

## Paste in Haste: The Fallout of AI Hallucinations in Court Filings and the New ARDC's Guide to Implementing AI

Mary F. Andreoni  
ARDC Ethics Education Senior Counsel

### Introduction

Every week brings a new headline: a lawyer or self-represented litigant is sanctioned for submitting filings containing “hallucinated” case law or statutes that were generated using Artificial Intelligence (“AI”). The issue first gained national attention in 2023 when a federal judge sanctioned two lawyers for submitting a brief containing citations to fictitious cases fabricated by ChatGPT—citations the lawyers failed to verify. What should have been an isolated occurrence turned out to be just the beginning of a continuing trend. Since 2023, researchers had identified an estimated 712 legal decisions worldwide involving hallucinated content, with about 90% occurring in 2025.

Courts have responded with standing orders and local rules, and bar associations have issued advisory opinions, including the American Bar Association (see ABA Formal Opinion 512 *Generative Artificial Intelligence Tools*, July 29, 2024). Despite this, the misuse of AI in court filings continues.

Relying on AI to draft motions or briefs without rigorously fact-checking the output risks more than embarrassment: it can result in sanctions, dismissal of claims, or disciplinary action. Long before AI arrived, we all learned in law school to track the history and treatment of each case cited to ensure that the case that we were relying on was still “good law.” The fundamentals of legal research haven’t changed, only the tools have. In a profession that requires precision and accuracy, the modern twist on an old proverb applies—paste in haste, repent at leisure. Even judges aren’t immune: two federal judges recently withdrew rulings after acknowledging that their staff used AI tools that produced fabricated citations.

This article explores why AI hallucinations happen, how to avoid them, and highlights the ARDC’s October 2025 release of *The Illinois Attorney’s Guide to Implementing AI*, a timely resource to help Illinois lawyers to better understand and use AI ethically and effectively in their law practice.

### The AI Citation Fallout

2023 became a pivotal moment in legal tech history. The first widely reported case of a lawyer sanctioned for citing fake AI-generated case law was *Mata v. Avianca, Inc.*, 678 F.Supp.3d 443 (S.D.N.Y. 2023). In that case, a lawyer submitted a legal brief that cited six fictitious cases. The opposing counsel flagged the citations as untraceable. When the judge requested verification, the lawyer reluctantly admitted that he had used AI in preparing the brief and, ironically, the lawyer even asked the AI program if the cases were real. Not surprisingly, the AI program confirmed that they were. The judge sanctioned the lawyer and his colleague, ordering them to pay a \$5,000 fine and mandating further legal education on the use of AI tools.

Rather than standing as a cautionary tale, the number of cases since *Mata* has only increased.

There's even a database maintained by a legal researcher that has cataloged over 700 global instances of AI hallucinations in court cases including fake citations. See [AI Hallucination Cases database](#), maintained by French lawyer and data scientist, Damien Charlotin.

While most cases so far have involved court-imposed sanctions, it was only a matter of time before the disciplinary authorities became involved. Recently, the Massachusetts Board of Bar Overseers publicly reprimanded a lawyer for submitting court pleadings containing fictitious case citations generated by an AI tool. This was in addition to the \$2,000 fine previously imposed in the related civil matter.

## Why AI Hallucinations Happen

Generative AI is only as reliable as the data it was trained on and the clarity of the prompt it receives. If the training data is outdated, biased, or incomplete or if the system misreads the user's intent, the output can be misleading or outright false.

To understand why these hallucinations happen, it helps to know how AI "thinks." AI doesn't possess factual knowledge in the traditional sense. It's trained to predict the next word based on patterns in massive datasets. That means it may invent citations or details that sound plausible but are entirely fabricated. Also, AI rarely admits uncertainty. Instead of saying "I don't know," it will typically generate an answer even if that answer is wrong. Finally, because the foundation of these systems is built on data that may be biased, incomplete, or outdated, especially in a constantly changing area like law, AI's outputs can reflect those flaws.

## New ARDC AI Guide to Using Generative AI

Despite growing awareness and stricter court response, many lawyers continue to fall into the trap of misusing generative AI. Avoiding these pitfalls begins with understanding how these tools work before applying them in practice.

To support that effort, the ARDC released *The Illinois Attorney's Guide to Implementing AI* (Oct. 2025), a practical resource for navigating the ethical use of AI in legal work. While tailored for solo and small firm lawyers, its insights apply to any lawyer or judge who is using, or considering using, AI in their legal work.

The *Guide* aligns with the Illinois Supreme Court's Policy on Artificial Intelligence, which permits AI use as long as lawyers uphold existing professional responsibilities. The Court's policy sets the foundation; the *Guide* focuses on how to integrate AI tools safely and ethically into legal practice. It explains how generative AI systems operate and presents a practical framework for assessing their appropriate use.

The framework centers on three essential steps:

- **Classifying the information being handled.** The *Guide* defines four categories of sensitivity of information, from general (non-confidential data entirely unrelated to any client matter) to sensitive personal data (financial, health, and other legally protected information).
- **Assessing the security level of the AI tool.** The *Guide* categorizes AI tools into four security levels based on eight AI safeguards outlined in the *Guide*, from public (minimal to no data protection); consumer-grade (some protection from basic controls like opt-outs from model training); business-class (stronger safeguards, though may lack advanced administrative features); and enterprise (highest protection across all safeguards). Included in the *Guide* are detailed explanations and checklists to help lawyers classify virtually any AI tool they are likely to encounter.
- **Aligning data sensitivity with tool security.** The *Guide* includes a decision matrix that helps lawyers match the level of the sensitivity of the data with the appropriate AI tool. Confidential data should never be processed using public AI tools, even with client consent. Business-class or enterprise tools may be acceptable if clients are informed and can opt out.

To support implementation, the *Guide* offers a Practice Resource Kit with checklists, sample policies, and communication strategies to help lawyers explain AI use to clients transparently. The *Guide* is a must-read for any legal professional looking to use AI wisely—and ethically. To download a copy of the *Guide*, go to the ARDC website ([www.iardc.org](http://www.iardc.org)).

## Best Practices for Verifying AI-produced Citations

Whether you're using AI to brainstorm arguments or draft entire briefs, every cited authority must be real, relevant, and accurately represented.

The following best practices can help ensure your AI-assisted citations are court-ready, ethically sound, and professionally defensible:

- **Cross-check with trusted legal databases:** Always verify citations using authoritative sources like Westlaw, LexisNexis, Bloomberg Law, or PACER. If a case or statute doesn't appear in these databases, it likely doesn't exist.
- **Use citation-verification tools:** Tools like Lexis+ AI's Protégé or Bloomberg Law's Brief Analyzer are designed to flag hallucinated citations and link to verified sources.
- **Read the full text of the case:** Don't rely on summaries or quoted excerpts alone. Review the full judicial opinion to confirm the quote is accurate, the content supports your argument, and the case is still good law.
- **Verify every citation individually.** Even if the first few AI-generated citations check out, verify *every* one individually.

- **Keep a verification log.** Maintain a brief record of where and how each citation was verified. This documentation can be critical if your sources are challenged by opposing counsel or the court.
- **Stay current on court rules.** Review local rules and standing orders before submitting AI-assisted documents. Many courts now require disclosure of AI use in filings or mandate verification of all cited authorities.
- **Educate your team:** Establish clear protocols for verification before any AI-assisted work is filed. Ensure paralegals, clerks, and associates understand the risks of AI-generated citations.
- **Keep up with ethical guidance:** The legal profession is rapidly evolving in response to AI. The ARDC's *Guide* offers a framework for responsible use, including recommendations on disclosure, verification, and client communication.
- **Practice using AI.** The best way to understand how AI works is to interact with it regularly. Ask it questions you already know the answers to. Test its limits with simple, silly, complex, or even philosophical prompts. The more you experiment, the better you'll grasp its strengths and its weaknesses.

The bottom line is to treat AI output like a first draft from a junior associate or law clerk — potentially useful but never ready for submission until you've personally verified every fact, quote, and citation.

## **Ethical Considerations in Using AI in Court Filings**

Using AI in legal work isn't just a matter of efficiency, it also requires careful attention to the ethical responsibilities set out in the Illinois Rules of Professional Conduct.

The duty of competence (Rule 1.1) requires lawyers to understand how AI tools function, including their limitations and potential for error. The duty of candor (Rule 3.3) obligates lawyers to personally verify that every cited authority exists, is accurate, and supports the intended argument.

The duty of supervision (Rules 5.1 and 5.3) requires senior lawyers to ensure that staff and subordinate lawyers using AI tools are doing so ethically and responsibly. And the duty to make meritorious claims (Rule 3.1) requires lawyers to conduct a reasonable inquiry that what is being filed has a good-faith basis in law and fact.

At the end of the day, lawyers are still accountable for the work they submit. AI may assist with drafting or research, but it doesn't absolve a lawyer of their professional obligations.

## **Long-term Risks of Overreliance on AI**

Overreliance on AI in legal practice can pose some serious long-term risks. One of the most concerning is the creation of fictitious case law. When AI tools generate briefs with fabricated

citations, those errors can slip into the legal record, potentially influencing future rulings and undermining the integrity of the judicial system.

Another growing issue is the rise of “workslop” or “AI slop”—AI-generated content that looks polished but lacks substance. These outputs often require human intervention to correct, refine, or completely redo, wasting time instead of saving it, and resulting in low-quality work product.

Poorly generated legal documents can also harm clients. Inaccuracies or omissions introduced by AI may lead to unfavorable rulings, procedural errors, or even sanctions.

Finally, as more self-represented litigants turn to AI for help, courts might be faced with added strain managing flawed filings, thereby complicating case resolution, and stretching already limited resources.

## Conclusion

From lawyers and self-represented litigants to judges, overreliance on AI and a lack of scrutiny around its outputs are leading to consequences that go far beyond a single brief or sanction. The legal system depends on accuracy, precedent, and trust and when AI-generated content undermines those pillars, it threatens the integrity of the justice process itself.

The path forward isn’t to reject AI but to use it wisely. As AI becomes more integrated into how lawyers write, research, and create, it’s essential to remember that the lawyer, not AI, is responsible for understanding the client’s objectives, advocating for their best interests, and making ethical decisions.

One of the most iconic lines from the original *Star Trek* series comes from Mr. Spock, the embodiment of logic, in an episode where a computer designed to replace the ship’s crew goes rogue and seizes control of the *Enterprise*: “Computers make excellent and efficient servants, but I have no wish to serve under them.” That episode is as relevant today as it was when it first aired in 1968. AI can assist and occasionally inspire, but no matter how efficient or precise these programs become, they must remain tools of human will—not masters of it.