

AI Hallucinations: What Lawyers Need to Understand

Like humans, artificial intelligence (AI) tools sometimes make things up. Lawyers must understand that hallucinations are common and that the ramifications of missing a hallucination – potentially sanctions, discipline, or losing a case – leave no choice but to verify every citation produced by an AI tool.

BY BRENT J. HOEFT

The ever-increasing uptake in usage of generative artificial intelligence (GAI) in the legal profession is understandable; it can be genuinely useful for lawyers in many areas of a law practice. GAI can outline and help flesh out legal arguments, summarize and aid in review of long, complex documents, help generate a checklist of issues to investigate, draft multi-document deal packages, conduct 50-state surveys on a legal topic, and much more. But its usage comes with several risks, including one that can be particularly detrimental to lawyers – hallucinations.

Hallucinations are GAI outputs that might read as confident, professional, and convincing but are wrong, unsupported, or misleading.¹ Hallucinations are not glitches in the system; they arise out of how GAI systems work. One type of GAI, and the kind most used in the legal profession, are large language models (LLMs), which are prediction-based language models.² At a very basic level, this means that these systems are predictive tools trained to predict the next most likely word. LLMs excel at producing persuasive output based on only predictive likelihoods, but they are not inherently built to ensure legal accuracy. When the LLM references information that is incomplete, unreliable, or conflicting, the LLM might fill in the gaps with text that looks like legal analysis.³


Therein lies the problem. If a lawyer isn't careful, fails to do the requisite due diligence, or is not knowledgeable enough in a practice area to catch the hallucination, using a GAI tool without the proper oversight can lead to serious

consequences. An attorney might violate court rules, be subjected to sanctions and disciplinary matters, and cost their client the case.

One common misunderstanding is that there is only one type of hallucination: the GAI tool makes up stuff that does not exist, be it a court, a case, or quotations. Therefore, a lawyer might think that so long as the existence of these cases and courts is verified, then that is all that



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.



needs to be done to ensure there are not any hallucinations. However, this is only one type of hallucination; it is important to be aware that there are others. Below are three common types of hallucinations lawyers should recognize.

Hallucination Type 1: Fabricated Authority

This type of hallucination generated the earliest headlines, is most known by attorneys, and is the easiest to describe: the AI tool simply makes stuff up. In these hallucination instances, the GAI output provides a case name, a citation, or even a convincing quotation that does not exist at all.

In *Mata v. Avianca Inc.*,⁴ an early, widely cited sanctions decision, a federal judge found that lawyers submitted filings to the court citing nonexistent judicial opinions with fake quotations and citations generated by ChatGPT. The court emphasized that while technology can assist legal work, attorneys retain a gatekeeping duty to ensure accuracy and compliance with Fed. R. Civ. P. 11.⁵ The court noted that fabricated authority is harmful on many levels because it wastes client resources, forces opposing counsel and courts to spend time disproving fiction, and can result in sanctions for the attorney and harm to the client and to attorney reputation.⁶ *Mata* is not the only case of its kind.

An attorney must independently verify every case citation or quotation in a motion or brief to the court. Often this type of hallucination arises from using a general foundational AI model (for example, ChatGPT, Claude, or Gemini) to do legal research. These foundational models are not legal research tools. They are not grounded in a database of good legal precedent. They may return output that looks like strong legal reasoning based on legal precedent. But “looks right and sounds right” is not verification. Independent verification using a trusted legal research tool is necessary.

ALSO OF INTEREST



AI Summit 2026

Interested in learning more about artificial intelligence? If yes, then this full-day, live seminar in Madison is for you!

AI Summit 2026 features presentations from a faculty of experts, including the State Bar of Wisconsin's Practice411™ manager Brent Hoeft and ethics counsel Sarah E. Peterson, as well as attorneys Aaron W. Brooks, Timothy D. Edwards, Fatima D. Pahlavan, and Kris Turner.

The summit is presented by State Bar of Wisconsin PINNACLE; qualifies for 4.5 CLE, 1.5 EPR, and 3 LPM credits; and will be held at the State Bar Center in Madison on Wednesday, May 20, from 8:30 to 4:45 p.m.

For more information and to register, please visit marketplace.wisbar.org. Product ID: CA3959M. WL



Hallucination Type 2: Real Case or Quotation, Wrong Proposition

In this situation, AI output contains a citation to a real case but for a proposition that the case does not support, or it overstates the holding, treats dicta as binding, or ignores key limiting facts. This is often more dangerous than phantom citations because the case the AI purports to be citing does exist, so a quick “does this case exist?” check will return a positive response and might lull the user into false confidence.

This kind of hallucination was at issue in *Whiting v. City of Athens*.⁷ The court found that plaintiff's counsel had misrepresented the district court's sanctions order and had also cited more than 24 fake citations, provided citations that lacked the language quoted in the brief, and provided citations that failed to support the cited proposition.⁸

This case is a good reminder to treat AI-generated content as a starting point. Attorneys are expected to undertake appropriate due-diligence verification and to make sure that the case is supported by legal precedent.

Hallucination Type 3: Grey-Area Overconfidence

In this type of situation, the GAI tool is not inventing a case or quotation, and it might not even be entirely “wrong” in what it is stating. Instead, the GAI tool presents a legal grey area as definitively settled, glossing over splits of authority, exceptions, fact sensitivity, or jurisdictional differences. This type of hallucination is even harder to identify without proper due diligence and verification because it reads like clean, confident analysis. LLMs are optimized to give a direct answer. When the underlying law is messy, they often “smooth the mess” into a simple rule.⁹

The *Whiting* case dealt with this type of hallucination as well. In *Whiting*, the court stated that counsel provided several citations that did not support the proposition being asserted and were used to give the impression that the law was clear on the issue when in fact it was not. The court considered this to be a substantive misrepresentation of the law because there were no acknowledgments made of circuit splits, limiting facts, contrary authority, or open questions in the case law.¹⁰

So, an AI-generated response might be incomplete because controlling limitations, key exceptions, or the most relevant authority is missing. This risk might be amplified in areas with fast-evolving or newly developing case law or areas that have small state-by-state variations.

When AI seems unusually certain in a complicated area, treat that as a red flag. Ask the AI tool to identify splits, minority views, and Wisconsin-specific nuance. After that, the lawyer must still verify using primary law.

Practical Tips for Verification of GAI Output

A simple, repeatable workflow can keep AI useful without letting it become a silent liability.

1) Use AI at the starting line and during the race, never at the finish. Let it suggest issues, organize facts, draft outlines, or generate a research plan. An attorney must independently verify all citations generated by AI for their accuracy. AI-generated content should never leave the law firm without first being verified by an attorney.

2) Prioritize RAG-based legal AI for research questions. If answering a question requires legal research, the attorney or staff member must use tools grounded in Westlaw, Lexis, or vLex/Fastcase content libraries rather than a general foundational model.¹¹ These legal research tools use retrieval augmented

generation (RAG). Instead of relying solely on what a model absorbed during training, the system retrieves relevant documents from a defined and limited collection and uses that retrieved text as “grounding” for the response.¹² Microsoft describes RAG as a three-step process: Retrieve → Augment → Generate.¹³ The process is designed to keep answers grounded in the sources specifically provided or authorized. When RAG is paired with authoritative and trusted legal databases from Westlaw, Lexis, or vLex/Fastcase, the AI is limited to vetted legal content rather than the uncurated entirety of the internet or the model’s own internal memory.

Caution: *Reduces does not mean eliminates.* Even the best RAG systems can still fail by retrieving the wrong materials, missing the most controlling authority, or producing an answer that is accurate yet incomplete (see hallucination types 2 and 3 above). This is why, even when using RAG-based legal AI tools, an attorney’s professional obligation remains the same: verify the authority, read the key passages, and confirm validity before relying on it.

3) Verify: pull and read the sources and check validity and context.

Lawyers must not rely on a citation or quotation unless the actual case, statute, or regulation is reviewed and verified. This ethical obligation exists no matter the original source of the

information – be it a law clerk or AI generated. The tool used is irrelevant. It is always the attorney who has due diligence and ethical obligations to verify accuracy and prevent misrepresentations to the court. True independent verification requires the attorney of record to do the following: 1) open and read every case cited to make sure that the proposition is accurate for the case in which it is cited, 2) verify every quotation used is true and accurate in the case cited, and 3) verify that all cited cases are still good law.¹⁴

4) Be aware of grey areas in the law.

If the AI-generated answer seems “too clean,” force the tool to address splits, minority views, and uncertainties and then go back to step 3 above: verify.

Conclusion

GAI tools have the potential to be game changers for how attorneys practice law. Hallucinations pose major issues, so lawyers must remain vigilant to the risks that are inherent in these AI tools. As always, with any technology, it is the lawyer’s ethical obligation to be competent in the technology used and to understand the benefits and risks of using that technology in representing clients.¹⁵ Ultimately, the responsibility for legal accuracy and integrity rests with the attorney, not with the technology used. **WL**

ENDNOTES

¹LexisNexis, *Retrieval Augmented Generation (RAG) for Trusted Generative AI* (Sept. 18, 2025), <https://www.lexisnexis.com/community/insights/professional/b/industry-insights/posts/rag-for-generative-ai> [hereinafter LexisNexis, RAG].

²*Id.*

³*Id.*

⁴*Mata v. Avianca Inc.*, 678 F. Supp. 3d 443 (S.D.N.Y. 2023), <http://case-law.vlex.com/vid/mata-v-avianca-inc-1056619281>.

⁵This rule and other Federal Rules of Civil Procedure can be viewed at LII, *Federal Rules of Civil Procedure*, <https://www.law.cornell.edu/rules/frcp> (last visited April 8, 2026).

⁶*Mata*, 678 F. Supp. 3d 443.

⁷*Whiting v. City of Athens*, Nos. 24-5918, 24-5919, 25-5424 (6th Cir. Mar 13, 2026), <http://case-law.vlex.com/vid/whiting-v-city-of-1114443436>.

⁸*Id.*

⁹LexisNexis, RAG, *supra* note 1.

¹⁰*Whiting*, Nos. 24-5918, 24-5919, 25-5424.

¹¹See Thomson Reuters, *Introducing AI-Assisted Research: Legal Research Meets Generative AI*, <https://legal.thomsonreuters.com/>

[blog/legal-research-meets-generative-ai/](https://www.lexisnexis.com/blog/legal-research-meets-generative-ai/); LexisNexis, *LexisNexis Enhances Lexis+ AI with new Features, AI Models, and Graph Technology to Further Drive High Quality, Trusted Answers for Legal Professionals* (July 22, 2024), <https://www.lexisnexis.com/community/pressroom/b/news/posts/lexisnexis-enhances-lexis-ai-with-new-features-ai-models-and-graph-technology-to-further-drive-high-quality-trusted-answers-for-legal-professionals>; vLex, *AI Engineered for Lawyers*, <https://vlex.com/vincent-ai>.

¹²LexisNexis, RAG, *supra* note 1.

¹³Microsoft Build 2026 (June 2-3, 2026), *Retrieval Augmented Generation (RAG) and Indexes*, <https://learn.microsoft.com/en-us/azure/foundry/concepts/retrieval-augmented-generation?view=foundation-classic>.

¹⁴Nexlaw Blog, *Using AI to Check AI Citations Is Not Verification* (April 1, 2026), <https://www.nexlaw.ai/blog/using-ai-to-check-ai-citations-not-verification/>.

¹⁵Wis. Sup. Ct. R. 20:1.1, <https://www.wicourts.gov/sc/rules/chap20a.pdf>. **WL**