



A California Law Firm Just Paid \$31K in Sanctions for AI Hallucinations

Lessons for California Attorneys



There was a time when the term hallucination belonged strictly to the realm of human perception. Now, it's part of the artificial intelligence (AI) lexicon, used when generative systems give incorrect information due to their probabilistic nature.

There was a time when the term hallucination belonged strictly to the realm of human perception. Now, it's part of the artificial intelligence (AI) lexicon, used when generative systems fabricate facts, such as inventing a Supreme Court case out of thin air.

As rapidly as generative AI tools such as ChatGPT and Google Gemini have streamlined workflows, cautionary tales have emerged in their wake. For California attorneys embracing AI for research and drafting, recent high-profile missteps serve as stark reminders: The risks are real, the consequences are escalating and the ethical obligations are nonnegotiable.



Here's what to know and how to mitigate the risks, including best practices for responsible use, verification and disclosure.

1 A 'collective debacle' in the Central District of California

When a team of experienced California litigators set out to craft what should have been a routine supplemental brief on a privilege issue in federal insurance lawsuit *Lacey v. State Farm General Insurance Co.*, they unwittingly walked into a labyrinth of AI-generated illusions.

It began innocently enough. A partner at a business litigation firm leveraged CoCounsel, Westlaw Precision and Google Gemini to help outline his arguments for the plaintiff. But the draft, passed from one lawyer to another — including co-counsel at another firm — was littered with hallucinations. At least nine out of 27 citations were wrong; two citations didn't exist at all and several quotes attributed to judicial rulings were false or inaccurate.

The unraveling was swift and public. When the special master in the case flagged two suspicious citations, the attorneys scrambled to submit a "corrected" brief with at least six more phantom authorities. The special master didn't mince words, calling the episode a "collective debacle" and "tantamount to bad faith." He criticized both the drafting firm for not verifying the AI's work and failing to disclose his use of it and the other firm for failing to check the research they received.

"Plaintiff's use of AI affirmatively misled me," the special master wrote. "I read their brief, was persuaded (or at least intrigued) by the authorities that they cited, and looked up the decisions to learn more about them — only to find that they didn't exist. That's scary. It almost led to the scarier outcome (from my perspective) of including those bogus materials in a judicial order. Strong deterrence is needed to make sure that attorneys don't succumb to this easy shortcut."

In addition to falling under a harsh media spotlight, the consequences were severe for the plaintiff's team. All versions of the brief were struck from the record, the attorneys' requested discovery relief was denied, the firms were ordered to pay \$31,100 in legal fees for the defendant and the attorneys were required to disclose the incident to their client.

The defense sought an additional \$25,000 in costs for preparing its own brief and attending the hearing, but the special master opted for an additional \$5,000, reasoning that full reimbursement wasn't necessary for deterrence. He also declined to impose sanctions or penalties against the plaintiff or individual attorneys, noting that "Their admissions of responsibility have been full, fair and sincere, and that "Justice would not be served by piling on them for their mistakes."



"I read their brief, was persuaded (or at least intrigued) by the authorities that they cited, and looked up the decisions to learn more about them — only to find that they didn't exist. That's scary. It almost led to the scarier outcome (from my perspective) of including those bogus materials in a judicial order. Strong deterrence is needed to make sure that attorneys don't succumb to this easy shortcut."

2 A judicial rebuke in Toronto

Another recent AI-assisted misstep resulted in a show-cause hearing, where the attorney in question was ordered to demonstrate why she shouldn't be held in contempt for supporting her legal arguments with cases that didn't exist.

In this case, a Toronto attorney found herself in ethical hot water after filing a *factum* (Canada's version of a legal brief) in a complex estate and family law matter. When an Ontario Superior Court judge tried to locate some of the referenced authorities, he was unable to do so, igniting suspicion that the attorney had used AI to draft the document and failed to check for accuracy. Other elements of the filing misinterpreted existing case law, according to the judge.

In a subsequent hearing, the attorney admitted using ChatGPT to produce the *factum*, which included fake cases. The judge dismissed the contempt proceedings, reasoning that the attorney had expressed deep regret, been "bombarded" by calls from reporters and colleagues, taken full responsibility and outlined steps she had taken to prevent future errors. These included withdrawing and correcting the *factum*, not billing the client for the work and committing to continuing professional development focused on legal ethics and AI use.

"The error was not delegating the *factum* or using generative AI to assist in drafting the *factum*," the judge wrote. "Rather, [her] failure arose when she signed, delivered and used the *factum* without ensuring that the cases were authentic and supported the legal arguments she was submitting to the court."



*"Rather, [her] failure arose when she signed, delivered and used the *factum* without ensuring that the cases were authentic and supported the legal arguments she was submitting to the court."*

3 A faux pas in Utah

Utah's first known case of an AI-generated legal hallucination led to court-ordered sanctions and the dismissal of a newly hired law clerk.

In what's now becoming a familiar story, the law firm submitted a court filing containing several miscited cases, including at least one reference that couldn't be found in any legal database. An investigation revealed that ChatGPT had created the nonexistent citation, "*Royer v. Nelson, 2007 UT App 74, 156 P.3d 789*" — but couldn't provide any further details other than that it involved a dispute between individuals named Royer and Nelson in the Utah Court of Appeals.

A recent law school graduate had prepared the filing while working as an unlicensed clerk, and he'd used ChatGPT without informing the firm's attorneys. The firm, which didn't have an AI policy at the time, submitted the document without catching the errors — a move that the judge overseeing the case said violated their duty to ensure the accuracy of everything filed with the court.

The judge described this as a failure of the lawyers' "gatekeeping responsibilities as members of the Utah State Bar," ordering the lead attorney to pay the opposing party's legal fees and to donate \$1,000 to a local legal aid organization. Acknowledging that the firm had quickly accepted responsibility and updated its internal policies, the judge declined to impose harsher penalties. The clerk was fired, and the firm has since established an AI policy.

4 Claude hallucinates in litigation over its training

Even AI companies aren't immune from the pitfalls of their own technology. As part of its defense in a California copyright lawsuit over the training of its AI chatbot Claude, Anthropic's legal team submitted expert testimony misciting an academic article.

The data scientist in question pointed the court to an article in *The American Statistician* to demonstrate how rarely Claude reproduces copyrighted song lyrics, but it transpired that the citation was a Claude-generated hallucination. The firm clarified that the expert had cited a real journal article, but the AI chatbot had introduced mistakes.

"Unfortunately, although providing the correct publication title, publication year and link to the provided source, the returned citation included an inaccurate title and incorrect authors," the attorney wrote in a subsequent declaration to the court.

The judge called the situation "very serious and grave," noting a critical distinction between a missing citation and one invented by AI. The firm stated it has since implemented additional levels of review to avoid a similar incident.



"Unfortunately, although providing the correct publication title, publication year and link to the provided source, the returned citation included an inaccurate title and incorrect authors."



Lessons Learned: Ethical Obligations in the Age of AI

These are not isolated incidents. Courts across North America are responding with sanctions, public reprimands and, in some cases, threats of contempt to underscore the seriousness of misusing AI-generated content in legal filings. The harms are multifaceted, including wasted resources for opposing parties and the court, erosion of public trust, lasting reputational damage for attorneys and their firms and the potential for erroneous judicial orders based on fabricated law.

Here are the key lessons for lawyers navigating the use of AI in practice.

1 Verify everything

Generative AI tools are powerful and efficient, but they can't substitute an attorney's legal judgment. Because large language models (LLMs) predict text based on patterns rather than retrieving verified facts, they will confidently produce responses that sound plausible but are entirely incorrect.

AI systems optimize for plausible-sounding responses rather than factual accuracy. So, attorneys must independently verify all AI-generated content, especially legal research and citations. Relying on AI without scrutiny is a breach of the duty of competence and diligence.

2 Be transparent about AI use

When AI is used in legal work, attorneys should disclose it — both internally and externally. Be open about the use of AI tools with colleagues and, where appropriate, clients and the court, especially if there's any risk that AI-generated content could be incorporated into filings. Remember, clients have a right to understand how their information is being handled, especially if it's being processed by technologies that may raise privacy or ethical concerns.

3 Understand the technology as it evolves

The [State Bar of California](#) and the [American Bar Association](#) have both issued guidance stressing that attorneys must understand the benefits and risks of any technology they use. This includes knowing how AI tools generate content, their limitations and the potential for hallucinations or data privacy issues.

4 Protect confidentiality

AI tools, especially public or cloud-based models, may retain or learn from user inputs. Never input confidential client information into systems that don't guarantee data privacy and security. Opt for enterprise-grade tools with robust protections or keep sensitive work out of an LLM.

Before integrating generative AI tools into any workflow, review:

- The platform's terms of service
- Whether it uses client data to train its models
- Its data privacy protections and jurisdiction of data storage
- Whether it offers options to opt out of model training or third-party data sharing

Ask your vendors:

- Will the platform store or process sensitive client financial or health data?
- How is the LLM trained and does it involve user data?
- Is data encrypted at rest and in transit? This is particularly important for client estate plans and tax details.
- Does the tool support HIPAA or other applicable compliance frameworks (for elder care/medical directives)?
- Can generated documents be saved and version-controlled securely for compliance audits?

Lessons Learned: Ethical Obligations in the Age of AI



5 Supervise and train attorneys

Firms must establish clear policies and provide regular training on AI use. Supervising attorneys are responsible for ensuring that all staff understand the risks and that AI-generated work is properly vetted before submission.

All attorneys and staff should understand:

- How to securely use AI tools during document drafting or client intake
- What types of information should never be entered into public-facing AI tools
- How to identify and correct AI “hallucinations,” such as incorrect case law or fictional legal strategies
- How AI outputs should be integrated into legal analysis — rather than taken at face value

6 Monitor evolving law and guidance

AI regulation is changing rapidly. For example, as of Jan. 1, 2026, California’s [AI Transparency Act](#) will require greater transparency and accountability for generative AI systems that produce images, audio and video. The law targets large AI developers to ensure their tools clearly label AI-generated content and allow users to detect when something has been created or altered by AI. This means briefs, memos or other legal documents generated (even in part) by AI tools covered under the law must contain disclosures.

At the federal level, the Biden administration issued an executive order in 2023 establishing new standards for AI safety, security and civil rights, as well as promoting responsible innovation. The Trump administration has since repealed that order and shifted focus to promoting U.S. dominance in AI, encouraging innovation and removing some restrictions (including the AI Bill of Rights provisions that limited certain uses of AI by tech companies such as Google).

Stay informed through [California State Bar updates and resources](#), [ABA opinions](#) and [legal news coverage](#) to ensure compliance.

Use AI as a tool, not a crutch

AI has the potential to significantly improve the speed, accuracy and efficiency of legal work. In fact, law firms that fail to integrate it effectively will risk falling behind. But the pitfalls are real. As recent cases show, the consequences of careless reliance on AI can include sanctions, reputational harm and broken client trust.

As these tools become more sophisticated, California attorneys must adopt them with intention, prioritizing transparency, ethical obligations and sound legal judgment. The next headline about an AI blunder doesn’t need to involve your firm — if your lawyers treat AI as an aid, not a substitute, by taking ownership of their analysis and advice.



► For more insights into AI and legal ethics, get in touch to [schedule a free demo](#).

CEB is a registered trademark of Continuing Education of the Bar - California (CEB). © The Regents of the University of California, 2025. All rights reserved.



1-800-232-3444



www.ceb.com

A California Law Firm Just Paid \$31K in Sanctions for AI Hallucinations