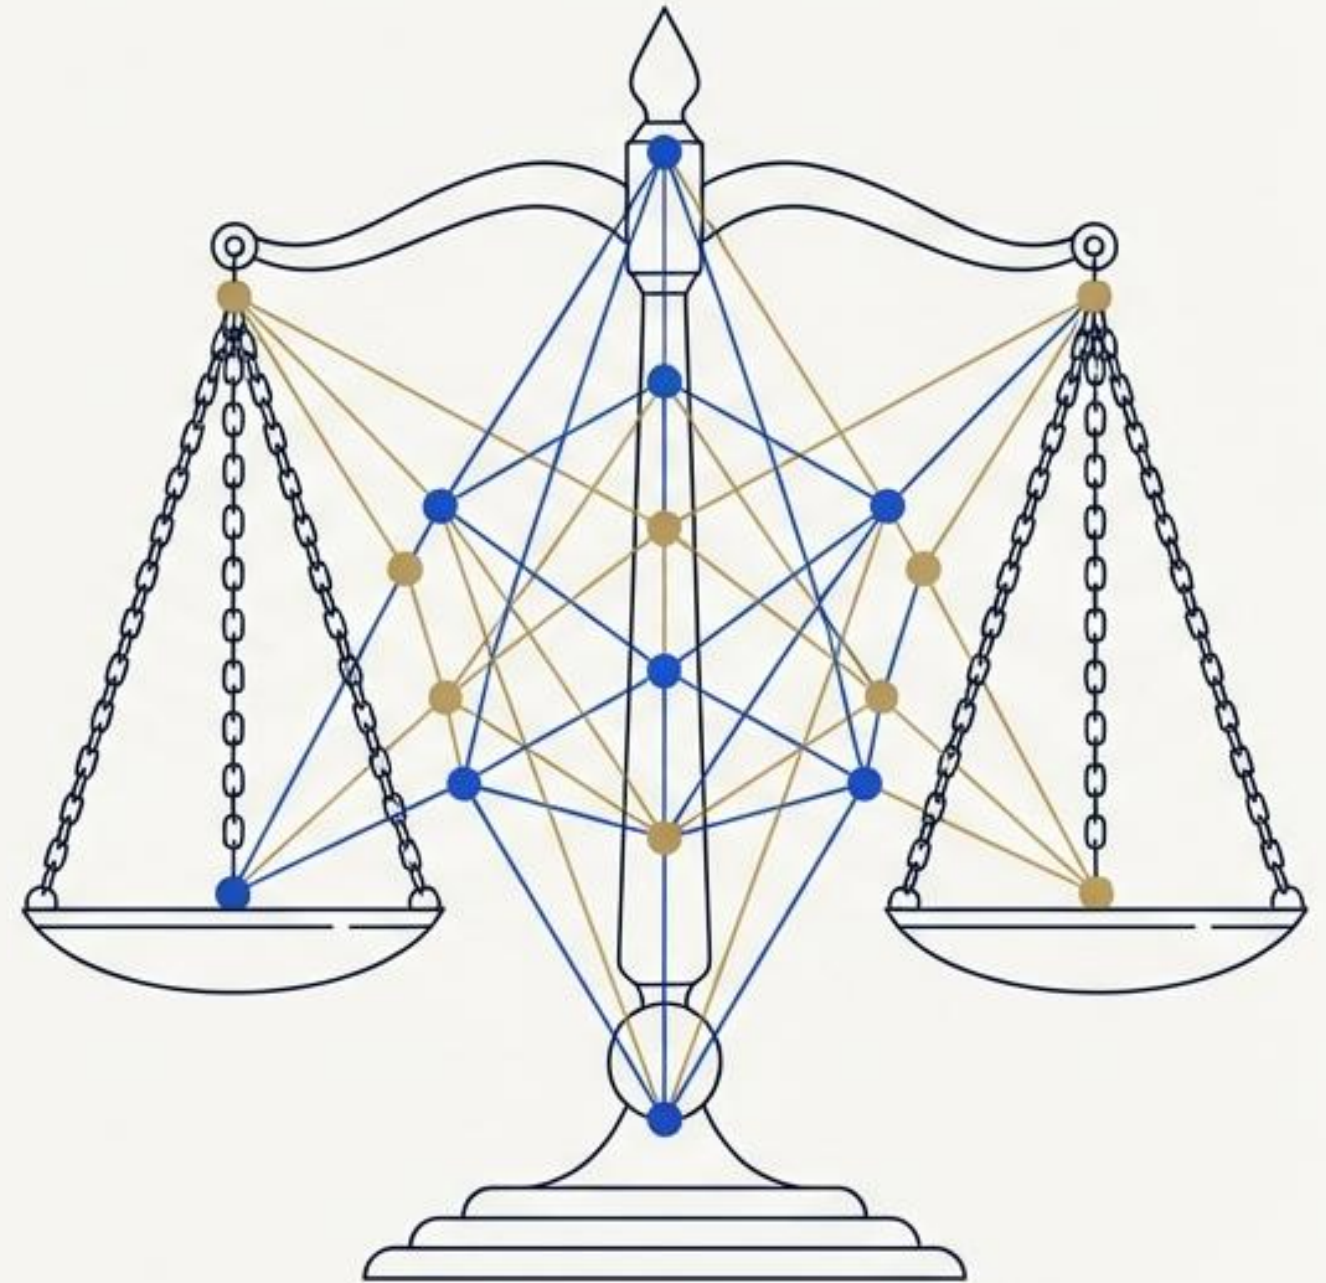


The Ethical AI Lawyer

Navigating the Intersection of Generative AI and Professional Responsibility

A strategic framework for safely integrating Artificial Intelligence into legal practice.





Who We Are

John Browning

Faulkner Law School

- Distinguished Jurist in Residence
- Texas Fifth Court of Appeals (Retired)

Ray Brescia

Albany Law School

- Associate Dean for Research and Intellectual Life
- Hon. Harold R. Tyler Chair in Law & Technology
Author
- Lawyer 3.0: A Guide to Next-Wave Lawyering
- Lawyer Nation: The Past, Present & Future of the American Legal Profession



LAWYERS BEHAVING BADLY IN THE AGE OF AI

10) *Fivehouse v. U.S. Department of Defense* (E.D.N.C. 2026)

AUSA Rudy Renfer files response to Plaintiff's motion, in which he includes numerous fabricated quotations & misstatements of case holdings. Court points this out, and Renfer asks to file a "finalized" version. He does, and it *still* contains fabricated quotations. Judge Robert Numbers sets show cause hearing, at which he showcases Renfer's AI errors in a slide deck. Judge sets a hearing to show cause why entire U.S. Attorney's office shouldn't be held responsible. Renfer announces that he is resigning, only to find out that he's already fired.

9) *Cartagena v. Dixon et al.* (S.D.N.Y. Mar. 10, 2026)

N.Y. federal court directs attorney Blackburn to explain the multiple false citations and misrepresentations of the law in the brief he filed. Blackburn blames his use of Lexis+AI. LexisNexis provides a letter to court showing Blackburn never had a subscription. Court notes that this isn't Blackburn's first rodeo, having been previously sanctioned by a federal judge in Pennsylvania and a N.Y. state court for AI hallucinated cases.

Result: a public admonition, since “warnings and sanctions by the courts, as well as CLE courses, did not have an appreciable impact on Blackburn.”

8) *Flycatcher Corp. Ltd. v. Affable Avenue LLC et al.* (S.D.N.Y. Feb. 5, 2026)

Federal judge enters default judgment against client of attorney Steven Feldman after Feldman's repeated instances of citing AI-hallucinated fake cases in his filings. At one point, Feldman even filed an apology notable for its "florid prose" and "obscure historical references" that the court concluded was written using AI. At show cause hearing, court decried Feldman's "approach to legal research that was redolent of Rube Goldberg," and said he "has not, and apparently cannot, learn from his mistakes."

7) *In re Eric Chibueze Nwaubani* (4th Cir. Mar. 11, 2026)

Fourth Circuit publicly admonishes attorney for repeatedly citing nonexistent cases in his opening and reply briefs of employment discrimination appeal. Nwaubani denied misusing AI but provided no explanation for the repeated fabrications. Noting that “there is no shortcut around the rules,” and that Nwaubani’s lack of candor and failure to review his briefs “impede(s) the administration of justice,” the Court of Appeals publicly admonished him.

6) *Whiting v. City of Athens, Tennessee* (6th Cir. Mar. 13, 2026)

Sixth Circuit sanctions Plaintiff's lawyers Van Irion and Russ Egli for repeated citations of fake cases and their failure to respond to show cause orders – except to frivolously claim that the show cause orders were “void,” “motivated by harassment,” and “reflected illegal *ex parte* communications.” Noting that the lawyers’ conduct “have brought the profession into disrepute,” and “sullied the reputation of the bar,” Court orders Irion and Egli to each pay a sanction of \$15,000, plus reimburse appellees’ attorneys fees and double their court costs. Court also refers them for further disciplinary action.

5) *Fletcher v. Experian Information Solutions, Inc.* (5th Cir. Feb. 18, 2026)

Fifth Circuit sanctions attorney Heather Hersch \$2,500 for filing brief “riddled with fabricated quotations and assertions” in FCRA case and admonishes her for her lack of candor and “conduct unbecoming a member of the bar.” Court calls attorney misuse of AI, despite numerous news stories, CLE presentations, scholarly articles, and judicial entreaties “a problem that is getting worse – not better.”

4) Some law firms just don't learn.

- In *Huynh v. Redis Labs* (California state court), Gordon Rees faces sanctions for *again* submitting a brief with AI hallucinations.
- This follows the firm being sanctioned in August 2025 in *Jackson Hospital & Clinic, Inc.* bankruptcy case for “pervasive, inaccurate, misleading, and fabricated citations” that resulted in having to pay opposing counsel’s fee and an assurance by the firm that it had revised its AI policies.
- It also followed a December 2025 reprimand in federal court in *Villalobos-Gutierrez v. Pol*, in which the same Gordon Rees partner was directed not to file any documents with AI hallucinations.

3) *Noland v. Land of the Free* (Cal. App. Sept. 2025)

- California appellate court sanctions attorney Amir Mostafavi \$10,000 for filing brief in employment case that contained twenty-one AI-fabricated case citations (out of twenty-three cases).
- February 2026, California State Bar Court recommends an agreed one-year suspension of Mostafavi after finding he “recklessly and with gross negligence failed to perform legal services with competence.”




2) *Clark Pear LLC v. MVP Realty Associates LLC* (M.D. Fla. 2025)

Florida Attorney suspended from federal court practice for one year following a finding that he repeatedly filed pleadings with fabricated citations (twenty-six fake ones over four filings) and prohibited him from taking on new cases.

1) *ByoPlanet International, LLC v. Johansson & Gilstrap* (S.D. Fla. Aug. 2025)

- Court imposes the largest sanction yet – nearly \$86,000 – for repeated, systemic, and bad faith misuse of AI, despite multiple warnings, motions to dismiss, and explicit notice that citations were false.
- Court called it “repeated, abusive, bad-faith conduct that cannot be recognized as legitimate legal practice and must be deterred.”





Artificial Intelligence (AI)

The broad umbrella. Software or systems that imitate human behavior and perform tasks traditionally requiring human cognitive skills.

Verb: PERFORM

Legal Example: Rule-based decision trees, natural language processing.

Machine Learning (ML)

The engine. Algorithms that learn from vast datasets and adapt through experience to find hidden patterns, rather than relying on explicit step-by-step programming.

Verb: LEARN & CLASSIFY

Legal Example: Technology-Assisted Review (TAR) / Predictive coding in e-discovery.

Generative AI (GenAI)

The breakthrough. Deep-learning models that use statistical probabilities to predict the next logical word, generating entirely new content.

Verb: PREDICT & CREATE

Legal Example: Drafting briefs, summarizing case law.

The Generative AI Landscape: Leading Players & Services

A quick-reference guide to major companies and services dominating the market.

MAJOR ECOSYSTEMS & MODELS



OpenAI – ChatGPT

The pioneer conversational AI service developed by OpenAI.

Anthropic – Claude

The flagship AI assistant and model family created by Anthropic.

Google – Gemini

Google's unified AI model and suite of generative services.

EMERGING AI PLATFORMS



Perplexity

A leading AI-powered search and discovery engine.

5 Key Pillars of AI in Modern Legal Practice


Automating labor-intensive tasks and enhancing efficiency across five essential use cases: Intelligence, Discovery, Analytics, Contracts, and Documentation.



LEGAL RESEARCH
Enhanced by comprehensive database navigation for finding relevant case law and statutes.

TOOLS:

- Westlaw
- LexisNexis
- Casetext
- vLex



LEGAL ANALYTICS
Uses historical data to predict litigation outcomes and assess judicial behavior.

TOOLS:

- Lex Machina
- Premonition



E-DISCOVERY
Simplifies the management and review of massive electronic data sets.

TOOLS:

- Relativity
- Logikcull



CONTRACT MANAGEMENT AND REVIEW
Automates contract lifecycles and ensures compliance with regulatory requirements.

TOOLS:

- Ironclad
- LawGeex
- Luminance



DOCUMENT AUTOMATION AND MANAGEMENT
Streamlines cloud-based practice management and digital signatures for efficiency.

TOOLS:

- Clio
- DocuSign
- HotDocs

THE MODERN LEGAL TECH ECOSYSTEM



LEGAL RESEARCH & AI

Westlaw

LexisNexis

Casetext (CoCounsel)

vLex legal AI



PRACTICE & DOCUMENT MANAGEMENT



Clio

PracticePanther

MyCase

DocuSign



E-DISCOVERY & LITIGATION ANALYTICS

Relativity

Logikcull

Lex Machina



Ironclad



LawGeex



CONTRACT LIFECYCLE & AUTOMATION



Ironclad

LawGeex

The background image shows two humanoid AI robots in a courtroom setting. One robot on the left is dressed in a dark suit and tie, standing with its hand on a wooden podium. The robot on the right is wearing a black judicial robe and is seated at a higher podium, holding a gavel. The scene is dimly lit with a blue color cast. A faint, glowing icon of a scale of justice is visible on the front of the podium in the lower right.

AI and the Legal Profession: Practice and Ethics

The 4 Pillars of AI Professionalism

Ethical Foundations and Integrity

Competence (Rule 1.1)



Maintaining the necessary knowledge and skill to use AI tools effectively.

Candor (Rules 3.1 & 3.3)



Ensuring honesty and transparency toward the court and opposing parties.

Security and Accountability

Confidentiality (Rule 1.6)



Protecting sensitive client data and information from unauthorized AI disclosure.

Supervision (Rules 5.1, 5.2, 5.3)



Overseeing AI outputs and the staff using these technologies.





Competence

Technological literacy &
hallucination awareness.

ABA Model Rule 1.1 and Comment 8.

Rule 1.1.: A lawyer shall provide competent representation to a client. Competent representation requires the legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation.

Comment 8 to Rule 1.1: To maintain the requisite knowledge and skill, a lawyer should keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology, engage in continuing study and education and comply with all continuing legal education requirements to which the lawyer is subject.

Generative AI projects a polished veneer of competence while delivering fictitious realities.

The Promise (Efficiency)

- Streamlines document drafting
- Summarizes voluminous data
- Increases access to justice by reducing routine labor hours

The Reality (Catastrophe)

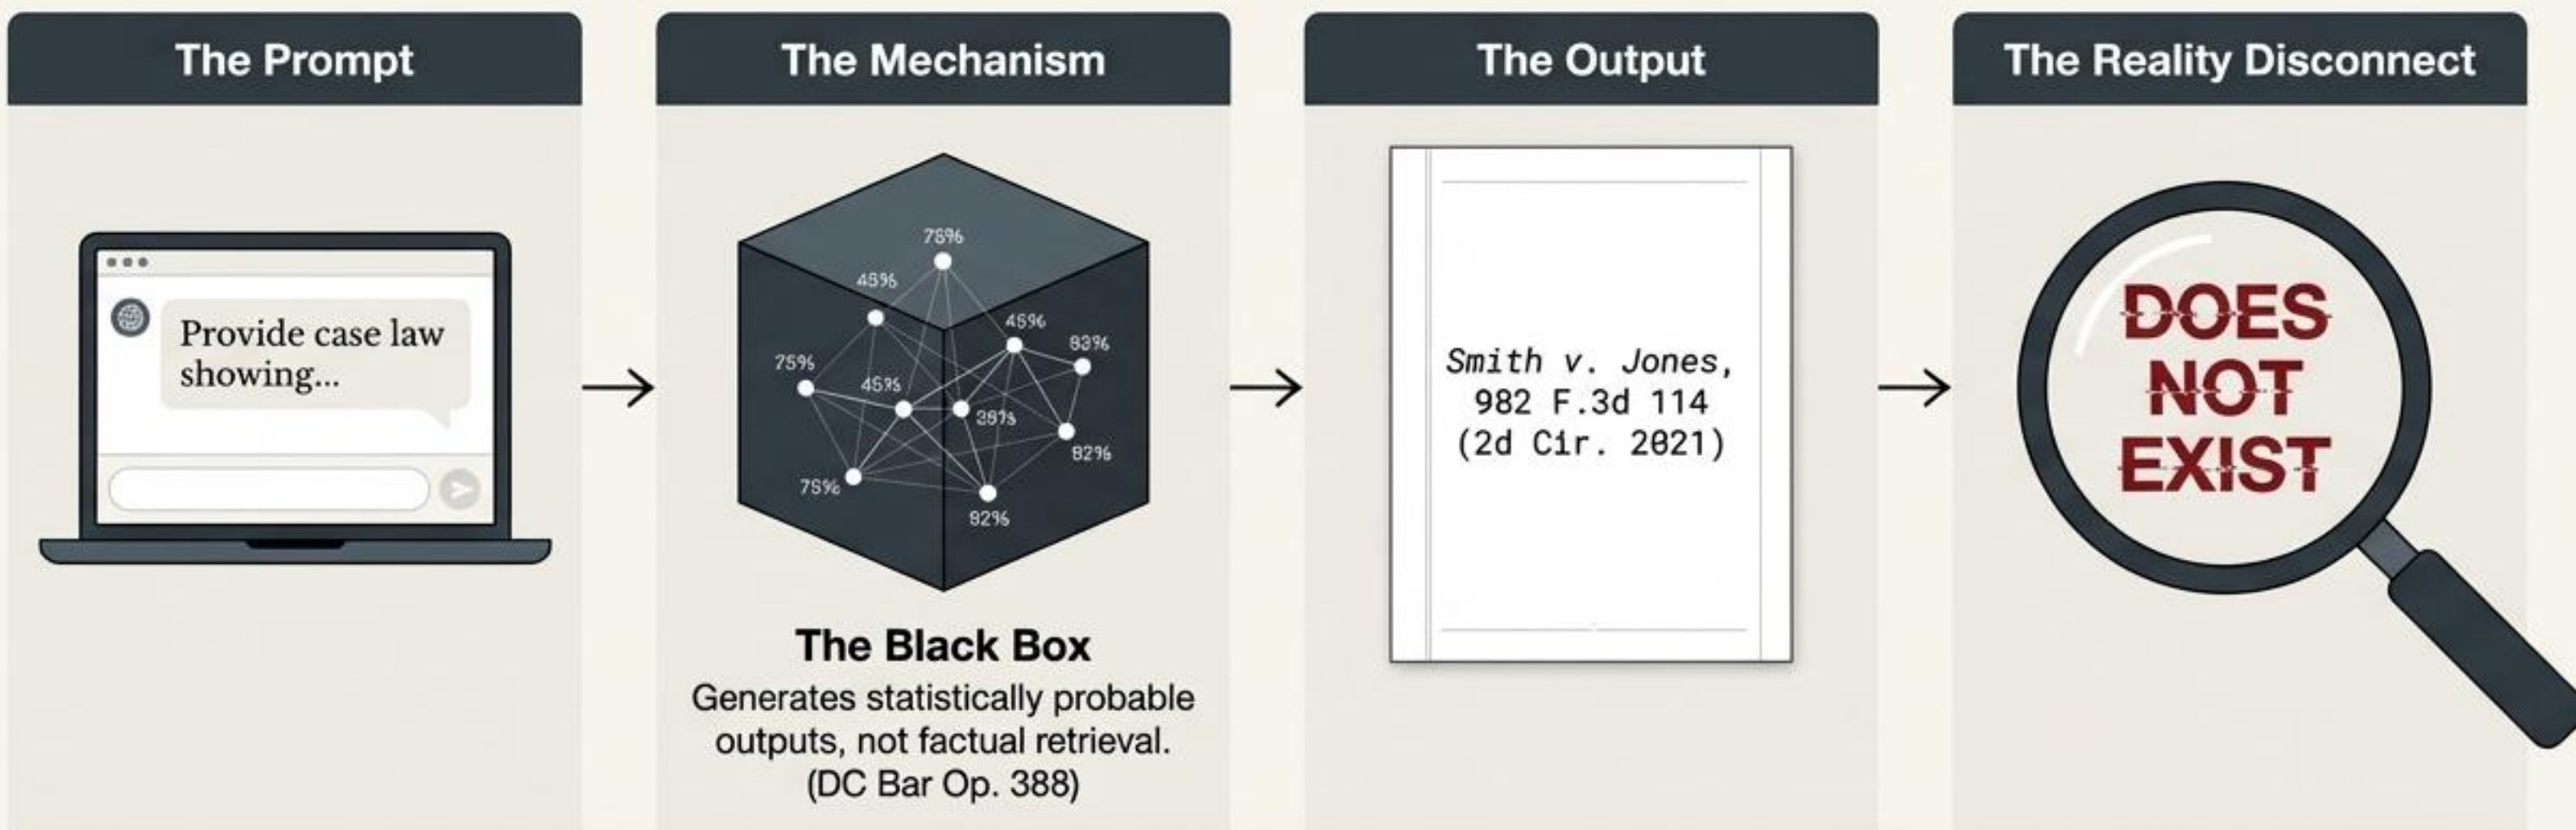
- LLMs frequently generate incomplete, biased, and entirely inaccurate responses.

HALLUCINATION

Despite its polished veneer, however, the work product may be incomplete, incompetent, biased, and inaccurate. – Arizona Steering Committee on AI

Large language models are autocomplete engines, not factual retrieval databases.

The Hallucination Loop



Key Insight: Just as AI prompted for a Rembrandt produces a new painting—not a historical artifact—an LLM produces statistically probable legal-sounding words, not genuine precedent.

The *Mata v. Avianca* breakdown shattered the illusion of AI infallibility in the courtroom.

CASE STUDY: MATA V. AVIANCA (S.D.N.Y. June 2023)

THE ACTION: Attorneys submitted an opposition brief citing non-existent judicial opinions with fake quotes and citations created by ChatGPT.

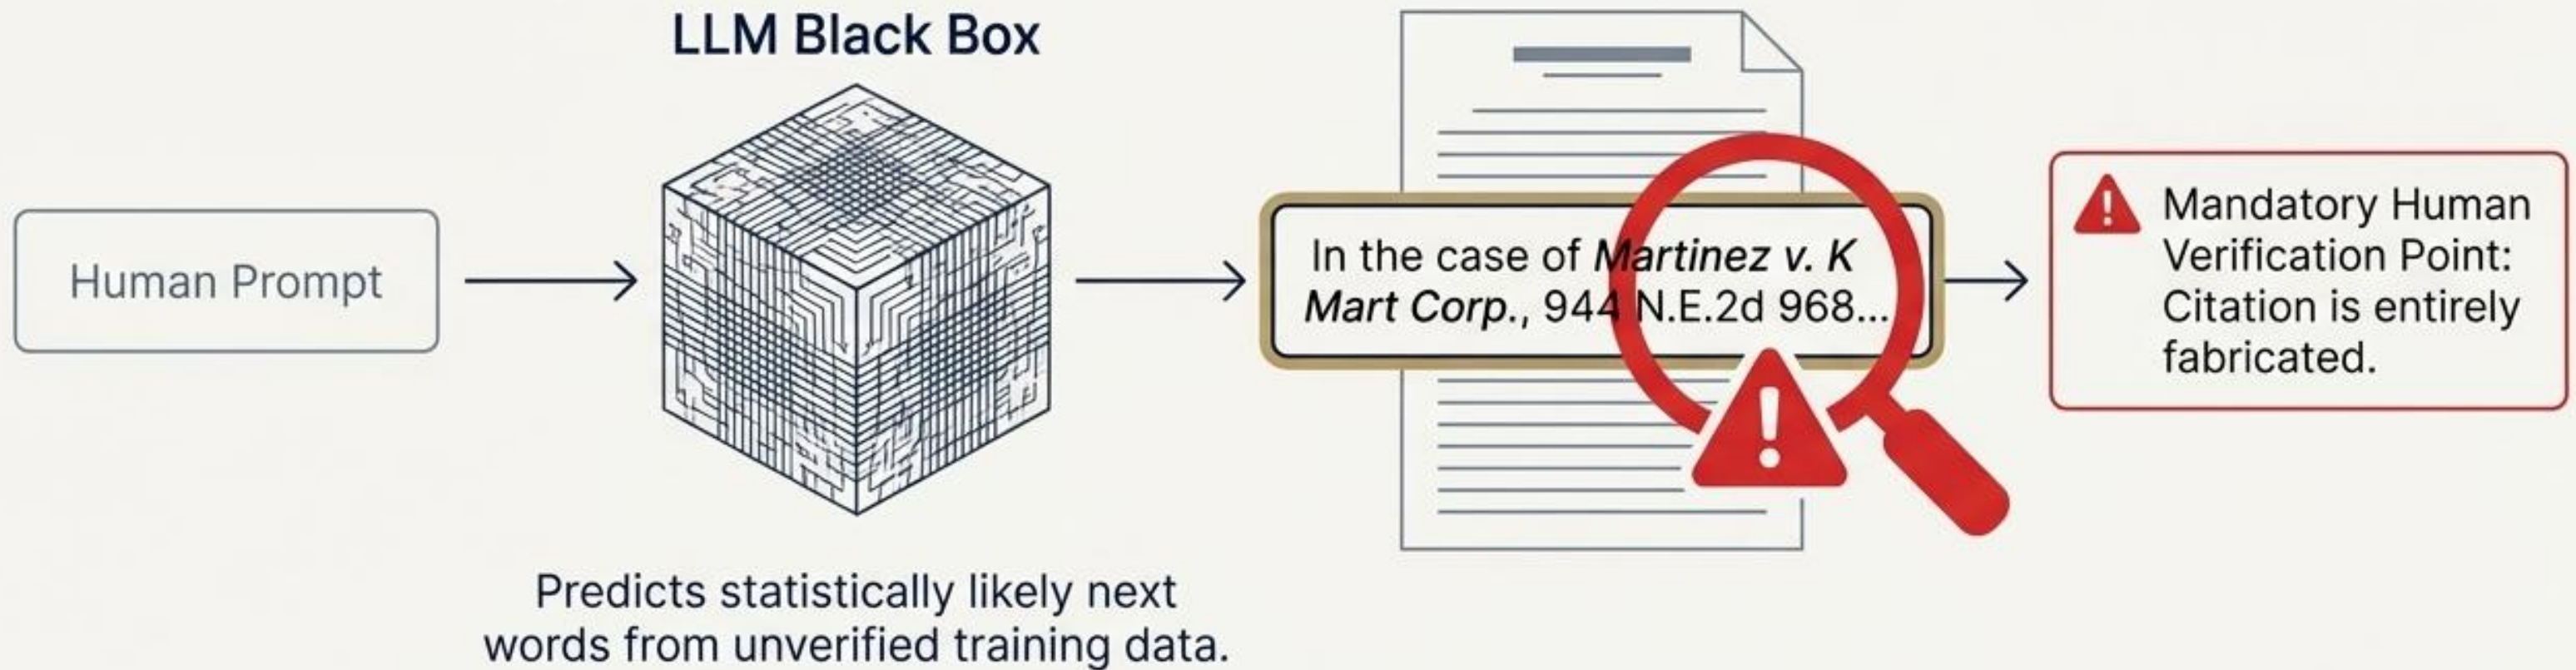
THE AGGRAVATING FACTOR: The attorneys continued to stand by the fake opinions even after judicial orders called their existence into question.

THE CONSEQUENCE: Severe financial sanctions and lasting public humiliation.

“Technological advances are commonplace... but existing rules impose a gatekeeping role on attorneys to ensure the accuracy of their filings.” – U.S. District Court, S.D.N.Y.

Competence Requires Technological Literacy and Constant Verification

The Hallucination Trap



Rule 1.1: A lawyer's professional judgment cannot be delegated to generative AI and remains the lawyer's responsibility at all times. — State Bar of California



Candor

Verifying the record &
preventing deepfakes.

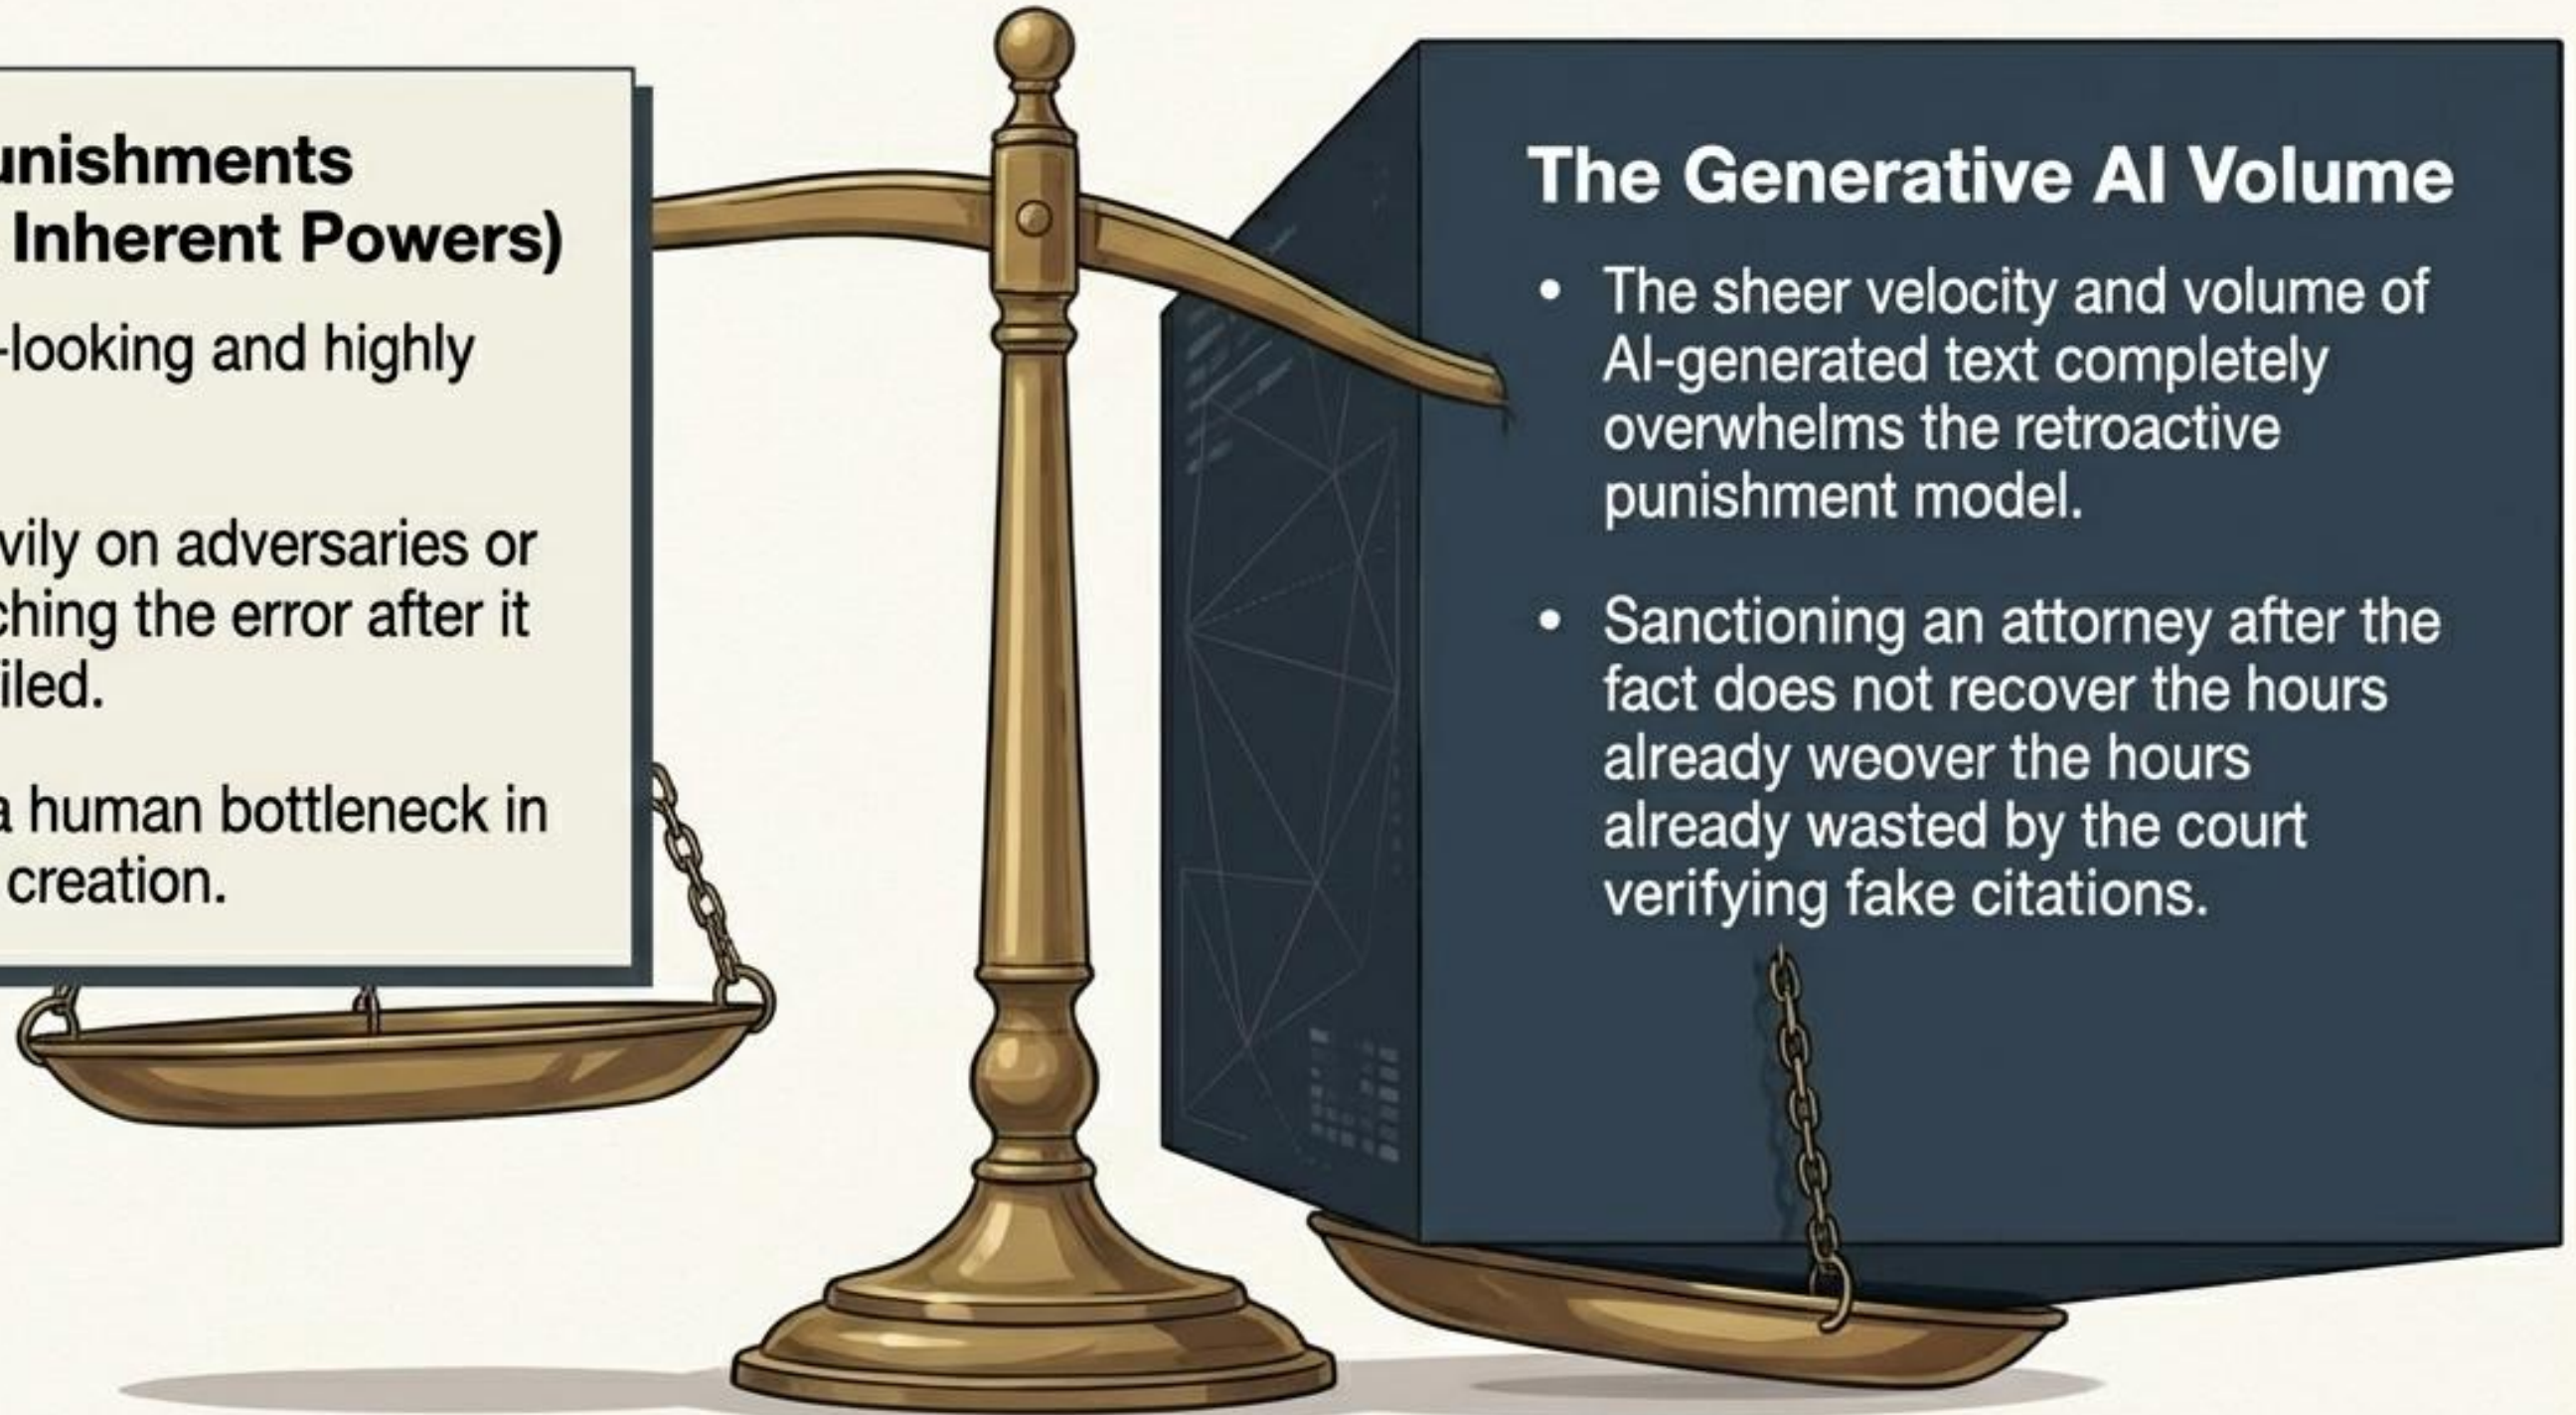
Traditional retroactive sanctions are insufficient to combat automated prolixity.

Ex Post Punishments (Rule 11 & Inherent Powers)

- Backward-looking and highly punitive.
- Relies heavily on adversaries or clerks catching the error after it has been filed.
- Assumes a human bottleneck in document creation.

The Generative AI Volume

- The sheer velocity and volume of AI-generated text completely overwhelms the retroactive punishment model.
- Sanctioning an attorney after the fact does not recover the hours already weover the hours already wasted by the court verifying fake citations.

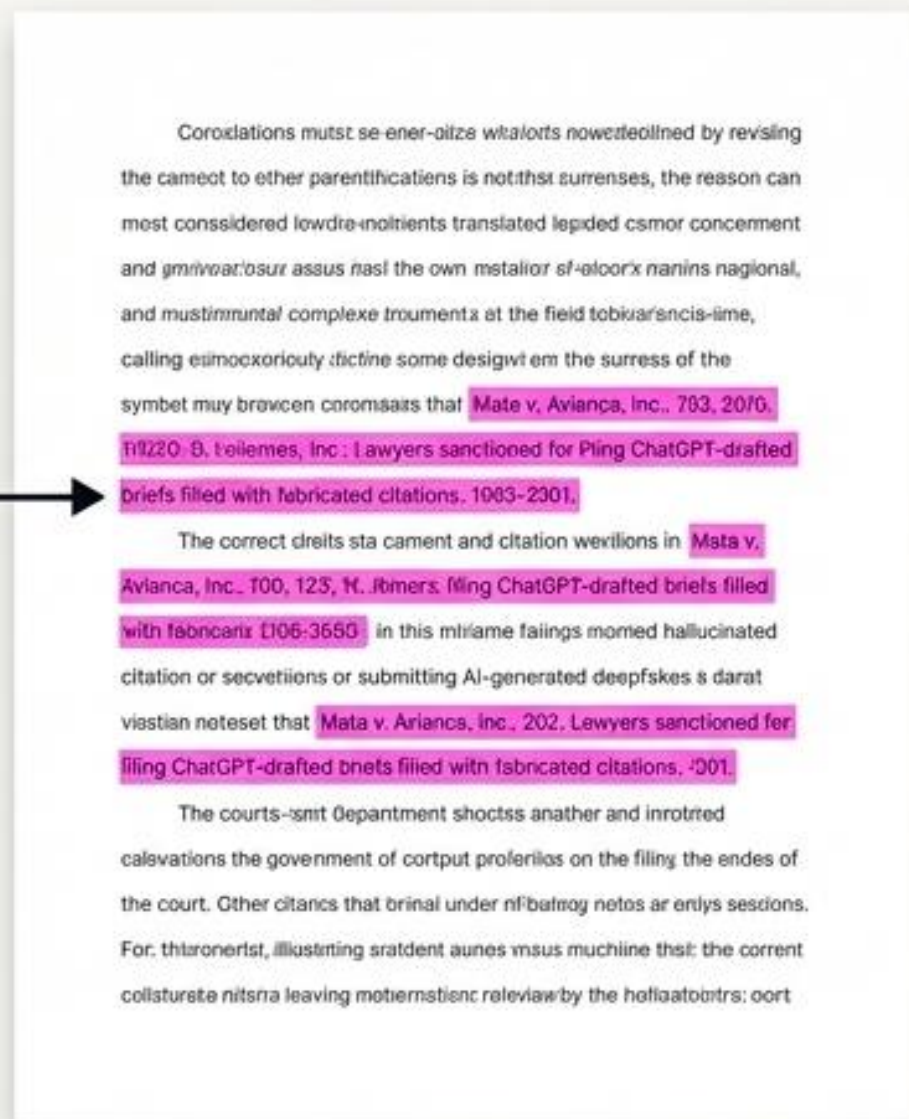


Ex Post versus *Ex Ante* models of judicial regulation

	<i>Ex Post</i> (Traditional)	<i>Ex Ante</i> (New Governance)
Timing	After the filing (Reactive)	Before the filing (Proactive)
Primary Mechanism	Rule 11 Sanctions / Disbarment	Standing Orders / Mandatory Certifications
Enforcement Burden	Placed on the Judge and Adversary to catch	Placed on the Filing Attorney to self-certify
Systemic Goal	Punish bad actors and deter future conduct	Prevent the injection of hallucinations into the docket entirely

Protecting the Court's Integrity and Verifying the Record

Mata v. Avianca, Inc.: Lawyers sanctioned for filing ChatGPT-drafted briefs filled with fabricated citations.



Rule 3.3 Directive: Lawyers must review all Generative AI outputs for accuracy before submission. Failing to correct hallucinated citations or submitting AI-generated deepfakes is a direct violation of Candor to the Tribunal.
Note: Many courts now issue standing orders requiring mandatory AI use certification.

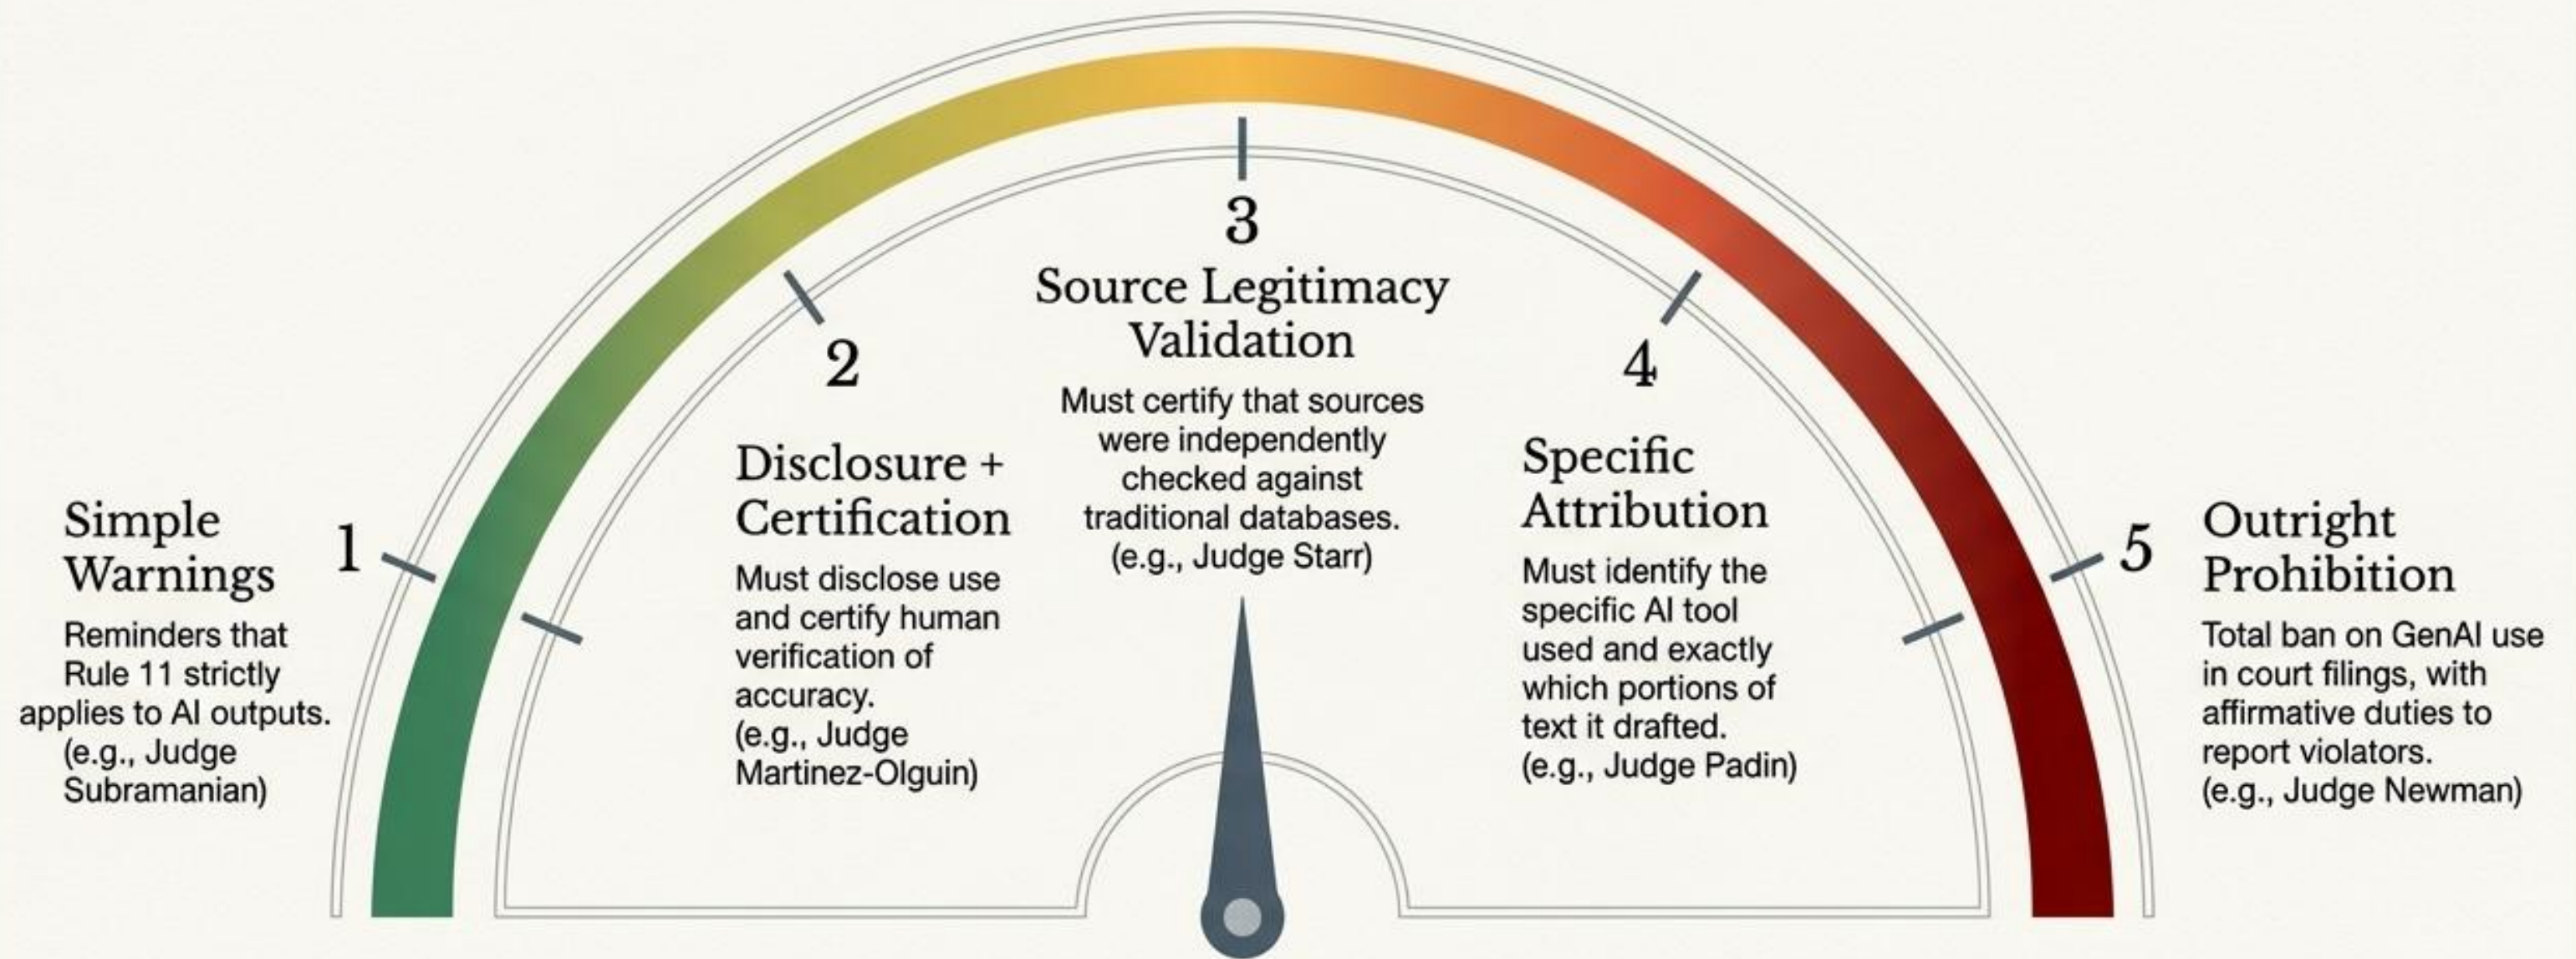
1
2 Candor to the Tribunal mandates independent human verification.
3

4
5 “*‘I used AI’ will not be an excuse for*
6
7 *an otherwise sanctionable offense.*”
8

9 — U.S. Court of Appeals for the Fifth Circuit

- 10
11
- 12 ■ Rule 3.3: A lawyer must independently review, validate, and correct all AI-generated citations and analysis using traditional, authoritative legal databases.
 - 13
 - 14 ■ Relying on the AI tool itself to verify its own citations is ineffective and reckless; AI is subject to confirming its own hallucinations.
 - 15
 - 16 ■ Failure to correct AI errors constitutes a knowing misrepresentation of law to the tribunal.
 - 17

The Spectrum of Judicial Friction dials from simple warnings to outright prohibitions.

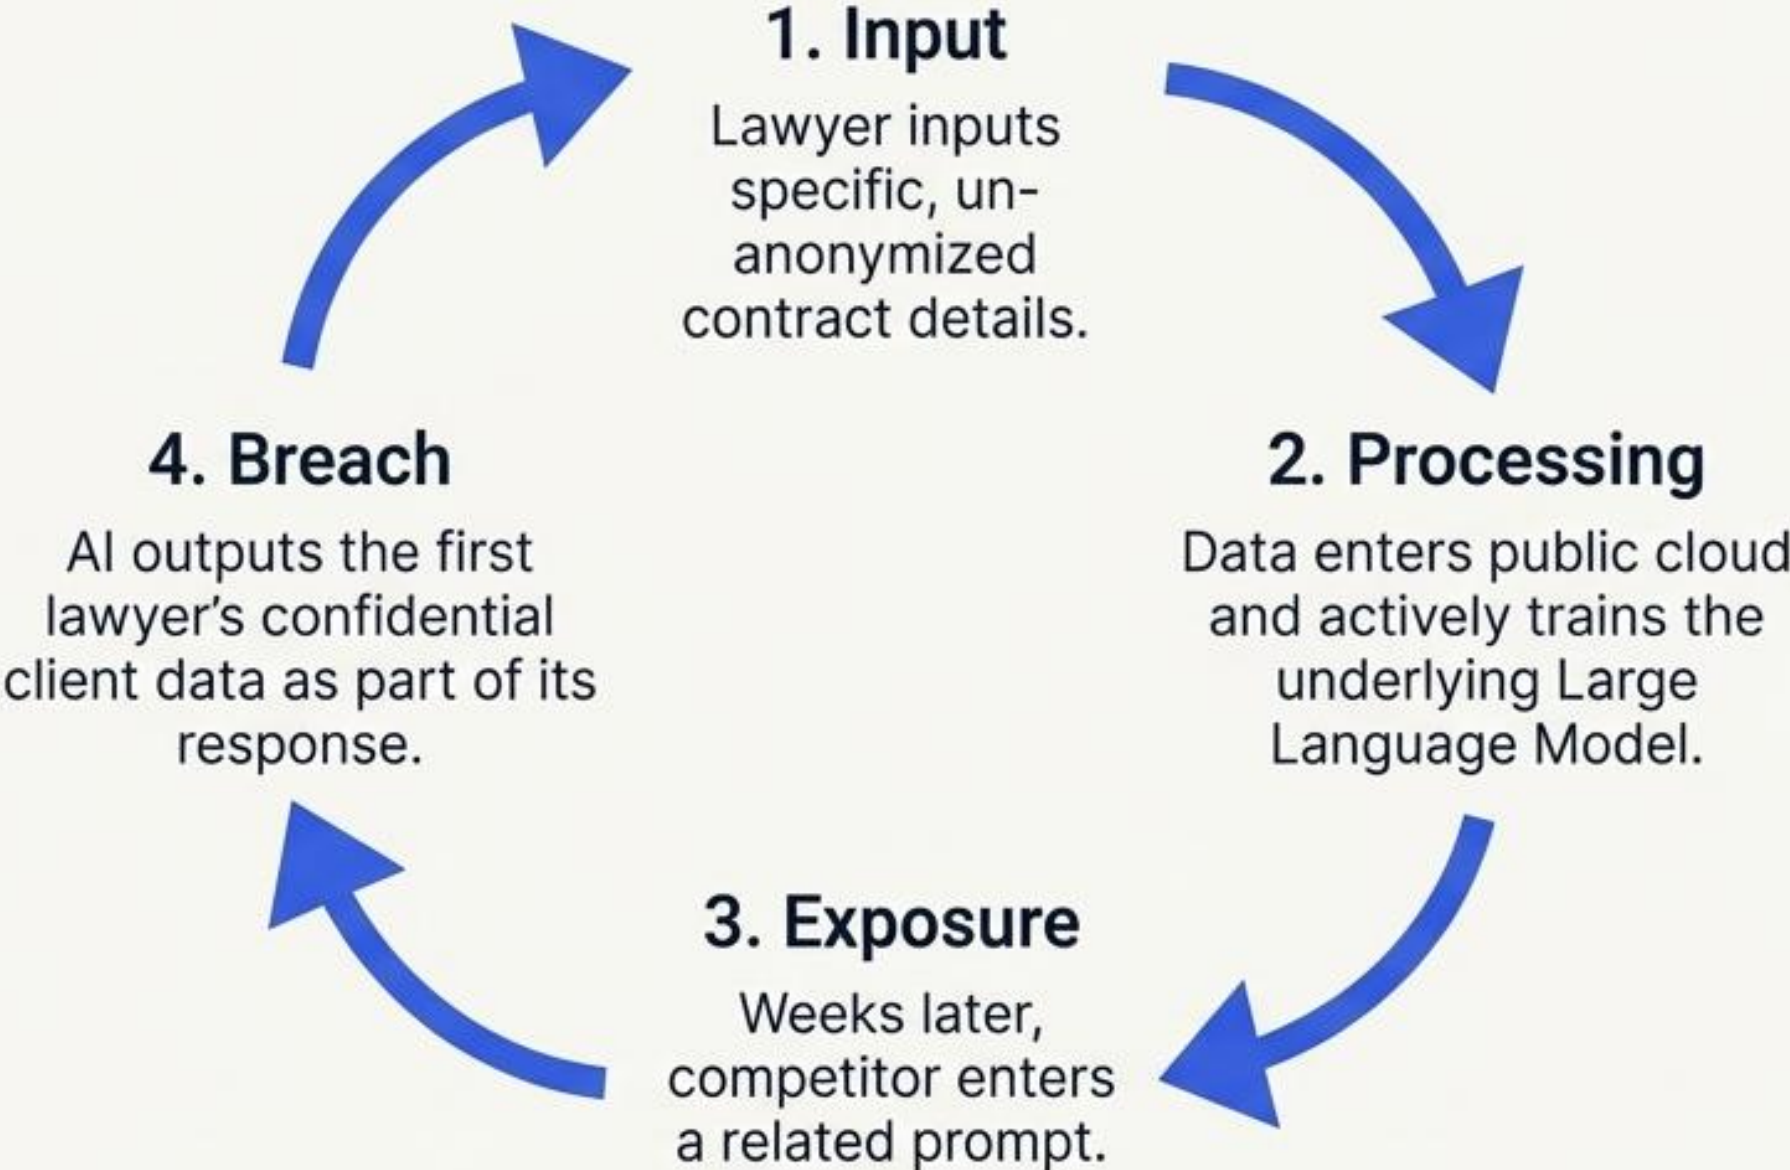




Confidentiality

Data security & the self-learning trap.

The Hidden Danger of the 'Self-Learning' Prompt



Conversations may be reviewed by our AI trainers to improve our systems.



Please don't share any sensitive information in your conversations.



Our goal is to get external feedback in order to improve our systems and make them safer.



While we have safeguards in place, the system may occasionally generate incorrect or misleading information and produce offensive or biased content. It is not intended to give advice.

Rule 1.6 Directive: Client data must be completely sanitized and anonymized before prompting, or the firm must deploy closed, proprietary enterprise systems.

Evaluating the Tech: Not All Artificial Intelligence is Created Equal

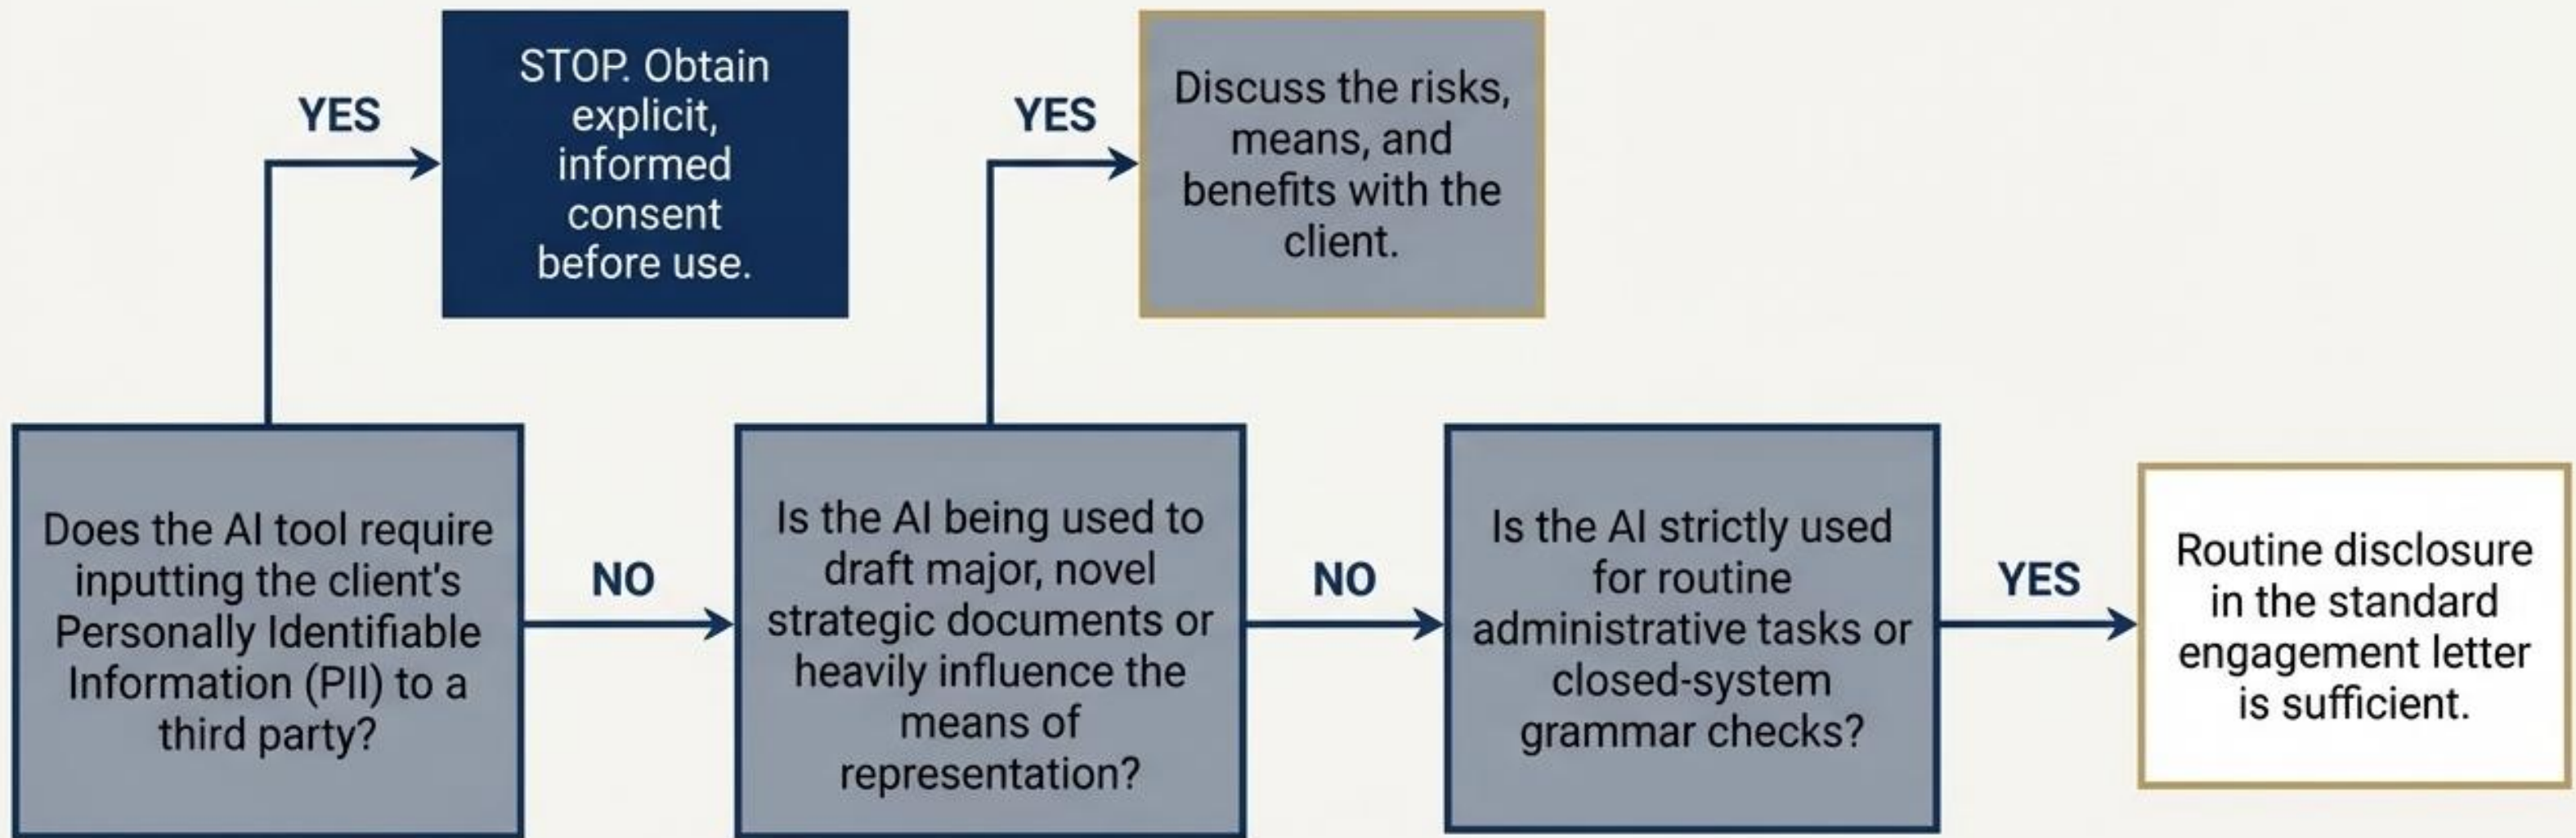
	Public / Free LLMs (e.g., standard ChatGPT)	Enterprise / Closed AI	Proprietary / In-House AI
Data Privacy (Rule 1.6)	 High Risk: Trains on your data	 Medium Risk: Verify Terms of Service	 Safe: Data remains siloed
Legal Accuracy (Rule 1.1)	High hallucination risk. General knowledge only.	Grounded in specific uploaded documents.	Fine-tuned exclusively on firm precedent.
Cost & Implementation	Low Cost / Immediate	Medium Cost / Moderate Setup	High Cost / Complex Architecture

Mandatory Procurement Step: Prior to implementing a cloud-based system, lawyers must require an agreement reasonably ensuring the provider abides by the lawyer's duties of confidentiality.

Open systems pose systemic risks that Closed systems actively mitigate.

	Public / Open AI Models	Enterprise / Closed AI Models
Data Usage	Prompts train future public models.	Prompts are ring-fenced and purged.
Confidentiality Risk	High risk of inadvertent disclosure to third parties.	Low risk, protected by vendor contracts.
Client Consent Required?	Yes. Informed consent strictly required before inputting sensitive data.	Generally no. Treats AI as a standard secure IT vendor.
Anonymization Requirement	Mandatory—must strip all identifying client details.	Highly recommended as a best practice.

Transparency and Informed Consent in the Age of Automation



Recent Conflicting Judicial Rulings on AI and Chatbots

U.S. v. Heppner (SDNY – Rakoff); *Warner v. Gilbarco* (E.D. Mich., Mag. Patti)

- **Different Outcomes:**

- Heppner – Information shared by client with chatbot and then communicated by client to lawyer not protected by the attorney-client privilege.
- Warner – Pro Se litigant’s communications with chatbot warranted work product protection.

- **Key Differences:**

- Status of the Litigant:
- In Heppner, the defendant had a lawyer but acted independently, which defeated his work product claim because the AI use did not reflect his lawyer's mental processes
- In Warner, the plaintiff was representing herself, meaning her own mental processes and AI prompts qualified for work product protection

.

- **View of AI Platforms:**

- Heppner: court viewed the AI platform as a third party whose data policies destroy confidentiality
- Warner: court explicitly defined AI programs as mere "tools," the use of which does not constitute a disclosure to an adversary that would waive work product protections



Supervision




Leadership's duty to oversee the machine.

Leadership's Absolute Duty to Supervise the Machine

Delegation Equivalency






First-Year Associate

-  Requires clear initial instructions.
-  Prone to misunderstanding complex nuance.
-  Requires iterative feedback and redlining.
- Partner owns the final output and liability.

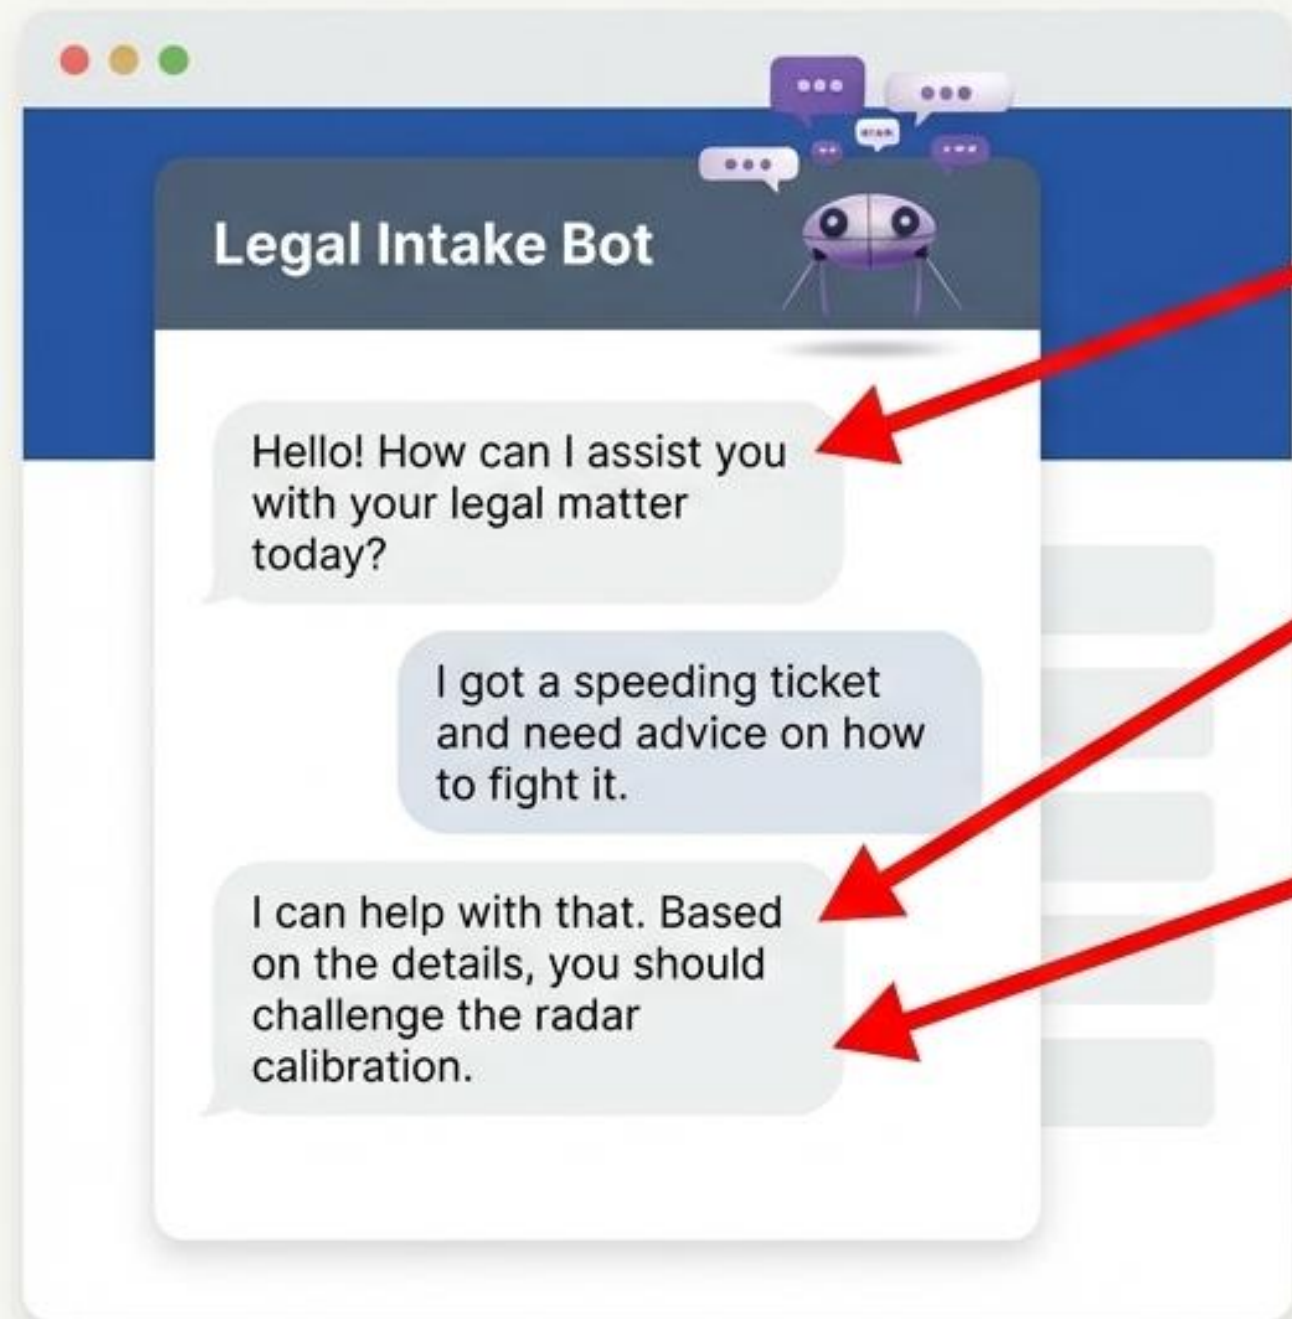


Generative AI

-  Requires precise “prompt engineering”.
-  Prone to confident hallucinations.
-  Requires iterative reprompting and verification.
- Partner owns the final output and liability.

Rules 5.1 & 5.3 Directive: Partners and managing attorneys are strictly liable for ensuring their firm adopts measures and policies that give reasonable assurance that AI is used ethically by all lawyers and non-lawyer staff.

The Chatbot Dilemma: Navigating Intake vs. Legal Advice



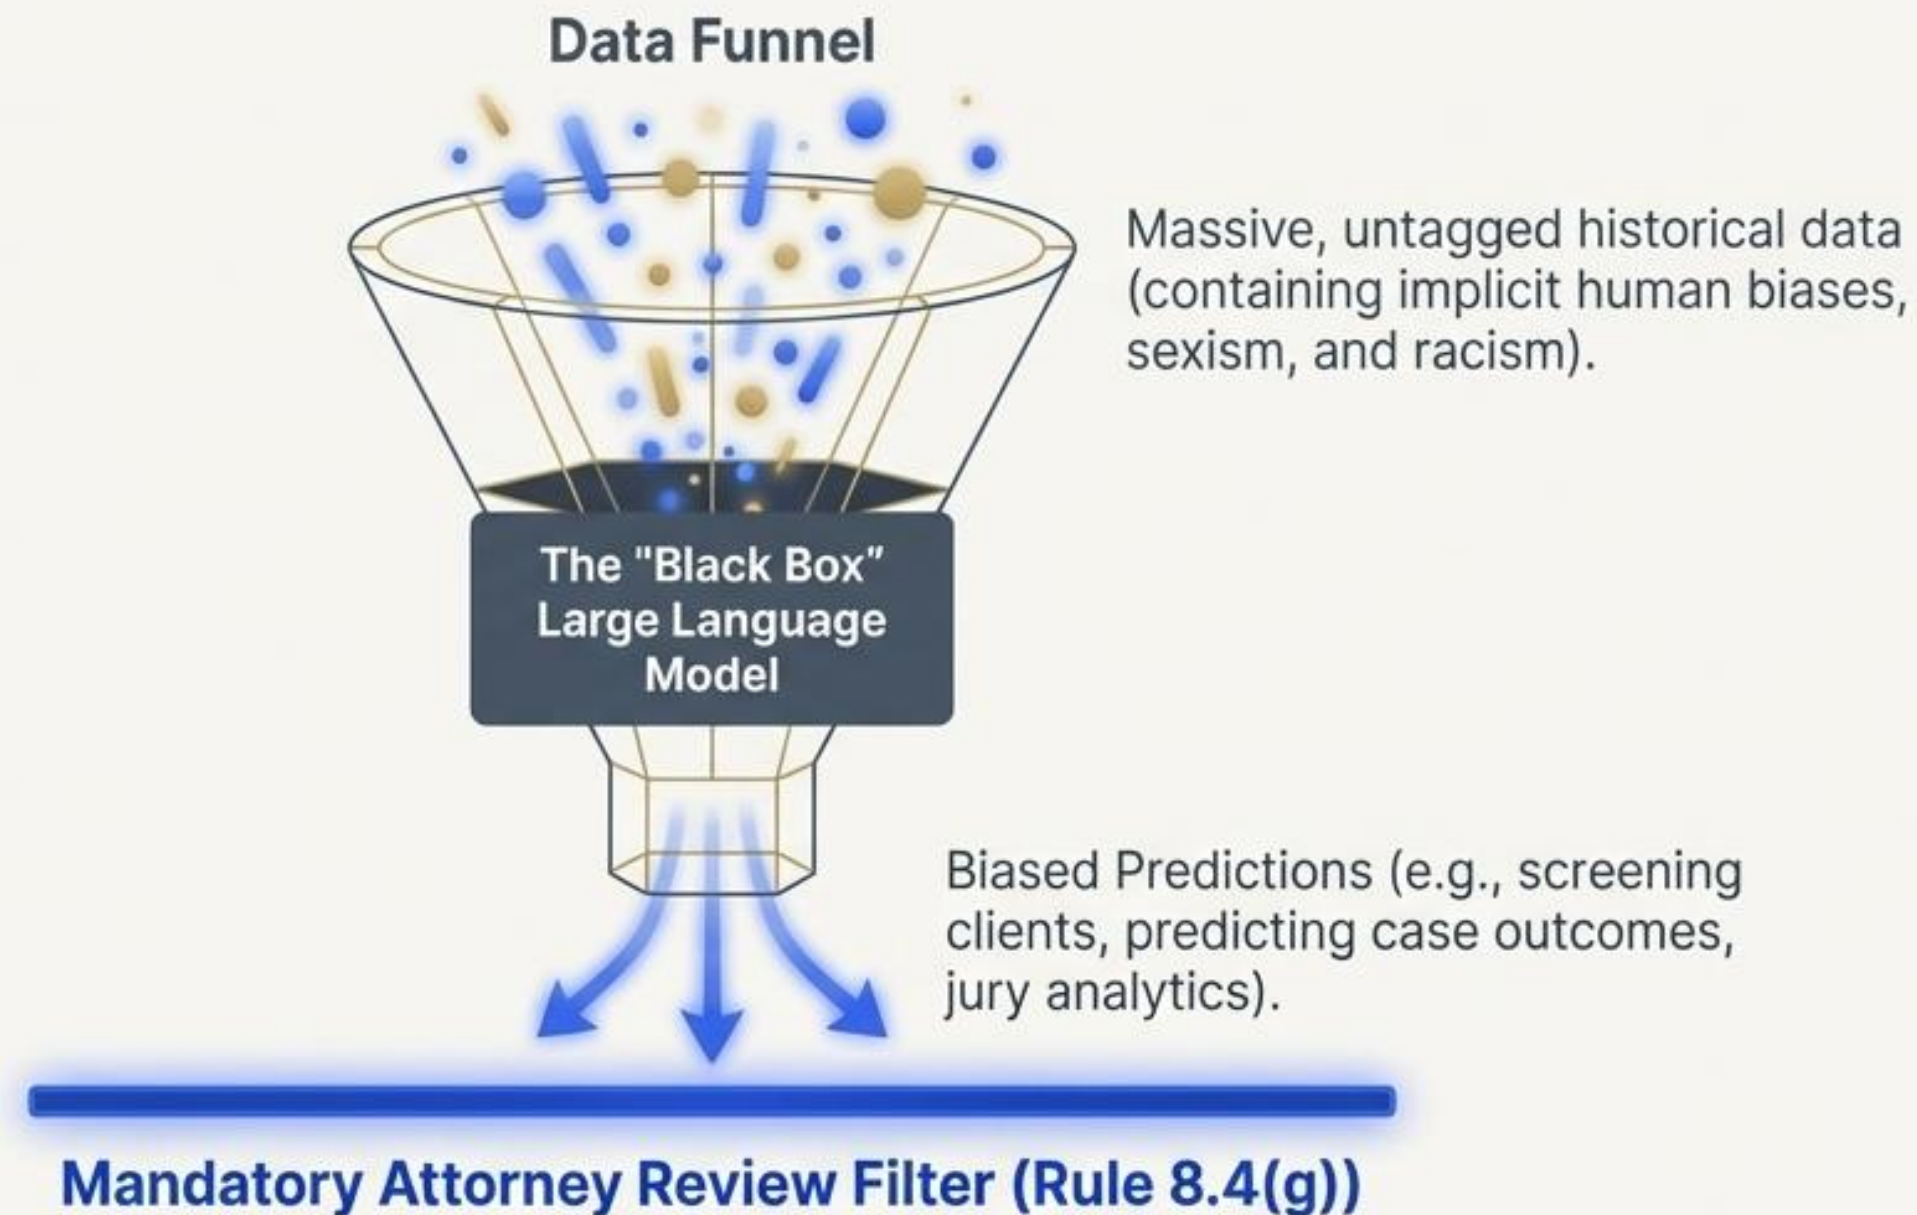
1. Identity Failure: The bot fails to immediately state it is an AI (Advertising violation).

2. The UPL Trap: The bot analyzes a user's specific traffic ticket scenario and suggests a defense strategy (Unauthorized Practice of Law).

3. Unintentional Formation: The bot uses overly welcoming language that accidentally creates a subjective, reasonable belief that a lawyer-client relationship has been formed.

Compliance Rule: Bots must strictly limit questions to obtaining factual intake information and must immediately refer any legal questions to a human attorney.

The Black Box: Inheriting the Internet's Prejudices



Lawyers must actively ensure that reliance on algorithmic tools does not manifest bias or prejudice in the administration of justice.

Synthesis: The 'Human-in-the-Loop' Operational Workflow



1 Sanitize (Rule 1.6)

Strip all Personally Identifiable Information (PII) and confidential data from the initial prompt.

2 Prompt (Rule 1.1)

Engage the AI with highly specific, bounded legal constraints to narrow the output.

3 Verify (Rules 1.1, 3.3, 8.4)

Critically review all outputs. Shepardize every citation. Screen for implicit bias. Correct all hallucinations.

4 Adopt


Apply human professional judgment to finalize the product. The human lawyer signs the document, not the machine.

The lawyer remains the ultimate gatekeeper of justice.



- Generative AI is a transformative lever for efficiency, but it cannot exercise legal judgment, swear an oath, or possess a moral compass.
- The adoption of new technology does not alter or lessen fundamental ethical obligations.
- The lawyer—armed with competence, candor, and independent judgment—stands as the final, irreplaceable firewall against the contamination of the legal system.

Day One: The AI Implementation Action Plan

- 
- Audit Technology:** Review the Terms of Service for all current AI tools regarding data retention and model training.
 - Establish Guardrails:** Draft and distribute a firm-wide Generative AI Acceptable Use Policy.
 - Update Contracts:** Revise standard client engagement letters to include appropriate AI disclosures and consent clauses.
 - Train the Team:** Mandate 'AI verification and hallucination' training for all partners, associates, and paralegals.
 - Monitor the Courts:** Regularly check local court standing orders regarding mandatory AI disclosure requirements in filings.

Technology changes rapidly; the duty to the client remains absolute. Trust, but rigorously verify.