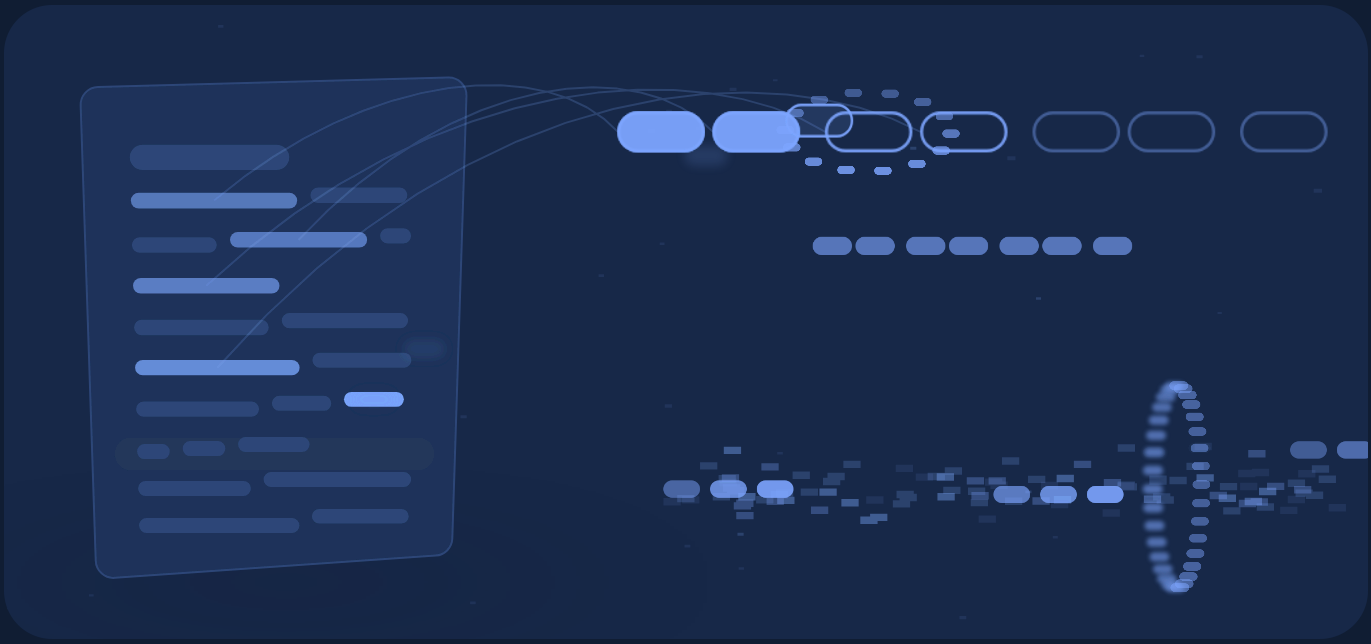
Rillor

RESEARCH REPORT

What court records show about fabricated citations in LLM output

Court opinions and sanctions orders from 2023 to 2026 show what judges found in hallucinated legal citations and what they ordered in response.



Rillor

Contents

01 Summary

02 Background

03 Method

04 Findings

05 Implications

06 Limits of this analysis

07 Sources

Summary

Court records from June 2023 to March 2026 describe three ways citations fail in filings drafted with generative tools, and remedies that run from penalties to disqualification and dismissal [4-20].

- A public database listed 2,286 court and tribunal decisions as of 9 November 2026 in which a court or tribunal found or implied that a party relied on hallucinated content: 1,550 from the United States, with self-represented litigants a party in 1,314 and lawyers in 907 [1]. Its author calls the count an undercount [3].
- Courts describe three failures: authorities that do not exist, quotations a real source does not contain, and real authorities cited for what they do not hold. In Al-Haroun v Qatar National Bank (High Court of England and Wales, June 2025), 18 of 45 citations were to cases that do not exist, and in Noland v. Land of the Free (California Court of Appeal, September 2025), 21 of 23 quotations in one brief were fabricated [5, 10].
- Remedies ranged from penalties to disqualification and dismissal. Mata v. Avianca imposed $5,000 in June 2023, Johnson v. Dunn disqualified three lawyers from the case in July 2025, and Couvrette v. Wisnovsky dismissed the plaintiffs' claims with prejudice in December 2025 [4, 13, 15].
- Two U.S. courts wrote in 2025 that the penalties imposed in earlier cases had failed to deter [15, 17]. No source reviewed measures whether sanctions deter.
- Court rules differ. The Fifth Circuit declined a special rule and relies on existing duties, the Northern District of Texas requires a disclosure on a brief's first page, and a New South Wales practice note requires authors who used generative tools to verify every citation, never solely with such a tool [21, 23, 25].
- A citation lookup can show that an authority exists, and courts named reading the cited source as the check for quotations and holdings [6, 10, 31].

Background

Mata v. Avianca, Inc. is the order later courts quote. The plaintiff's lawyers cited decisions that did not exist in a submission of 1 March 2023, and the Clerk of the Eleventh Circuit confirmed that one of them "is not an authentic ruling of the Court" [4]. On 22 June 2023 the Southern District of New York found that the respondents "continued to stand by the fake opinions after judicial orders called their existence into question", and found bad faith "based upon acts of conscious avoidance and false and misleading statements to the Court" [4]. The opinion separated the tool from the duty: "there is nothing inherently improper about using a reliable artificial intelligence tool for assistance. But existing rules impose a gatekeeping role on attorneys to ensure the accuracy of their filings" [4]. Its list of harms, from the other side's wasted time and money to a future litigant's claimed doubt about a real ruling, was later quoted by the UK First-tier Tribunal and the Court of Appeals of Georgia [4, 7, 12].

Damien Charlotin started a public database of such decisions in April 2025 [3]. In November 2025 he estimated that before 2025 "we maybe had two or three cases a month", and in a February 2026 edit he put the daily average at five [3]. Both are his estimates. By our count of its CSV export, the database lists 16 decisions from 2023, 61 from 2024, 865 from 2025 and 1,344 from 2026 to 9 November [2].

The Federal Court of Canada defines hallucination as "facts, citations, and other content generated by AI that are not true, and have been fabricated by AI in response to a prompt or request" [24]. This report reads the court records: what judges found, what they ordered and the rules courts and bars wrote. The check itself, step by step, is the subject of our methods note on checking LLM citations against their sources (https://rillor.com/insights/checking-llm-citations-against-sources).

Method

What was examined. Court decisions, sanctions orders, court rules, practice notes and bar guidance on fabricated or misrepresented citations in court filings, and the public database that collects such decisions.

Records read. Seventeen decisions from the United States, the United Kingdom, Canada and Australia, dated 22 June 2023 to 23 March 2026 [4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20]. Each was read in full from a public copy: the court's own, CourtListener's, the National Archives' or the copy the database hosts, which also served where court sites refused automated access. The rules and guidance are Federal Rule of Civil Procedure 11 [22], instruments of four courts [21, 23, 24, 25] and the State Bar of California's guidance [26], with a Senate committee release [27], three studies of legal hallucination [28, 29, 30] and a citation lookup service's documentation [31].

The database. Counts were read on 10 November 2026 from the page last updated 9 November 2026 [1]. It covers "only cases where the court or tribunal has explicitly found (or implied) that a party relied on hallucinated content or material", plus some decisions where use was "alleged but not confirmed" [1]. It "does not track the (necessarily wider) universe of all fake citations or use of AI in court filings" [1], and its author calls it "necessarily an undercount" [3].

Counting. Page counts are filters. A decision can name several kinds of party and carry several failure categories, so the counts overlap and are never shares of one total. Figures marked "by our count" come from the CSV export downloaded 10 November 2026: 2,286 rows, all of them dated [2]. Its party field matches the page filter [1, 2]. In the monetary field, "1 is a placeholder, not the real figure" where the amount is unknown [1]. The amounts include fee awards (Finding 3), so this report calls them monetary awards. Other currencies were not converted.

Naming. Parties appear as the case captions name them. Tools and products go unnamed, and quotations are shortened where they name a person or a product.

Findings

1. Decisions by party, country and year

Self-represented litigants appear in 1,314 decisions, lawyers in 907, judges in 34 and experts in 18 [1]. Among U.S. decisions, by our count, those naming a lawyer rose from 208 in 2025 to 382 in 2026 to 9 November, and those naming a self-represented litigant from 311 to 581 [2]. The author describes the cases as "either self-represented litigants, or lawyers who are surprised that an existing tool suddenly has an AI component that hallucinates" [3].

Year of decision	Decisions	U.S. decisions
2023	16	11
2024	61	38
2025	865	528
2026, to 9 November	1,344	973

Source: our count of the CSV export [2].

The U.S. share ran between 61% and 72% each year (derived) [2]. Monthly counts for 2026 run from 184 in March to 69 in October [2]. The author finds cases through referrals, "dedicated scrapers and bots" and keyword searches [3], so the latest months may be incomplete, and no source reviewed says the pace fell.

After the United States, the page lists Canada with 245 decisions, Australia 112, the United Kingdom 90, Israel 58 and Brazil 41, among 46 entries in its jurisdiction filter, which include domestic and international arbitration and WIPO [1]. The author links the U.S. weight to "an excellent legal data ecosystem", while in many countries legal data "is hoarded by legal editors, or subject to many artificial frictions" [3]. By legal field, contract cases appear in 462 decisions, administrative 290, civil rights 220, employment 208 and tort 201 [1].

2. Fabricated authorities, misrepresented holdings and false quotes

The High Court of England and Wales described what generative tools can produce: they "may cite sources that do not exist. They may purport to quote passages from a genuine source that do not appear in that source" [5]. The database tags each decision with the failures found.

Category (page filter)	Decisions
Fabricated	1,872
Misrepresented	968
False quotes	606
Outdated advice	38

Source: [1], as of 9 November 2026.

The categories sum to 3,484 against 2,286 decisions (derived) [1]. By our count, 1,037 decisions carry more than one category and 910 carry only Fabricated [2]. By what was invented or misused, the page lists case law in 2,047 decisions, legal norms in 262, exhibits and submissions in 166, doctrinal work in 57, and overturned case law and repealed law in 19 each [1].

Single filings show the scale. Motions in limine in Wadsworth v. Walmart "cited nine cases, but eight did not exist" [11]. In Al-Haroun, 18 of 45 citations were to cases that do not exist, and the real ones "in many instances ... did not contain the quotations that were attributed to them, did not support the propositions for which they were cited" [5]. An appellee's brief in Shahid v. Esaam carried "11 bogus case citations out of 15 total" [12]. In Noland v. Land of the Free, 21 of 23 case quotations in the opening brief were fabricated, and "many of the cases plaintiff cites do not discuss the topics for which they are cited" [10]. Three briefs in Couvrette v. Wisnovsky cited fifteen nonexistent cases and attributed fabricated quotations to eight real authorities, and the court left the fake names out "as doing so could inadvertently amplify the lie that these cases exist" [13]. The Fifth Circuit's opinion in Fletcher v. Experian lists, in its text, the 16 fabricated quotations and 5 "additional serious misrepresentations of law or fact" named in its show-cause order [20]. Invented sources need not be cases: in Kohls v. Ellison, an expert's declaration cited two nonexistent academic articles [18].

The errors came from different places. In Wadsworth, the cases came from a platform that, the court said, "appears to be an in-house database" of the plaintiffs' firm [11], and in Dayal from a tool "incorporated in the legal practice management software he subscribes to" [9]. The database warns that a tool's mention "does not necessarily mean that tool was responsible" [1].

3. What courts ordered

Date	Decision	Order
22 June 2023	Mata v. Avianca, S.D.N.Y.	$5,000 penalty, letters to the plaintiff and to each judge falsely named as an author [4]
30 January 2024	Park v. Kim, 2d Cir.	Referral to the Grievance Panel, a copy of the decision to the client [6]
20 February 2024	Zhang v. Chen, B.C. Supreme Court	Costs of the extra effort borne by counsel personally [8]
27 August 2024	Dayal, Federal Circuit and Family Court of Australia	Referral to the Victorian legal regulator [9]
10 January 2025	Kohls v. Ellison, D. Minn.	Expert declaration excluded [18]
24 February 2025	Wadsworth v. Walmart, D. Wyo.	Pro hac vice admission revoked, $3,000 for the drafter, $1,000 for each of two signers [11]
6 May 2025	Lacey v. State Farm, C.D. Cal.	Briefs struck, discovery relief denied, $31,100 to the defense [16]
20 May 2025	Ko v. Li, Ontario Superior Court	Contempt proceeding dismissed on conditions [19]
28 May 2025	HoosierVac, S.D. Ind.	$6,000, against $15,000 recommended [17]
6 June 2025	Ayinde and Al-Haroun, High Court	No contempt proceedings, judgment sent to the profession's bodies [5]
30 June 2025	Shahid v. Esaam, Ga. Ct. App.	Trial court order vacated, $2,500, "the most the law allows" [12]
23 July 2025	Johnson v. Dunn, N.D. Ala.	Public reprimand, disqualification, referral to the state bar [15]
12 September 2025	Noland v. Land of the Free, Cal. Ct. App.	$10,000, opinion served on the client and the State Bar [10]
12 December 2025	Couvrette v. Wisnovsky, D. Or.	Briefs struck, $15,500, claims dismissed with prejudice [13]

Date	Decision	Order
18 February 2026	Fletcher v. Experian, 5th Cir.	$2,500 [20]
23 March 2026	Couvrette v. Wisnovsky, D. Or.	$94,704.38 in fees and costs, split 15% and 85% between local and lead counsel [14]

Rule 11 frames the federal orders. A sanction "must be limited to what suffices to deter repetition of the conduct or comparable conduct by others similarly situated", and may be a nonmonetary directive, a penalty paid into court or, on motion, the other side's fees [22]. "Absent exceptional circumstances, a law firm must be held jointly responsible for a violation committed by its partner, associate, or employee", as Mata applied [4, 22]. The rule bars a monetary sanction "against a represented party for violating Rule 11(b)(2)" [22]. In Couvrette, the money fell on counsel, and the plaintiffs, who "share responsibility with their attorney", lost their claims [13, 14].

Fault and conduct shaped the remedy. Mata found bad faith [4], while Wadsworth held that "A finding of subjective bad faith is not required to impose sanctions" under its standard [11]. Zhang v. Chen refused special costs because counsel had no intention "to deceive or misdirect" [8]. Lacey sanctioned no individual lawyer because "their admissions of responsibility have been full, fair, and sincere" [16]. Noland declined to order sanctions payable to opposing counsel because the respondents "did not alert the court to the fabricated citations" [10].

Deterrence in the courts' words. Johnson v. Dunn held that "Fabricating legal authority is serious misconduct that demands a serious sanction", called for "substantially greater accountability than the reprimands and modest fines that have become common", and wrote that "If fines and public embarrassment were effective deterrents, there would not be so many cases to cite" [15]. In HoosierVac, the magistrate judge's recommendation found that penalties imposed on other attorneys "have evidently failed to act as a deterrent", and the district judge set $6,000 as the amount that "suffices to deter" [17]. Couvrette adopted an Oregon Court of Appeals valuation of "$500 per non-existent case and $1,000 per fabricated quotation" and warned that "With a known price tag, an attorney could decide" the risk was worth taking [13]. The High Court held that "save in exceptional circumstances, admonishment alone is unlikely to be a sufficient response" [5]. No source reviewed measures whether any sanction deterred.

The database's monetary figures. By our count, 394 decisions record an amount, the placeholder included: 267 of the 907 with a lawyer as a party (29.4%) and 123 of the 1,314 with a self-represented party (9.4%) [2]. Among 206 U.S. decisions with an amount above the placeholder, the median is $2,500 and the mean $7,568 [2]. The amounts reach $110,204, Couvrette's $15,500 sanction plus its $94,704.38 fee award (derived) [2, 13, 14]. The database flags a professional sanction in 169 decisions, 163 of them with a lawyer as a party [2].

4. Self-represented litigants, experts and courts' own orders

Self-represented litigants. In Harber v The Commissioners for His Majesty's Revenue and Customs, decided on 4 December 2023, the self-represented appellant gave "the names, dates and summaries of nine First-tier Tribunal ('FTT') decisions", and "none of those authorities were genuine" [7]. The tribunal accepted that she "did not know how to check their validity by using the FTT website or other legal websites" [7]. Its decision "would have been the same" without them, yet relying on authorities that are not genuine "is a serious and important issue" [7]. The database records no sanction on her [2]. Rule 11(b) binds "an attorney or unrepresented party" [22]. Dahl and colleagues wrote that "the risks are highest for those who stand to benefit from LLMs the most", naming pro se litigants and those without access to traditional legal resources [28].

Experts. In Kohls v. Ellison (D. Minn., 10 January 2025), an expert's declaration "inadvertently included citations to two non-existent academic articles, and incorrectly cited the authors of a third article" [18]. He "typically validates citations with a reference software when he writes academic articles but did not do so" here, and the court excluded the declaration [18].

Courts' own orders. In Shahid v. Esaam (30 June 2025), the trial court's order "relied upon non-existent case law", and the Court of Appeals of Georgia vacated it, making "no factual finding as to who (or what) inserted the fictitious cases" [12]. A Senate Judiciary Committee release of 23 October 2025 reports that two federal district judges, in Mississippi and New Jersey, "admitted ... that their staff used generative AI to draft error-ridden orders", which "misquoted state law, referenced individuals who didn't appear in the case and attributed fake quotes to defendants" [27]. One judge now requires a second independent review and printed copies of all cited cases, and the other bars law clerks and interns from using such tools to draft opinions or orders [27]. In Lacey, the Special Master looked up the authorities that had persuaded him and found they did not exist: "That's scary" [16]. The database lists 34 decisions with a judge as a party [1].

5. Courts outside the United States

England and Wales. In R (Ayinde) v London Borough of Haringey and Al-Haroun v Qatar National Bank, decided on 6 June 2025, a Divisional Court used its Hamid jurisdiction, "the court's inherent power to regulate its own procedures and to enforce duties that lawyers owe to the court" [5]. It held that freely available generative tools "are not capable of conducting reliable legal research", and that lawyers who use them "have a professional duty therefore to check the accuracy of such research by reference to authoritative sources" [5]. A lawyer "is not entitled to rely on their lay client for the accuracy of citations of authority or quotations" [5]. Its powers run from "public admonition of the lawyer" through wasted costs orders and "referral to a regulator" to "referral to the police" [5]. In Ayinde, the defendant sought wasted costs because the claimant's lawyers "had cited five fake cases" [5], and the database records a wasted costs order and £11,000 at the hearing of 3 April 2025 [2]. The court found the threshold for contempt met and decided "not to initiate contempt proceedings or to refer the case to the Law Officers" [5]. It called for "practical and effective measures" by leaders such as "heads of chambers and managing partners" and by regulators, and directed that the judgment go to the Bar Council, the Law Society and the Council of the Inns of Court [5].

Canada. In Zhang v. Chen (Supreme Court of British Columbia, 20 February 2024), a notice of application cited two cases, both "discovered to be non-existent" [8]. The court held that citing fake cases "is an abuse of process and is tantamount to making a false statement to the court", and ordered that the "additional effort and expense is to be borne personally" by counsel, set at four half days [8]. In Ko v. Li (Ontario Superior Court of Justice), counsel relied on two nonexistent cases in oral argument on 1 May 2025, and the problem came to light when the judge "could not find any reference online to the cases" [19]. On 20 May 2025 the court dismissed the contempt proceeding on conditions: professional development courses, and no bill to the client for the research, factum writing and attendance [19].

Australia. In Dayal [2024] FedCFamC2F 1166 (27 August 2024), a solicitor tendered a list and summary of authorities that did not exist, and the "accuracy of the document produced was not verified by the solicitor" [9]. The court referred the matter to the Office of the Victorian Legal Services Board and Commissioner, a "decision for referral not intended to be punitive" [9].

6. Rules courts and bars wrote

No special rule. The Second Circuit held in Park v. Kim that "the duties imposed by Rule 11 require that attorneys read, and thereby confirm the existence and validity of, the legal authorities on which they rely" [6]. A rule on these tools was "not necessary to inform a licensed attorney ... that she must ensure that her submissions to the Court are accurate" [6]. The Fifth Circuit's proposed rule would have required filers to certify that no generative program was used, or that "all generated text, including all citations and legal analysis, has been reviewed for accuracy and approved by a human" [6]. "Thirteen comments were submitted by judges, professors, and practitioners", many reasoning that existing rules "already impose an obligation", and the court declined the rule [20]. Its undated notice reminds filers that filings "shall be carefully checked for truthfulness and accuracy as the rules already require" [21]. In Fletcher, on 18 February 2026, the court recalled that it had concluded "existing rules were sufficient to deter misconduct related to generative AI use" [20].

Disclosure. In the edition of the Northern District of Texas's local civil rules effective 1 September 2026, "A brief prepared using generative artificial intelligence must disclose this fact on the first page", and a brief without the disclosure "certifies that no part of the brief was prepared using generative artificial intelligence" [23]. The Federal Court of Canada's notice of 7 May 2024 requires a declaration "when the role AI plays in the preparation of materials for the purpose of litigation resembles that of a co-author", and urges "verification of any AI-created content" [24].

Verification. A New South Wales practice note applies to all proceedings from 3 February 2025 [25]. Where generative tools were used for written submissions, it requires the author to verify "that all citations, legal and academic authority and case law and legislative references: (a) exist, (b) are accurate, and (c) are relevant to the proceedings" [25]. "Such verification must not be solely carried out by using a Gen AI tool or program" [25]. The note also bars "Gen AI" from "generating the content of affidavits, witness statements" [25].

Bar guidance. The State Bar of California's 2026 guidance replaces its November 2023 version and addresses "the unique challenges presented by the use of agentic AI" [26]. It says "The lawyer's duty of candor to the tribunal cannot be delegated to AI", that lawyers "must not permit AI systems to autonomously file documents", and that the risk of hallucinated outputs "can also be compounded when using agentic AI" [26]. Noland pointed to the 2023 version and said "even a superficial review of the literature would have alerted counsel to this issue" [10].

7. Measured hallucination rates in published studies

The studies measure models and tools on test questions, never filings.

- **General-purpose models.** Dahl and colleagues found legal hallucinations at 58% to 88% across the general-purpose models tested, "when these models are asked specific, verifiable questions about random federal court cases" [28]. The models "cannot always predict, or do not always know, when they are producing legal hallucinations" [28]. Zhang v. Chen cited an earlier version of this study [8].
- **Commercial legal research tools.** A preregistered 2024 evaluation found fewer hallucinations than in general-purpose chatbots, and the commercial legal research tools tested each hallucinated 17% to 33% of the time [29].
- **Detection.** A 2026 benchmark found "over 1,000 filings containing fabricated citations" and built "a dataset of 1,300 brief excerpts containing injected errors" [30]. Its headline agentic configuration "achieves 84.4% recall and a 55.0% F1 score", and "all models struggle with subtle error categories" [30].

Each figure measures a different task on systems of its time, and none converts into a rate per court filing.

8. Which check targets which failure

Failure	Check that targets it	What the check leaves open
Authority that does not exist	Look the citation up in an authoritative source or a lookup service [5, 24, 31]	Whether a real authority says what the filing claims
Quotation the source does not contain	Find the quoted words in the source text [10, 25]	Whether the quotation supports the proposition
Real authority cited for what it does not hold	Read the cited passage against the proposition [6, 25]	Subtle errors, which all models tested struggled with [30]
Authority that is not a case	Look it up where it is held [18, 31]	The lookup service documented here skips statutes and law journals [31]

Existence. Free Law Project's lookup service "can be useful as a guardrail to help prevent hallucinated citations" [31]. Its statuses separate a citation found (200), a valid citation it could not find (404), a reporter not in its system (400) and a citation matching more than one item (300) [31]. It is "throttled to 60 valid citations per minute" and "does not look up statutes, law journals, id, or supra citations" [31]. A found citation shows that a reporter, volume and page lead to a decision. It says nothing about what the decision holds.

Quotations and holdings. Courts named reading as the check. Noland held that no paper filed in any court "should contain any citations" that "the attorney responsible for submitting the pleading has not personally read and verified" [10]. Wadsworth said a researcher "must still read the case to ensure the excerpt is existing law" [11], a test that also bears on the overturned case law and repealed law the database finds in 19 decisions each [1]. The New South Wales note separates the three questions, exist, accurate and relevant, and rules out a generative tool as the only verifier [25]. The five steps of the check, from existence to faithfulness, are in our methods note on checking LLM citations against their sources (https://rillor.com/insights/checking-llm-citations-against-sources).

Implications

For lawyers and self-represented filers. The records place the check on whoever signs or files [4, 6, 11, 22]. A lawyer cannot rely on the client for citations [5], and in California the duty of candor "cannot be delegated to AI" [26]. Disclosure duties differ by court [21, 23, 24].

For firm leaders and regulators. The High Court called on leaders and regulators to act [5]. The errors came from tools firms already use: a platform the court said "appears to be an in-house database" [11] and a tool built into practice management software [9]. Clients bore costs: claims dismissed with prejudice [13] and discovery relief denied [16]. Johnson v. Dunn disqualified counsel but found no hardship to the client [15].

For builders of legal research and drafting tools. A lookup and a reading test different failures (Finding 8). California's guidance says lawyers "must not permit AI systems to autonomously file documents" [26], New South Wales rules out a generative tool as the only verifier [25], and in the detection benchmark "all models struggle with subtle error categories" [30].

For teams that turn legal and regulatory text into records. Courts checked citations by going to the source: the Mata court through the Clerk of the Eleventh Circuit [4], the Lacey Special Master [16] and the judge in Ko v. Li [19]. A record that keeps its source passage allows the same check. Rillor's document and regulation intelligence work turns rules, contracts and filings into records a system can check, with each field linked to the passage it came from and a person reviewing every record.

Limits of this analysis

- The database is an undercount by its author's account. It holds only decisions where a court found or implied reliance on hallucinated content, plus some alleged cases [1, 3]. No source gives the number of fabricated citations in all filings or a rate per filing.
- Counts change daily. Those here are from the update of 9 November 2026, viewed 10 November 2026 [1].
- Party types and failure categories overlap. Figures marked "by our count" are this report's arithmetic on the CSV export [2].
- The U.S. weight may reflect access to court records [3], and the latest months may be incomplete.
- The monetary field includes fee awards, uses 1 for unknown amounts and holds nine currencies, none converted here [1, 2].
- The seventeen decisions were chosen to show each kind of finding and remedy, not drawn as a sample, and several were read in the copy the database hosts.
- ABA Formal Opinion 512 and the judicial guidance of England and Wales refused automated access and are not covered. The two judges' letters and orders and the federal courts' interim guidance were not opened, so only the Senate release's words are used [27].
- No source reviewed measures whether sanctions deter.
- This report describes what courts did and said. It is research, not legal advice.

Sources

1. AI Hallucination Cases Database. Damien Charlotin, last updated 9 November 2026, viewed 10 November 2026. https://www.damiencharlotin.com/hallucinations/

2. AI Hallucination Cases Database, CSV export. Damien Charlotin, last updated 9 November 2026, viewed 10 November 2026. https://www.damiencharlotin.com/hallucinations/hallucinations/download.csv

3. Hallucinations Case Database: FAQ. Damien Charlotin, Artificial Authority, 14 November 2025, edited February 2026, viewed 8 November 2026. https://artificialauthority.ai/p/hallucinations-case-database-faq

4. Mata v. Avianca, Inc., No. 22-cv-1461 (PKC), Opinion and Order on Sanctions, ECF No. 54. U.S. District Court for the Southern District of New York, 22 June 2023, viewed 8 November 2026. https://storage.courtlistener.com/recap/gov.uscourts.nysd.575368/gov.uscourts.nysd.575368.54.0.pdf

5. R (Ayinde) v London Borough of Haringey, and Al-Haroun v Qatar National Bank QPSC, [2025] EWHC 1383 (Admin). High Court of Justice, King's Bench Division, Divisional Court, 6 June 2025, viewed 8 November 2026. https://caselaw.nationalarchives.gov.uk/ewhc/admin/2025/1383

6. Park v. Kim, No. 22-2057. U.S. Court of Appeals for the Second Circuit, 30 January 2024, viewed 8 November 2026. https://www.damiencharlotin.com/documents/243/Park_v._Kim_US_2nd_Cir._CA_January_30_2024.pdf

7. Harber v The Commissioners for His Majesty's Revenue and Customs, [2023] UKFTT 1007 (TC). First-tier Tribunal (Tax Chamber), 4 December 2023, viewed 8 November 2026. https://caselaw.nationalarchives.gov.uk/ukftt/tc/2023/1007

8. Zhang v. Chen, 2024 BCSC 285. Supreme Court of British Columbia, 20 February 2024, viewed 8 November 2026. https://www.damiencharlotin.com/documents/772/Zhang_v._Chen.pdf

9. Dayal [2024] FedCFamC2F 1166. Federal Circuit and Family Court of Australia (Division 2), 27 August 2024, viewed 8 November 2026. https://www.damiencharlotin.com/documents/361/In_Re_Dayal_Australia_August_27_2024.pdf

10. Noland v. Land of the Free, L.P., B331918, 114 Cal.App.5th 426. Court of Appeal of the State of California, Second Appellate District, 12 September 2025, viewed 8 November 2026. https://storage.courtlistener.com/pdf/2025/09/12/noland_v._land_of_the_free_l.p..pdf

11. Wadsworth v. Walmart Inc., No. 2:23-CV-118-KHR, Order on sanctions, ECF No. 181. U.S. District Court for the District of Wyoming, 24 February 2025, viewed 8 November 2026. https://storage.courtlistener.com/recap/gov.uscourts.wyd.64014/gov.uscourts.wyd.64014.181.0_1.pdf

12. Shahid v. Esaam, A25A0196. Court of Appeals of Georgia, 30 June 2025, viewed 8 November 2026. https://www.damiencharlotin.com/documents/530/Shahid_v._Esaam_Georgia_30_June_2025.pdf

13. Couvrette v. Wisnovsky, No. 1:21-cv-00157-CL, Opinion and Order, ECF No. 215. U.S. District Court for the District of Oregon, 12 December 2025, viewed 8 November 2026. https://www.damiencharlotin.com/documents/1170/Couvrette_v._Wisnoksy_USA_12_December_2025.pdf

14. Couvrette v. Wisnovsky, No. 1:21-cv-00157-CL, Opinion and Order, ECF No. 225. U.S. District Court for the District of Oregon, 23 March 2026, viewed 8 November 2026. https://www.damiencharlotin.com/documents/1790/Dkt_225_Opinion_and_Order.pdf

15. Johnson v. Dunn, No. 2:21-cv-1701-AMM, Sanctions Order, ECF No. 204. U.S. District Court for the Northern District of Alabama, 23 July 2025, viewed 8 November 2026. https://storage.courtlistener.com/recap/gov.uscourts.alnd.179677/gov.uscourts.alnd.179677.204.0.pdf

16. Lacey v. State Farm General Insurance Co., No. 2:24-cv-05205-FMO-MAA, Order of the Special Master, ECF No. 119. U.S. District Court for the Central District of California, 6 May 2025, viewed 8 November 2026. https://www.damiencharlotin.com/documents/311/Lacey_v._State_Farm_General_Insurances_Co._D._Cal._May_6_2025.pdf

17. Mid Central Operating Engineers Health and Welfare Fund v. HoosierVac LLC, No. 2:24-cv-00326-JPH-MJD, Order on Report and Recommendation, ECF No. 121. U.S. District Court for the Southern District of Indiana, 28 May 2025, viewed 8 November 2026. https://www.damiencharlotin.com/documents/453/Mid_Cent._Operating_Engrs_Health_v._Hoosiervac_S.D._Indiana_USA_May_28_2025.pdf

18. Kohls v. Ellison, No. 24-cv-3754 (LMP/DLM), Order, ECF No. 46. U.S. District Court for the District of Minnesota, 10 January 2025, viewed 8 November 2026. https://storage.courtlistener.com/recap/gov.uscourts.mnd.220348/gov.uscourts.mnd.220348.46.0.pdf

19. Ko v. Li, 2025 ONSC 2965. Ontario Superior Court of Justice, 20 May 2025, viewed 8 November 2026. https://www.damiencharlotin.com/documents/451/Ko_v._Li_Ontario_SC_Canada_May_28_2025.pdf

20. Fletcher v. Experian Information Solutions, Inc., No. 25-20086. U.S. Court of Appeals for the Fifth Circuit, 18 February 2026, viewed 8 November 2026. https://www.damiencharlotin.com/documents/1539/Fletcher_v._Experian_USA_18_February_2026.pdf

21. Court Decision on Proposed Rule. U.S. Court of Appeals for the Fifth Circuit, undated, viewed 8 November 2026. https://www.ca5.uscourts.gov/docs/default-source/default-document-library/court-decision-on-proposed-rule.pdf?sfvrsn=5967c92d_2

22. Federal Rules of Civil Procedure, Rule 11. Legal Information Institute, Cornell Law School, viewed 8 November 2026. https://www.law.cornell.edu/rules/frcp/rule_11

23. Local Civil Rules, LR 7.2(f). U.S. District Court for the Northern District of Texas, edition effective 1 September 2026, viewed 8 November 2026. https://www.txnd.uscourts.gov/sites/default/files/documents/CIVRULES.pdf

24. Notice to the Parties and the Profession: The Use of Artificial Intelligence in Court Proceedings. Federal Court (Canada), 7 May 2024, viewed 8 November 2026. https://www.fct-cf.gc.ca/Content/assets/pdf/base/FC-Updated-AI-Notice-EN.pdf

25. Practice Note SC Gen 23: Use of Generative Artificial Intelligence (Gen AI). Supreme Court of New South Wales, 28 January 2025, viewed 8 November 2026. https://supremecourt.nsw.gov.au/documents/Practice-and-Procedure/Practice-Notes/general/current/PN_SC_Gen_23.pdf

26. Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law. State Bar of California, Standing Committee on Professional Responsibility and Conduct, 2026, viewed 8 November 2026. https://www.calbar.ca.gov/Portals/0/documents/ethics/Generative-AI-Practical-Guidance.pdf

27. Grassley Releases Judges' Responses Owning Up to AI Use, Calls for Continued Oversight and Regulation. U.S. Senate Committee on the Judiciary, 23 October 2025, viewed 8 November 2026. https://www.judiciary.senate.gov/press/rep/releases/grassley-releases-judges-responses-owning-up-to-ai-use-calls-for-continued-oversight-and-regulation

28. Dahl, Magesh, Suzgun and Ho. Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models. Journal of Legal Analysis 16(1), arXiv 2401.01301 v2, 21 June 2024, viewed 8 November 2026. https://arxiv.org/abs/2401.01301

29. Magesh, Surani, Dahl, Suzgun, Manning and Ho. Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools. arXiv 2405.20362 v1, 30 May 2024, viewed 8 November 2026. https://arxiv.org/abs/2405.20362

30. Liu, Stammbach and Henderson. Who Checks the Citations? Benchmarking Legal Hallucination Detection. arXiv 2606.21155 v3, 4 October 2026, viewed 8 November 2026. https://arxiv.org/abs/2606.21155

31. Citation Lookup and Verification API. Free Law Project, CourtListener, viewed 8 November 2026. https://www.courtlistener.com/help/api/rest/citation-lookup/

About this report

Published by Rillor on 11 November 2026. Superintelligence (SI) tools assist research and drafting, and every figure cites its source. Rillor builds and secures SI systems, supplies the data they learn from and delivers the infrastructure they run on. We apply SI to financial and compute markets.

Cite as: Rillor. What court records show about fabricated citations in LLM output. 11 November 2026. https://rillor.com/insights/fabricated-citations-court-records

Notices

Superintelligence (SI) outputs can be wrong. Important decisions should include human review.

Work with Rillor

Send a sample of the rulebooks, regulations, contracts or filings you work from and the fields you need from them. Contact us at rillor.com/contact?type=document-intelligence or contact@rillor.com. More on document and regulation intelligence: rillor.com/si/document-intelligence