The Misuse and Misunderstanding of Generative Artificial Intelligence in Litigation

Dan D. Kohane[1]
Michael J. Williams[2]
Hurwitz Fine P.C., Buffalo, NY

The use of artificial intelligence must be accompanied by the application of actual intelligence in its execution.[3]

Generative artificial intelligence ("AI") is the latest innovative weapon in a litigator's arsenal to zealously and efficiently represent their clients. While useful as a research and drafting aid when properly used, overreliance on this developing technology becomes a double-edged sword to the neglectful practitioner. Electronically fabricated or deceptive case citations, quotations or facts, commonly referred to as "hallucinations", present an ongoing danger to our courts and the efficient administration of justice. Courts have responded by emphasizing that attorneys—not their computers—retain nondelegable and longstanding duties of reasonable inquiry as to all data presented and for candor to the tribunal should error occur. This article examines the growing problem, how a select subset of courts applies these obligations to the emerging generative AI technology and why ignorance of the hallucination danger is no defense.

Prevalence and Pitfalls of Generative AI Models

A 2024 LexisNexis survey reported that while 80% of responding Fortune 1000 executives expect that generative AI will reduce their billing from outside counsel, 78% of *legal* executives identified hallucinations as a significant obstacle to AI adoption. The concern is justified: a 2025 article in the Journal of Empirical Legal Studies found that LexisNexis' Lexis+ AI was the highest performing system while accurately answering only 65% of the researchers' queries. Westlaw's AI-Assisted Research reportedly boasted only a 42% accuracy rate and hallucinated at twice the rate of the other systems tested. No executive should wager on bet-the-corporation litigation against those odds.

Hallucinations occur in many unexpected forms. While completely manufactured case citations or quotations are the most frequently reported, they also manifest as false factual data on which analysis relies, applying out of context data inappropriately, illogical conclusions based on otherwise factual data, or responses accurate to the user's prompts which are distorted or unresponsive to the user's intent. All of these result from generative AI's imperfect deep-learning models, which are algorithms that imperfectly mimic human learning and decision-making.

---

[1] Dan D. Kohane is a Senior Member at Hurwitz Fine P.C. where he chairs the firm's Insurance Coverage and Extracontractual Liability teams. He is has served, for almost four decades, as Adjunct Professor of Insurance Law at the University at Buffalo School of Law. He was one of the original Founding Fellows of the American College of Coverage Counsel

[2] Michael J. Williams is a Member at Hurwitz Fine P.C. where he leads the firms Child Victims Act & Sexual Misconduct Coverage and Defense teams , as well as its Indian Law & Policy Team. He is an active litigator, focusing his practice or general liability including products liability, prescription medications and medical device claims.

[3] *Willis v. U.S. Bank N.A.*, 783 F. Supp. 3d 959, 959 (USDC, ND Texas (2025)

Those algorithms became more active with development of Large Language Models which gather tremendous amounts of data—from the internet complete with its misinformation and biases-- to predict the next element in a sequence to generate content. Stated differently, the model guesses at what the next appropriate word should be when generating a response with no ability to ascertain the truth or falsity of the prediction. This process is entirely invisible to the user-attorney, and any predicative misstep can generate facially plausible but completely erroneous results.

Industry-standard providers such as LexisNexis and Westlaw engage in constant development to improve performance. Recently, Retrieval-Augmented Generation—effectively limiting the universe of data applied to generate contextually appropriate and more accurate responses—have been boasted to reduce hallucinations. However, even these more sophisticated models cannot ascertain whether a source states a definitive rule of law or one hotly contested, particularly where courts themselves disagree as to the correct answer and the model applies only a miniscule subsection of the available data. The model simply lacks an attorney's specialized background knowledge in the subject matter and ability to synthesize relevant data within the context of fact patterns, rules of law, and resulting legal holdings across a myriad of jurisdictions and over time. The fact that our legal search engines present their generative AI responses as if they were as reliable as reported case law is a trap for the unwary or inexperienced practitioner.

A Singular Case Study

Both attorneys and generative AI models are imperfect. Many briefs contain scrivener's errors misreporting a case name or citation, omitting a critical word or misplacing a quotation mark. In a recent Southern District of New York case, *Jiminez-Fogarty v. Fogarty*, 1:24-cv-08705 (April 29, 2026), the court generously distinguished such errors by excluding instances of erroneous volume and page citations, completely incorrect case names, and even references addressing the general issue but presented in a "completely mischaracterized" or "grossly misleading" manner. The court detailed the egregious errors it permitted, then itemized seven "completely made up" citations, and detailed its own efforts attempting to identify any basis whatsoever for the offending attorney's references.

Rather than proceeding on its own findings, the court ordered the offending attorney an opportunity to show cause why they should not be sanctioned under Fed. R. Civ. P. 16(f), 28 U.S.C. § 1927, or the Court's inherent power, including requesting "a complete and detailed description of the process of the drafting of the two memoranda of law" and "a detailed and complete description of the role of any individual or computer system that was involved in the drafting process." While the offending attorney had clearly violated their duty of reasonable inquiry as to the arguments presented, they chose to additionally fail in their duty of candor before the court by responding to the itemized hallucinations with "[c]ase number exists; incorrect case name" or similar responses, despite these being untrue.

Opposing counsel then moved the court for Fed. R. Civ. P. 11 sanctions including attorneys' fees and costs. The court then filed a second order to show cause and again received a completely inadequate response. The Southern District lamented that while "most attorneys" candidly explain their error including admitting their reliance on AI platforms, the offending attorney offered "only

airy generalities and conclusory statements" while failing to answer the most basic question: "what was the source of the fabricated citations?" The court also found the offending attorney's assurance of manually checking each citation to be "obviously false" because that surely would have revealed the *eleven* errors to which counsel themselves admitted.

The offending attorney also blamed their errors on using "established legal research platforms", to which the court accurately responded that the proposition was "utterly devoid of evidence". While the court offered no evidence supporting its own reliance on the databases used to check the offending attorney's work, the reliability of official reporters remains the time-tested backbone of legal practice. LexisNexis and Westlaw generative AI platforms may produce hallucinated results as discussed above, but every court and practitioner relies on their industry-leading and accurate reporting of actual case law.

Having given the offending attorney a generous reading of their brief, and two opportunities to explain their errors, the court held Rule 11 sanctions appropriate for failing to conduct a reasonable inquiry of their own case citations, their bad faith unresponsiveness to the court and because the same attorney continued to file documents featuring hallucinated cases including before the sanctioning court. Ultimately, the court ordered a $2,500 fine and notice to the client to all other courts in which the offending attorney was appearing that included a copy of the Southern District's *twenty-page opinion* detailing their misconduct.

Hallucinated Case Citations Constitute Recklessness and Not Honest Error

While *Jiminez-Fogarty* is among the latest examples of hallucinatory error by counsel, litigators must recognize that the epidemic is no longer limited to mostly unpublished trial court opinions of which few would know other than the parties to the litigation. Appellate courts have also written on sanctions issued by lower courts, or for misconduct by hallucination on the appeals themselves. Every practitioner is accordingly on notice not just from articles in periodicals such as this one, but from the courts with authority to affirm sanctions levied against them or their clients. Court systems and bar associations have also proactively taken steps to remind their practitioners of their duties.

New York provides one of the most recent examples with 22 NYCRR §§ 161.1 to 161.4 (effective June 1, 2026). The Unified Court System undertook an extensive review of generative AI use in legal briefing, accepted public comments, and promulgated rules for those litigating on the electronic frontier. These can be boiled down to three policy decisions: (1) the use of AI in papers submitted to the court "should not be prohibited", (2) attorneys "should not be required" to disclose to the court if they used AI tools, and (3) individual judges may exercise their discretion to implement a rule requiring attorneys using AI to "carefully review the paper and independently ensure" it contains no factual or legal hallucinations. However, existing rules of professional responsibility already require attorneys to conduct a reasonable inquiry as to all data presented to the court--regardless of the tools they employ to draft their briefing--resulting in these rules covering little new ground.

Turning to the federal appellate courts finds a similar approach. The Second Circuit Court of Appeals has held that the Federal Rule of Civil Procedure Rule11 "reasonably inquiry" standard

requires "that attorneys read, and thereby confirm the existence and validity of, the legal authorities on which they rely," as there is "no other way to ensure that the arguments made based on those authorities are 'warranted by existing law,' . . . or otherwise 'legally tenable.'" *Park v. Kim*, 91 F.4th 610, 614-615 (2nd Cir. 2024) (quoting *Cooter & Gell v. Hartmarx Corp*., 496 U.S. 384, 393 (1990)). When an attorney relies on fabricated authority, they necessarily reveal that they "made no inquiry, much less the reasonable inquiry required by Rule 11." *Park*, 91 F.4th at 615; *see also Ioio v. City of New York*, 2026 U.S. App. LEXIS 7593, *7 n. 7 (2nd Cir. Mar. 16, 2026). In other words, standards set decades ago govern new technologies such as generative AI.

The same standard applies for briefing presented to the federal Courts of Appeals themselves. In *Fletcher v. Experian Info. Sols., Inc*., 168 F.4th 231 (5th Cir. 2026), the court applied Fed. R. App. P. 38--Rule 11's appellate equivalent—with an analysis that began by documenting hallucinogenic proliferation. The *Fletcher* court referenced a database assembled by a lawyer and data scientist who tracked 239 cases of hallucinations by attorneys in the United States by the date the opinion was drafted. *Id*. at 234. Only four months later, at the time of this article's submission for publication, that list stands at 372. For those litigating against *pro se* parties, an additional 602 instances are presently identified. As the Fifth Circuit stated: "It is a problem that is getting worse—not better." *Fletcher*, 168 F.4th at 234-235; *see also Ford v. Bank of N.Y. Mellon*, 2025 U.S. App. LEXIS 7985, *2 n. 1 (5th Cir. Apr. 4, 2026).

Fletcher's appellate reply brief itself presented "16 instances of fabricated quotations and 5 additional serious misrepresentations of law or fact". The offending attorney responded to the order show asserting that they "relied on publicly available versions of the cases, which [they] believed were accurate." The Fifth Circuit then researched for itself the alleged sources (Google Scholar, CourtListener, Justia, and FindLaw) and found that none of them included the hallucinated quotations. The offending attorney's other two reported sources, Casetext and vLex, were both generative AI products. Counsel further compounded their error by failing to candidly admit their misapplication of generative AI in their briefing. The court ultimately cautioned practitioners that admitting error would have likely limited sanctions, while misleading, evading and violating duties as an officer of the court resulted in a personal $2,500 sanction.

Relying on *Fletcher*, *Park* and other opinions, the Fourth Circuit has directly expressed the sentiment all practitioners should respect: sanctionable misconduct "applies to submitting a brief with nonexistent cases no matter how it is done, whether through generative AI or not." *In re Nwaubani*, 2026 U.S. App. LEXIS 7501 (4th Cir. Mar. 11, 2026). The Sixth Circuit agrees, holding that "[a]ttorneys have an ethical obligation to verify the citations and propositions they submit to courts; that obligation reflects duties of competence and candor that apply no matter the tools attorneys use." *United States v. Farris*, 2026 U.S. App. LEXIS 9704 (6th Cir. Apr. 3, 2026). Notably in *Farris*, the offending attorney was not monetarily sanctioned after promptly admitting their erroneous reliance on unverified generative AI provided by non-attorney staff, reminding all practitioners that candor counts.

The Sixth Circuit also provides a cautionary tale to offending attorneys who fail both their duties of reasonable inquiry and candor. In *Whiting v. City of Athens*, 2026 U.S. App. LEXIS 7479 (6th Cir. Mar. 13, 2026), the court reviewed briefing containing two dozen cases that "did not exist, did not include the quoted language [counsel] claimed they did, or did not discuss or support the

proposition for which" they were cited. Invited to explain their errors, the offending attorneys instead argued that the order to show cause was void, motivated by harassment towards them and reflected illegal *ex parte* communications with the court. The court rejected these arguments along with counsel's position that revealing "how the briefs were cite-checked" violates the work-product and attorney-client-privilege doctrines. Under both Rule 38 and the court's inherent authority, the Sixth Circuit found the two unrepentant offending attorneys had abused the adversary system by relying on "fake opinions", sanctioned them $15,000 each, required them to pay their opponent's attorneys fees and double their costs, awhile also referring them for disciplinary proceedings.

The Tenth Circuit has recently joined the federal appellate courts delving into generative AI hallucinations. In *Picon-Diaz v, Bondi*, 2026 U.S. App. LEXIS 4417 (10th Cir. Feb. 13, 2026), the court identified offending counsel as relying "multiple times throughout his brief on what clearly appears to be a case fabricated by the use of a generative artificial intelligence tool." Rather than sanction the error, the court instead chose to warn counsel "and all attorneys practicing before this court" of their responsibility to present only "real cases that contain quotations attributed to them and arguably stand for the propositions for which they are cited." *Id*. at *14; *see also Dodds v. Bridges*, 2026 U.S. App. LEXIS 4159 (10th Cir. Feb. 11, 2026) (generalized warning to *pro se* litigants and counsel alike).

Contemporaneously, in *Amarsingh v. Frontier Airlines, Inc*., 2026 U.S. App. LEXIS 3928 (10th Cir. Feb. 9, 2026), the Tenth Circuit fined offending counsel $1,000 for presenting hallucinated cases and quotations while blaming "an embedded algorithm" in ChatGPT and admitting they failed to verify the data. The court declined the relief requested by offending counsel's opponent: dismissal of the appeal. Though finding counsel's conduct "reckless" for failing to confirm their citations and quotations, the court was persuaded by remorse, somewhat qualified candor and the perceived "novelty at the time of the infraction (January 24, 2025) of the use of AI by lawyers". *Id*. at *19. Considering that *Park* was decided on January 30, 2024, and relies on an even earlier opinion addressing the issue, offending counsel in *Amarsingh* may be among the last permitted a "novelty" exception to longstanding rules of reasonable inquiry merely because the error results from AI hallucinations.

Conclusion

Clients, carriers, and law firms rightfully expect counsel to use the most efficient tools available to deliver high-quality, cost-effective services, including innovations such as generative AI. But the attorney using the tool remains responsible for the result. As one court has aptly put it, "the use of artificial intelligence must be accompanied by the application of actual intelligence in its execution." *Mid Cent. Operating Eng'rs Health & Welfare Fund v. Hoosiervac LLC*, No. 2:24-CV-326, 2025 U.S. Dist. LEXIS 31073, 2025 WL 574234, at *4 (S.D. Ind. Feb. 21, 2025), *report and recommendation adopted as modified*, 2025 U.S. Dist. LEXIS 100748, 2025 WL 1511211 (S.D. Ind. May 28, 2025). Attorneys must always conduct their reasonable inquiry into the facts and law presented to all courts before which they appear. In the event of error, they must protect their client by candidly responding to a court's inquiry. No attorney should ever cost a client their legal rights for something as elementary as failing to verify their case citations.