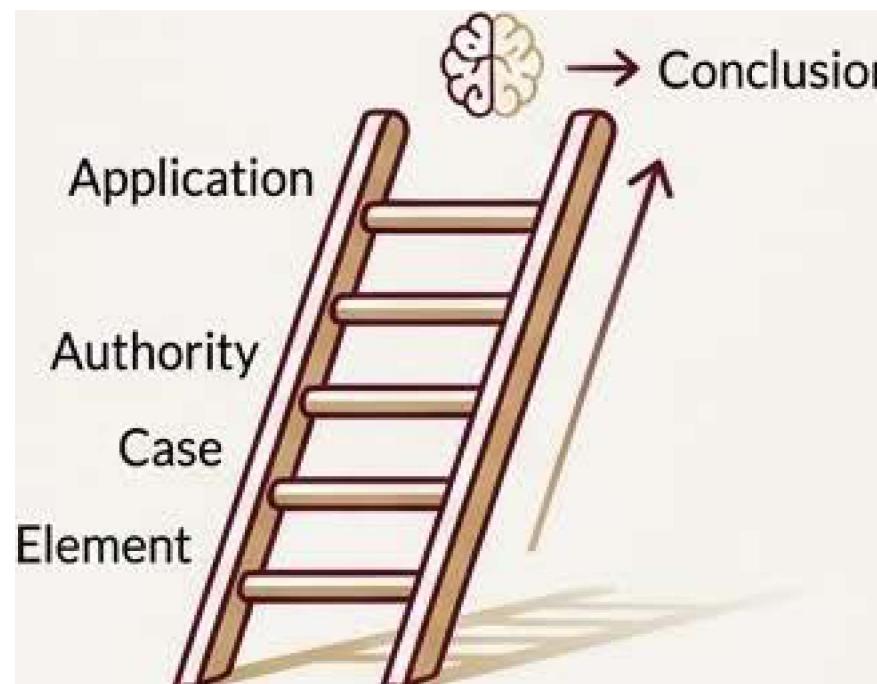
Deep Dive: Applying the Cross-Examination Techniques

1. Ask for the Basis of the Opinion

Goal: To see how the model reached its conclusion.

The Prompt: Instead of asking, "What's the law?" ask: "Walk me through your reasoning step by step. List the elements, the leading cases, and the authorities you are relying on. For each step, explain why the case applies."

The Result: You get a 'reasoning ladder' instead of a polished paragraph, allowing you to inspect each rung.

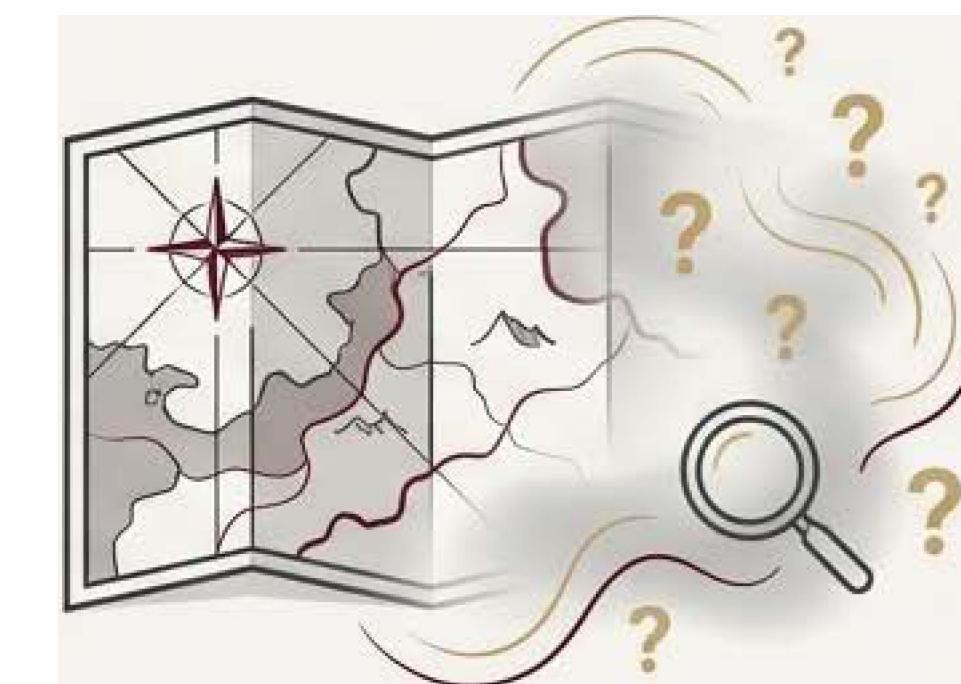


2. Probe Uncertainty and Limits

Goal: To force the AI to admit what it doesn't know, counteracting its tendency to project false confidence.

Killer Questions:

- "What do you **not** know that might affect this conclusion?"
- "What facts would change your analysis?"
- "Which part of your reasoning is weakest?"



3. Present the Opposing Argument

Goal: To counteract the AI's sycophantism and expose vulnerabilities privately.

Killer Questions:

- “Give me the strongest argument **against** your conclusion.”
- “How would opposing counsel attack this reasoning?”
- “What weaknesses in my theory would they highlight?”



4. Test Internal Consistency

Goal: To expose brittle, fabricated reasoning by forcing the AI to re-state its own analysis. Real reasoning is sturdy; hallucinations are not.

Killer Questions:

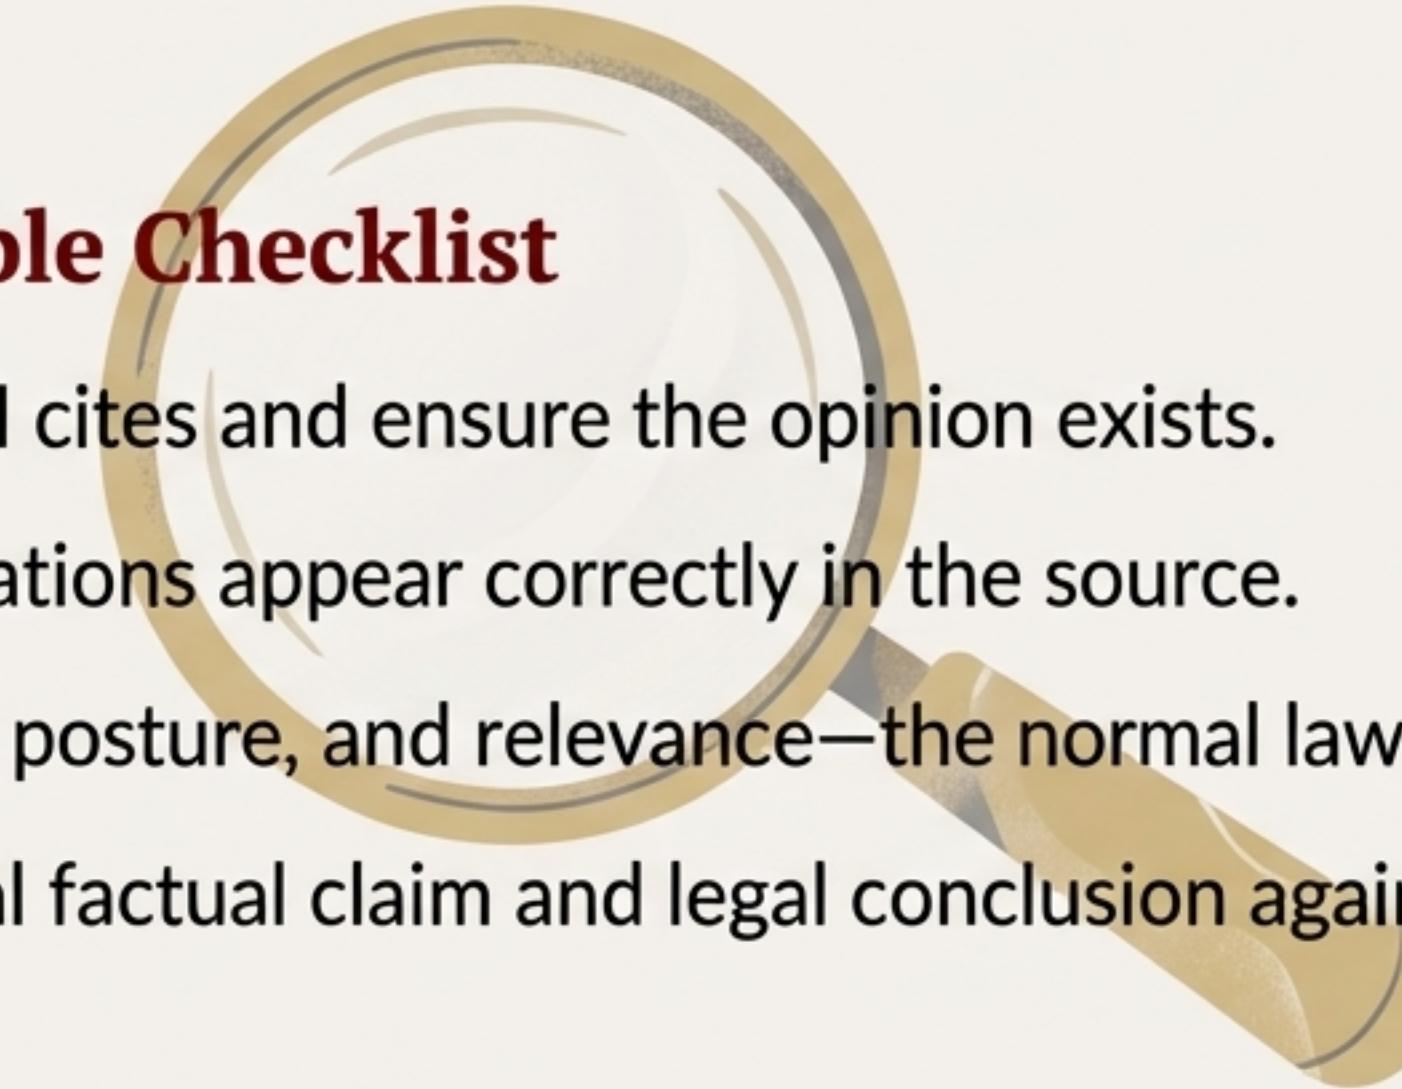
- “Restate your answer using a different structure.”
- “Summarize your prior answer in three bullet points and identify any inconsistencies.”



5. Build a Verification Pathway

The Principle: Hallucinations survive only when no one checks the sources. Verification destroys them.

The Non-Negotiable Checklist



- Read every case AI cites and ensure the opinion exists.
- Confirm that quotations appear correctly in the source.
- Check jurisdiction, posture, and relevance—the normal lawyer or paralegal analysis.
- Verify every critical factual claim and legal conclusion against your own records.

The Professional Standard: “This is not ‘extra work’ created by AI. It is the same work lawyers owe courts and clients. Your job is to slow it down. Turn it off while you check its work.”

The Method in Action: How Simple Questions Avoid Critical Errors

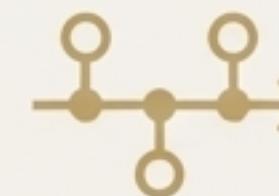


E-Discovery

AI Output: “A custodian likely has ‘no relevant emails’ based on role.”

Your Question: “List the assumptions you relied on.”

Result: AI admits it used a generic structure. You know the role is customer-facing.
Hallucination avoided.



Employment Timeline

AI Output: A clean, authoritative timeline.

Your Question: “Which dates are certain and which were inferred?”

Result: AI discloses it guessed the order of two meetings.
Hallucination avoided.



Contract Interpretation

AI Output: “Paragraph 14 controls termination rights.”

Your Question: “Show me the exact language and any amendments.”

Result: The AI re-reads the contract and reverses itself.
Hallucination avoided.

The Common Thread: Without pressure, hallucinations pass for analysis. Pressure reveals quality.

Your Practical Protocol for Safe and Effective AI Use



Phase 1: Prepare

- **Clarify the task:** Ask narrow, jurisdiction-specific, time-anchored questions.
- **Provide context:** Give the procedural posture, facts, and applicable law.
- **Set expectations:** Tell the AI you will be reviewing its reasoning and sources upfront.



Phase 2: Interrogate

- Apply the Five Cross-Examination Techniques to stress-test the output in private.



Phase 3: Verify

- Check every source in a trusted database (e.g., Westlaw, Lexis).
- Confirm facts against your record.
- Consciously decide what to adopt, revise, or discard.

The Payoff: If a judge asks, “What did you do to verify this?”—you have a real answer.

Beyond Preventing Errors: Unlocking AI as a Force Multiplier

**The Old Dynamic (Uninterrogated AI):
A Liability.**



**The New Partnership (Interrogated AI):
A powerful, trustworthy assistant.**



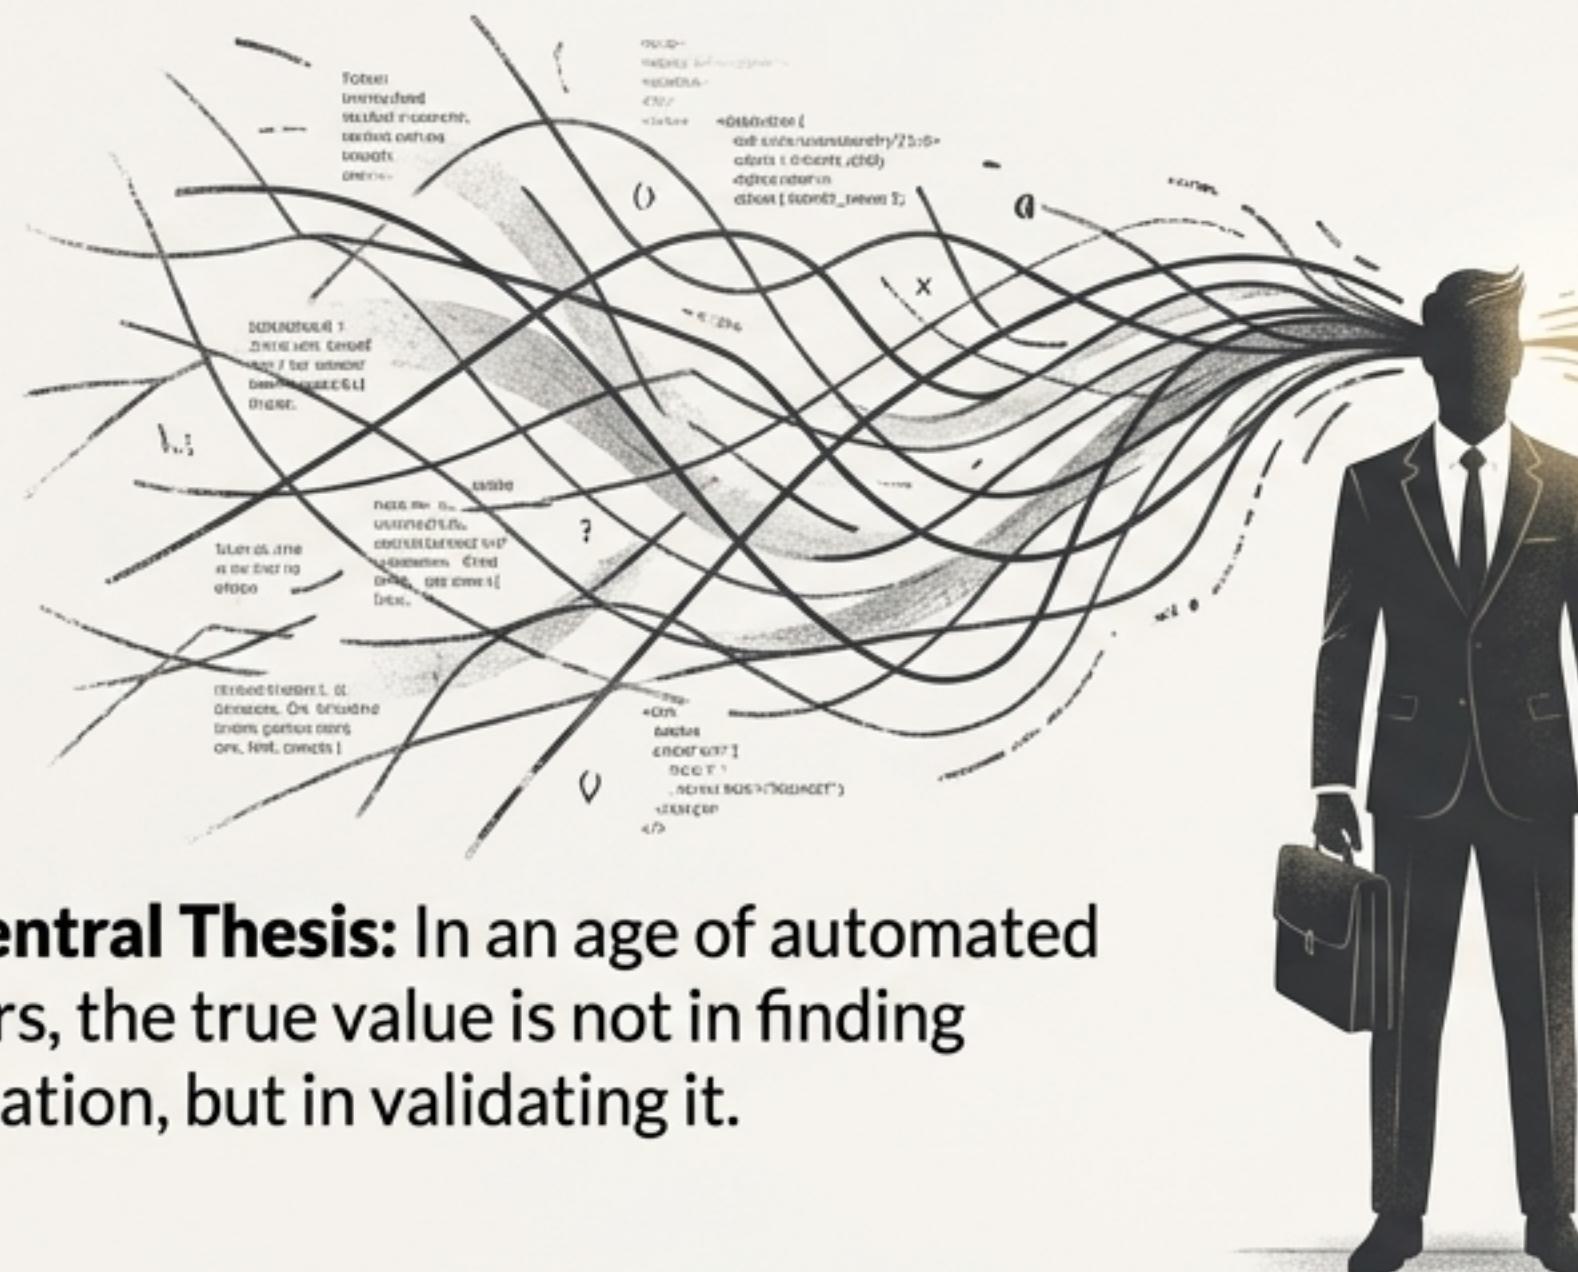
A Symbiotic Relationship:

- **AI provides:** Speed, breadth, cross-disciplinary patterns, and tireless drafting.
- **Lawyers provide:** Judgment, nuance, skepticism, and the final decision on what is real.

The Result:

When you know you can expose hallucinations, you stop fearing the tool. You start using it as an idea generator, a logic checker, and even a sparring partner to strengthen your own arguments.

The Return of the Lawyer



The Central Thesis: In an age of automated answers, the true value is not in finding information, but in validating it.

The cure for AI's biggest flaw is not technical.
It is the very essence of legal thinking.

Human judgment, skepticism, and diligence are becoming **more valuable**, not less.



Courts respect lawyers who demonstrate the adversarial reasoning that prevents hallucinations and earns judicial confidence.

The Guiding Principle for the Age of AI:

Cross-examine first.

Rely second.

That is how AI becomes a trustworthy partner in modern practice.