



When AI ‘Lies’: The Legal Risks of Hallucinations

The legal profession is built on accuracy, precedent, and trust. When AI hallucinations infiltrate legal work, it threatens to undermine these foundations in several ways.

BY BRENT J. HOEFT

Earlier this year Judge Scott Schlegel, a Louisiana State Appellate Court judge, warned of the high stakes “if a judge misuses AI. Judges are the ultimate arbiters of justice. Their decisions, unlike a single pleading or brief, carry the full weight of the law.”¹ The concern is no longer hypothetical. It was bound to happen. After the unending cases of lawyers submitting pleadings and briefs containing hallucinations to courts across the country, last month we saw judges fall victim to generative artificial intelligence (GenAI) hallucinations. These alarming incidents highlight the evolving issue of GenAI hallucinations in the legal industry.

Recent Court Incidents Highlight the Danger

One of the earliest and most notorious instances of hallucinations in a filing to a court was the *Mata v. Avianca* case in May 2023, in which the attorney, relying on ChatGPT, cited six nonexistent cases in a court brief.² The submission included detailed references to fake decisions complete with names, citations, and even a summary of precedent related to airline disputes. The opposing counsel and the judge couldn’t find the cases anywhere. The judge reprimanded the lawyers and later fined them for violating their duty of candor. Despite this early cautionary tale, attorneys submitting court filings containing hallucinations have continued. Now, GenAI hallucinations have infiltrated court proceedings from the bench’s side.

In *Shahid v. Esaam*, a Georgia divorce case, it wasn’t only a filing but also a court’s order that had AI hallucinations.³ The husband’s attorney submitted to the court a proposed order that contained numerous hallucinated case citations. The trial judge signed the proposed order, thereby adopting two fake precedents as part of

the judge’s reasoning. On appeal the issue was discovered and the wife’s attorney argued that the trial order was “void on its face” because it relied on cases that “do not exist.” On appeal the husband’s lawyer doubled down, citing 11 more nonexistent or irrelevant cases in their brief, including one to support a request for attorney fees. The Georgia Court of Appeals vacated the trial court’s decision entirely once it confirmed the citations were made up. Appellate Judge Jeffrey A. Watkins noted the filing irregularities “suggest that they were drafted using generative AI” and sanctioned the husband’s attorney.

In a separate instance of a judge citing hallucinations in an opinion, Judge Julien Neals of the U.S. District Court for the District of New Jersey had issued a written opinion on a request to dismiss a lawsuit in a securities case.⁴ Attorneys pointed out major errors in the opinion. The judge in his ruling had denied a motion to dismiss, but



Brent J. Hoeft, Cleveland State Univ. College of Law 2006, is the State Bar of Wisconsin’s practice management advisor and manager of the Practice411™ practice management program. If you have questions about technology, practice management, or the business aspects of your practice, call (800) 957-4670 or email practicehelp@wisbar.org. Access the digital article at www.wisbar.org/wl.



his written opinion contained strange mistakes, including quotations attributed to cases that didn't appear in those cases, references to a wrong court, and descriptions of precedent that were plainly incorrect. The lawyers in a related case filed a letter detailing these issues, effectively alerting the judge that parts of his analysis looked fictitious. Within days, Judge Neals took the extraordinary step of withdrawing his own opinion. The incident was widely reported as a rare example of a court having to retract a decision due to possible GenAI hallucinations.

Defining Hallucinations

Generative AI systems (for example ChatGPT, Google Gemini, Anthropic's Claude, and other large language models) create predictive text-based outputs arising from patterns in training data. When these models do not know an answer, they can and do make up answers that sound plausible to fill in the blanks or gaps in data. In AI lingo, this is called a "hallucination," meaning the AI has confidently generated information that is false or nonexistent. In the legal context, hallucinations typically take on a few different forms: phantom cases and courts, invented quotations or attribution, and misstated holdings or law.

Phantom Cases and Courts

The GenAI model might produce an official looking case citation that is completely made up. Often these fake citations mix party names that sound real with an incorrect court or docket number. Sometimes the court or judge doesn't exist. If a lawyer does not catch it, such phantom cases or precedents can make their way into filings.

Most commonly, lawyers associate hallucinations with this type of phantom case and court hallucination, probably because of the *Mata v. Avianca* case discussed above. This is also the type of hallucination that the legal research companies claim to be able to eliminate by using their generative AI products, which are bounded by their databases

containing credible legal documents. While that is true to an extent, made-up cases and courts are not the only kind of hallucination. Other harder to detect hallucinations are also common and are crucial for attorneys to be aware of and understand.

Invented Quotations or Misattribution

The GenAI tool might output a fabricated quotation, attributing it to a real case or judge. This is harder to spot because the case itself might be real, but the quoted language is not found in the opinion. As noted above, U.S. District Judge Julien Neals withdrew an opinion after people pointed out that it contained multiple quotations that were attributed to cases that did not contain the quoted passages. Misattribution can also mean citing the right case but saying it was from the wrong court or year. This also happened in Judge Neals' opinion, in which he referenced a real case from a New Jersey court but indicated it was from a New York court.

Misstated Holdings or Law

Sometimes GenAI output describes a real case or law incorrectly. It might claim a case held the opposite of what it truly held or it might confuse the facts. In Judge Neals' withdrawn opinion, in addition to the fake quotations, several case outcomes were mischaracterized. The opinion cited precedents supposedly supporting the plaintiffs, but some of those cases had ruled the opposite and in favor of defendants. In other examples, a GenAI tool might miss the narrow scope of a particular holding and summarize it more generally to match more closely to the user prompt. These errors are less immediately obvious than a completely fake case or court and can be equally damaging because they distort the legal analysis.

GenAI models are designed to produce fluent text based on predictive analysis, not to verify facts. If prompted with "I need you to give me cases that

support X position," the AI will comply by outputting convincing text learned from patterns in real cases, but it may invent details to best fit the request. Unless specifically integrated, public-facing GenAI models have no built-in legal database to cross-check citations, so if it "thinks" a citation looks plausible, it will present it confidently. Without human verification, these imaginary legal authorities can pass as the real thing.

Effect of Hallucinations in the Legal System

The legal profession is built on accuracy, precedent, and trust. When AI hallucinations infiltrate legal work, these foundations might be undermined in several ways:

- **Erosion of court trust and integrity.** Every time a fake case or quotation makes its way into a brief or a judgment, it chips away at the court's confidence. Judges might start doubting all submissions, which could hurt the credibility of diligent attorneys. And court orders that themselves contain hallucinations risk creating a perception of the breakdown of judicial integrity.

- **Wasted resources and delays.** Hallucinations force courts and opposing parties to spend time and money substantiating every assertion. These incidents clog up dockets and distract from the merits of cases.

- **Sanctions, malpractice, and professional liability.** Courts have shown they will sanction attorneys for AI-generated mistakes, refer them to disciplinary authorities, and even suspend the attorneys from the case. At a minimum, even if not formally sanctioned, the lawyer suffers reputational harm from being the subject of a published court opinion or news story about their lack of oversight and failure to verify the contents of their work product.

- **Client trust and relationships.** Clients expect their attorneys to use the right tools for their matter, but they also expect due diligence and zealous representation. If a client learns that



their lawyer submitted work filled with errors drafted by GenAI, their confidence in the attorney will plummet as will the attorney's reputation.

Understanding GenAI hallucinations and the different ways in which they can occur is part of a lawyer's duty of technical competence. Hallucinations are a substantial risk when using generative AI. Ignoring this risk converts GenAI from a powerful tool to a potential liability.

How Lawyers (and Judges) Can Avoid Hallucination Pitfalls

GenAI is here to stay, offering potential to streamline legal research and drafting. However, attorneys and judges must use it wisely and ethically. Here are some recommended best practices:

- **Develop protocols to double-check every citation and quote.** Always verify AI-generated case citations and quotations using reputable sources like Westlaw, Lexis, vLex, or official reporters. If you use it, you need to verify it. If your legal research process begins with only using GenAI, require all results to be independently verified. There are products now on the market that aid in the verification process, but they are unlikely to catch all forms of hallucinations discussed.⁵ So, although these tools can be an effective step in the verification process, they should not be relied upon as the sole means of verification.

- **Use reliable tools and demand sources.** As of June 2025, there are at least 638 legal-specific GenAI tools currently on the market, according

to a tracker from Legaltech Hub.⁶ Not all GenAI legal tools are equal. Legal-specific AI platforms integrated into Westlaw, Lexis, or vLex provide citations for every statement and reduce, but do not eliminate, the risk of hallucinations. General AI models like ChatGPT are more likely to fabricate information because of their broad access to the internet and information. If using general AI tools, always prompt them to show sources and then follow up and verify them.

- **Maintain human oversight and judgment.** AI is useful for first drafts or to aid in review, but human lawyers must be the ultimate editors. Review AI-generated text critically, just as you would an intern's or other assistant's work. If something looks off, stop and investigate it further. Supervising attorneys must be aware of when GenAI has been used in the creation of work product and must make sure that it has been verified.

- **Stay educated and follow emerging guidance.** The ABA's Formal Opinion 512 and other ethics opinions emphasize that using GenAI requires technical competence and adherence to ethical duties. Lawyers must understand AI's benefits and risks, including the potential for hallucinations.

- **Be candid and correct mistakes promptly.** If an oversight occurs and a hallucination is missed, inform the court and opposing counsel immediately and correct the record.

- **Follow the Judicial Guidelines from the Judicial Conference if you are a judge.**⁷ These guidelines offer a

framework for the ethical and effective use of AI. Emphasis should be placed on judicial independence and ensuring verification of all AI outputs.

The Takeaway

This is the worst GenAI will be. Going forward it will only get better and one day the hallucination problems might be fixed. But today is not that day. The potential for GenAI in the legal field is undeniable. It is also undeniable that failing to understand the risks posed by improper use of GenAI can result in dire consequences. Not verifying GenAI outputs and having a fake case or quotation getting into a court filing or order is not just an unfortunate error; it has the power to mislead, alter the course of a case, and result in sanctions or miscarriages of justice.

Lawyers and judges must approach GenAI tools cautiously and be aware that these systems, however impressive, do not *know* the law. GenAI can, and does, just make things up. By combining technological competence with good, old-fashioned legal diligence, attorneys can harness GenAI's benefits while avoiding its risks and pitfalls. In practice, this means no citation goes unchecked, no quote goes unverified, and the lawyer stays firmly in control of the work product. **WL**

ENDNOTES

¹Hon. Scott Schlegel, *The Higher Stakes of AI Misuse in the Judiciary*, Legal Tech, May 13, 2025, <https://substack.com/home/post/p-163470559>.

²*Mata v. Avianca Inc.*, 678 F. Supp. 3d 443 (2023), <https://national.claa.org/wp-content/uploads/2025/05/Mata-v-Avianca-Inc.pdf>.

³Joe Patrice, *Trial Court Decides Case Based On AI-Hallucinated Caselaw*, Above the Law, July 1, 2025, <https://abovethelaw.com/2025/07/trial-court-decides-case-based-on-ai-hallucinated-caselaw/>; see also *Shahid v. Esam*, 2025 Ga. App. LEXIS 299 *, <https://caselaw.findlaw.com/court/ga-court-of-appeals/117442275.html>.

⁴Justin Henry, *Judge Scraps Opinion After Lawyer Flags Made-Up Quotes*, Bloomberg L., July 24, 2025, <https://news.bloomberglaw.com/business-and-practice/judge-withdraws-pharma-opinion->

[after-lawyer-flags-made-up-quotes/](https://news.bloomberglaw.com/business-and-practice/judge-withdraws-pharma-opinion-after-lawyer-flags-made-up-quotes/); *In Re CorMedix Inc. Sec. Litig.*, No. 2:21-cv-14020 (D.N.J. July 22, 2025).

⁵See, e.g., Cite Check AI, <https://citecheck.ai/>; Clearbrief, <https://clearbrief.com/>.

⁶LTH GenAI Legal Tech Map: June 2025, Legaltech Hub, <https://www.legaltechnologyhub.com/contents/lth-genai-legal-tech-map-june-2025/>.

⁷Hon. Dixon Jr., Hon. Allison H. Goddard, Maura R. Grossman, Hon. Xavier Rodriguez, Hon. Scott U. Schlegel & Hon. Samuel A. Thumma, *Navigating AI in the Judiciary: New Guidelines for Judges and Their Chambers*, 26 Sedona Conf. J. 1 (forthcoming 2025), https://www.thesedonaconference.org/sites/default/files/meeting_paper/7.1%2520Navigating%2520AI%2520in%2520Judiciary_O.pdf. **WL**

