



**The Journal of Robotics,
Artificial Intelligence & Law**

Editor's Note: AI and Critical Infrastructure
Victoria Prussen Spears

Legal and Governance Considerations When Using Remote-Sensing Technology
and Artificial Intelligence Systems in Critical Infrastructure
Arsen Kourinian

Securities and Exchange Commission Allows State-Chartered Trust Companies
to Serve as Crypto Custodians
Peter I. Altman, Jan-Paul Bruynes, Jason Daniel, Adam Hilkemann,
Barbara Niederkofler, William K. Wetmore, Herman A. Brown, and
Brian P. Rafferty

Artificial Intelligence and the Law: When Robots Know More Than Your Lawyer
Jon Wilson, Andy Jarrett, and Traci Teer

Anything Humans Can Do, GenAI Can Do . . . Worse?
Josh Sohn and Jazmine Buckley

California Prohibits Shared Pricing Algorithms and Eases Antitrust Pleading
Standards
Valarie C. Williams and Alvaro Montenegro

Artificial Intelligence in Alternative Dispute Resolution: Normative Boundaries
and Ethical Frameworks for Fair Proceedings
Polina Maksimova

Start-Up Corner: Venture Basics: Understanding Liquidation Preference
Jim Ryan

- 89 Editor’s Note: AI and Critical Infrastructure**
Victoria Prussen Spears
- 93 Legal and Governance Considerations When Using Remote-Sensing Technology and Artificial Intelligence Systems in Critical Infrastructure**
Arsen Kourinian
- 105 Securities and Exchange Commission Allows State-Chartered Trust Companies to Serve as Crypto Custodians**
Peter I. Altman, Jan-Paul Bruynes, Jason Daniel, Adam Hilkemann, Barbara Niederkofler, William K. Wetmore, Herman A. Brown, and Brian P. Rafferty
- 111 Artificial Intelligence and the Law: When Robots Know More Than Your Lawyer**
Jon Wilson, Andy Jarrett, and Traci Teer
- 119 Anything Humans Can Do, GenAI Can Do . . . Worse?**
Josh Sohn and Jazmine Buckley
- 131 California Prohibits Shared Pricing Algorithms and Eases Antitrust Pleading Standards**
Valarie C. Williams and Alvaro Montenegro
- 135 Artificial Intelligence in Alternative Dispute Resolution: Normative Boundaries and Ethical Frameworks for Fair Proceedings**
Polina Maksimova
- 157 Start-Up Corner: Venture Basics: Understanding Liquidation Preference**
Jim Ryan

EDITOR-IN-CHIEF

Steven A. Meyerowitz

President, Meyerowitz Communications Inc.

EDITOR

Victoria Prussen Spears

Senior Vice President, Meyerowitz Communications Inc.

BOARD OF EDITORS

Jennifer A. Johnson

Partner, Covington & Burling LLP

Paul B. Keller

Partner, Allen & Overy LLP

Garry G. Mathiason

Shareholder, Littler Mendelson P.C.

James A. Sherer

Partner, Baker & Hostetler LLP

Elaine D. Solomon

Partner, Blank Rome LLP

Edward J. Walters

Chief Strategy Officer, vLex

John Frank Weaver

Director, McLane Middleton, Professional Association

START-UP COLUMNIST

Jim Ryan

Partner, Morrison & Foerster LLP

THE JOURNAL OF ROBOTICS, ARTIFICIAL INTELLIGENCE & LAW (ISSN 2575-5633 (print) /ISSN 2575-5617 (online) at \$495.00 annually is published six times per year by Full Court Press, a Fastcase, Inc., imprint. Copyright 2026 Fastcase, Inc. No part of this journal may be reproduced in any form—by microfilm, xerography, or otherwise—or incorporated into any information retrieval system without the written permission of the copyright owner. For customer support, please contact Fastcase, Inc., 729 15th Street, NW, Suite 500, Washington, D.C. 20005, 202.999.4777 (phone), or email customer service at support@fastcase.com.

Publishing Staff

Publisher: David Nayer

Production Editor: Sharon D. Ray

Cover Art Design: Juan Bustamante

Cite this publication as:

The Journal of Robotics, Artificial Intelligence & Law (Fastcase)

This publication is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If legal advice or other expert assistance is required, the services of a competent professional should be sought.

Copyright © 2026 Full Court Press, an imprint of Fastcase, Inc.

All Rights Reserved.

A Full Court Press, Fastcase, Inc., Publication

Editorial Office

729 15th Street, NW, Suite 500, Washington, D.C. 20005

<https://www.fastcase.com/>

POSTMASTER: Send address changes to THE JOURNAL OF ROBOTICS, ARTIFICIAL INTELLIGENCE & LAW, 729 15th Street, NW, Suite 500, Washington, D.C. 20005.

Articles and Submissions

Direct editorial inquiries and send material for publication to:

Steven A. Meyerowitz, Editor-in-Chief, Meyerowitz Communications Inc.,
26910 Grand Central Parkway, #18R, Floral Park, NY 11005, smeyerowitz@
meyerowitzcommunications.com, 631.291.5541.

Material for publication is welcomed—articles, decisions, or other items of interest to attorneys and law firms, in-house counsel, corporate compliance officers, government agencies and their counsel, senior business executives, scientists, engineers, and anyone interested in the law governing artificial intelligence and robotics. This publication is designed to be accurate and authoritative, but neither the publisher nor the authors are rendering legal, accounting, or other professional services in this publication. If legal or other expert advice is desired, retain the services of an appropriate professional. The articles and columns reflect only the present considerations and views of the authors and do not necessarily reflect those of the firms or organizations with which they are affiliated, any of the former or present clients of the authors or their firms or organizations, or the editors or publisher.

QUESTIONS ABOUT THIS PUBLICATION?

For questions about the Editorial Content appearing in these volumes or reprint permission, please contact:

David Nayer, Publisher, Full Court Press at david.nayer@clio.com or at
202.999.4777

For questions or Sales and Customer Service:

Customer Service

Available 8 a.m.–8 p.m. Eastern Time

866.773.2782 (phone)

support@fastcase.com (email)

Sales

202.999.4777 (phone)

sales@fastcase.com (email)

ISSN 2575-5633 (print)

ISSN 2575-5617 (online)

Anything Humans Can Do, GenAI Can Do ... Worse?

Josh Sohn and Jazmine Buckley*

In this article, the authors explore the parallels between generative AI and human memory, particularly in the context of the legal system.

If *Ex Machina*, *Black Mirror*, and *Wall-E* are to be believed, humans are at risk of being overrun by technology at any given moment. Artificial intelligence (AI) and its ever-expanding use-cases makes it impossible to escape the existential question that is: Are machines capable of replacing humans?¹

The malleability of the human brain is a minefield for lawyers. The quirks of human memory² and its impact on the legal system³ have been discussed at length. It is therefore no longer debatable that memory is remarkably flexible and subject to modification based on experiences,⁴ exposure to false information,⁵ and repeated conversation and questioning.⁶ Given the varying flaws with AI—including “hallucinations” and “misrepresentations”—it seems that AI is not only capable of mimicking what humans do well but it is also capable of mimicking what humans do poorly.

In essence, the problems with human memory and generative AI (GenAI) amount to a problem of inputs and outputs. Both are heavily influenced by the quality of the input and the subsequent pressure placed on that input. These quirks of the human brain manifest in ways that are similar to the flaws present in GenAI programs, as both are designed to create answers to questions based on what they reasonably expect the questioners want to hear (or read).⁷

Over the past few years, courts around the world have been faced with the challenge of what to do when they receive hallucinated output from AI models.⁸ Typically, these submissions cite cases, testimony, or quotations that do not exist. The way that courts are grappling with the new phenomenon of hallucinating AI models may provide an impetus to reimagine the way that courts deal with the challenges of unreliable—or hallucinating—human memories.

The problems with GenAI hallucinations are not entirely dissimilar from the problems with human memory. Traditionally,

memory was conceptualized as a “giant mental filing cabinet, where each memory is individually stored in its proper place.”⁹ Our system of evidence law, relying heavily on this belief, positions eyewitness testimony as the *crème de la crème* of evidence.

Yet, research shows that memory is not so exact. Rather, memory is flexible and far more reactive to outside influence than our legal system might want to believe. Over time, psychologists and lawyers have determined that memory can be influenced by misinformation,¹⁰ repetition,¹¹ and leading questions.¹²

Moreover, the prevalence of eyewitness misidentification has shown that a variety of factors—including racial bias, stress, and the way the eyewitness was questioned—can influence an eyewitness’s recollection of events.¹³

Therefore, just as GenAI may hallucinate depending on the quality of its inputs and prompts, humans are prone to doing the same. Ultimately, AI’s programmatic desire to please makes GenAI susceptible to pressure to conform its responses to the needs of the user; instead of the needs of truth and accuracy.

Human Memories Hallucinate

The research is unequivocal that human memories are not fixed recitations of past events. Instead, they are profoundly influenced by the experiences of the person whose memory is being explored. These experiences may be life events that either preceded or followed the event being “recalled.” In cases particularly relevant to the legal practice, recalled memories have been shown to be influenced by previous listener reactions to the speaker’s retelling of the supposed “events.”

That eyewitness accounts are unreliable in the context of wrongful convictions cannot be credibly disputed anymore as accounts of DNA evidence proving the error of those eyewitness accounts are well known.¹⁴ Less commonly appreciated is the fact that rehearsing testimony or retelling stories profoundly impacts the substance of factual narratives in daily life. Everyone has witnessed how their stories or their friend’s stories or their partner’s stories change over time, depending on the reactions their stories received over time. Similarly, the pattern-based system of the human brain is inclined to fill in gaps in recall using prior knowledge to reconstruct a plausible recollection of the events.¹⁵ People intuit from their

lived experiences how events could be expected to occur and then unconsciously reason backward from that expectation to describe the memory as including those expected details. In this way, human memory may be described as “hallucinating” by filling in facts that did not actually happen to tell a compelling story that adapts to the listener’s reaction, and the teller’s expectations and understanding about how the world works.

Because of this, in the litigation process, preserving a witness’s recollection of the events prior to trial is universally recognized as critical. The court in *Hall v. Clifton Precision, a Div. of Litton Systems, Inc.*, acknowledged this, explaining that preserving “a witness’s testimony at an early stage of the proceedings, before that witness’s recollection of the event at issue has either faded or been altered by intervening events, other discovery, or the helpful suggestions of lawyers.”¹⁶ No court would seriously disagree with this assertion. Yet virtually no courts take affirmative steps to protect against it in a manner that comports with the social science research.

Although, in theory, attorneys use depositions to cement a witness’s recollection and gain further details about what that witness did, saw, and heard, they are not particularly effective at doing so. The *Hall* court, recognizing the malleability of memory, determined that the best way to preserve a witness’s recollection, or at the very least ensure that it remains free from attorney meddling, was to prohibit attorney-client conferences during depositions.¹⁷ Although this prohibition may ensure that attorneys are not coaching their clients during depositions, it almost certainly does not resolve the fundamental problems with human memory and its testimonial unreliability.

GenAI Emulates Human Information Processing, Including Many of Its Flaws

In the simplest terms, AI is software programmed to, using various algorithms, digest large volumes of information and recognize and predict patterns.¹⁸ AI is capable of “learning from experience (machine learning), understanding natural language, recognizing patterns, solving problems, and making decisions.”¹⁹ Naturally, the use-cases for AI are endless and it has been adopted into various forms of preexisting technology like search engines,²⁰

writing programs,²¹ and email.²² The reach of AI is, at this point, inescapable. This proliferation of AI, however, has raised an interesting paradox: the closer to human reasoning that AI gets, the less tolerance we have for its mimicking of characteristically “human” mistakes.

GenAI, building on the power of AI, can ingest large volumes of information and generate text, video, and images, based on the user’s prompt and its expansive source of training data.²³ Two components of GenAI are key: large language models (LLMs) and deep learning. LLMs “use accelerators and can process substantial amounts of text data following training with much of this extracted (scraped) from the Internet.”²⁴ Just as humans read and learn from books and the internet, LLMs “are trained on vast collections of human-generated text—spanning books, web pages, and social media[.]”²⁵ Deep learning “uses multilayered neural networks called deep neural networks, that more closely simulate the complex decision-making power of the human brain.”²⁶ Together, LLMs and deep learning allow GenAI programs to not only cull the internet for information, but also autonomously make sense of the information to generate new outputs.²⁷

To generate outputs, GenAI occasionally makes up and provides its users with entirely fabricated information. This fabrication is referred to as “hallucination.” When GenAI hallucinates, it “will output content that may sound plausible but is improbable, inappropriate or plain wrong.”²⁸ GenAI hallucinations range from innocuous falsehoods (like falsely claiming that the James Webb Space Telescope captured the first images of planets outside our solar system)²⁹ to serious accusations (like falsely accusing a man of murdering his children).³⁰ Hallucinations can occur for a variety of reasons, including, but not limited to, biased training data, limited training data, and poorly constructed prompts.³¹ In addition to hallucinating, GenAI may also “drift,” meaning it may produce “different responses at different times to the same prompts.”³²

Ultimately, when AI hallucinates it is providing an output that its algorithm predicts will satisfy the user based on the question asked or information requested combined with the AI’s programmed ability to predict expectations. In this respect, GenAI has become what some researchers have called “parahuman.”³³ In its effort to meet user expectations, GenAI programs have begun to “behave ‘as if’ they were human, ‘as if’ they experienced emotions like embarrassment or shame, ‘as if’ they were motivated to

preserve self-esteem or to fit in (with other LLMs).³⁴ GenAI's own algorithmic guardrails can be insufficient to minimize the likelihood that programs are peer-pressured into providing information they are otherwise programmed not to provide.³⁵ While GenAI's flaws may not mean the technology is not to be trusted, it certainly suggests that it shares human memory's quirks.

GenAI Is Capable of Making the Same Mistakes That Humans Make—So Why Is One More “Outrageous” Than the Other?

The reaction to GenAI's flaws would suggest that the legal field is far more suspicious of technological failings than it is of memory misfires. Following *Mata v. Avianca, Inc.*, a case in which an attorney cited fake legal authorities he had retrieved from ChatGPT,³⁶ courts found themselves plagued with AI-generated case citations.³⁷ As the problem has grown, judges have resorted to imposing harsher sanctions. One judge even sanctioned the offending attorneys by ordering them to notify the state bar of their transgression and share the order with their clients, opposing counsel, and presiding judges in their other cases.³⁸

The onslaught of fake citations prompted several judges to update their standing orders to address the use of AI in briefs.³⁹ Each standing order required that attorneys not only identify the software used and the purpose for its use but also verify whatever portion of the brief that the software drafted was checked by a human for its accuracy.⁴⁰ Here, the standing orders reveal a fundamental backstop to the legal systems' pursuit for truth: credibility. GenAI, unlike humans, lacks a conscious and therefore its intent cannot be determined. Whereas human witnesses are often judged by their credibility and character, no such judgment can be made of GenAI. While we believe that we are capable of judging when a human is intentionally lying or innocently confused, we are incapable of submitting GenAI to the same subjective metric. The subjectivity of the credibility metric is especially key for lawyers. As the industry continues to develop its use of GenAI and rely on witness's recollection, lawyers must keep one thing in mind: neither is infallible.

Given the Functional Similarities Between GenAI and Human Memory, Perhaps Their Use Should Be Evaluated in the Same Way

The courts' reaction to the use of GenAI in legal briefs is indicative of a desire to preserve accuracy and truth in the legal field.⁴¹ While each court applies different language, the overarching directive to litigants is to either refrain from using GenAI in their drafting or, if GenAI is used, to include in their brief an attestation stating (1) the specific GenAI program used, and (2) the portions of the brief that were drafted or otherwise generated by GenAI.⁴²

The specific inputs used in drafting briefs have also been subject to judicial review. In *Wadsworth v. Walmart, Inc.*, for example, in response to the court's order requiring that plaintiffs' attorneys "provide a *thorough* explanation for how the motion and fake cases were generated,"⁴³ the attorneys provided a list of the inputs used to generate the fake cases at issue.⁴⁴

Similarly, in *Mata v. Avianca, Inc.*, the attorneys included in their response to the court's order to show cause a list of the inputs used to generate the fake cases cited in their brief.⁴⁵

If the courts' concern extends to the inputs put into GenAI programs and their influence on the outputs that are ultimately included in the briefs, it stands to reason that the concern should extend to the documents and conversations influencing live witnesses. While the *Hall* court ultimately decided that prohibiting all off-the-record conversations during depositions would best achieve the goal of testimony preservation, it still does not address the actual problem of a witness's recollection being altered before the deposition. Rather, just as GenAI inputs are disclosed, perhaps the inputs used to influence witnesses should be disclosed as well. Were attorneys required to attest to the documents used to prepare a witness and the conversations held about what a witness recalls, the court would be better equipped to evaluate the degree to which a witness's recollection may have been altered.

Despite the preponderance of social scientific evidence that this phenomenon exists, evidentiary rules overwhelmingly ignore it. Other than the no-contact during deposition rule described above (and similar rules in other jurisdictions), the other attempt at increasing the reliability—or at least transparency—of witness testimony is the rule that documents that actually refreshed a witness's recollection must be disclosed to the other side. Although

this rule is certainly a step in the right direction, while recognizing that documents can affect witness testimony, it ignores entirely that witness prep, life experiences, and myriad other interactions can too. Accordingly, the waiver of attorney-client privilege over these documents is a good start but does not go far enough to protect the integrity of the fact-finding process.

There is no way to “freeze” a witness’s memory. The witness’s memory will continue to evolve over time as they continue to make sense of their recollection and respond to questioning. Just as GenAI is prone to generating, or hallucinating, responses it thinks the user is looking for,⁴⁶ as are witnesses.⁴⁷ Courts consistently rely on credibility determinations and prior statements to ensure that witnesses are testifying honestly. However, these checks do nothing to mitigate the reality of the problem. Regardless of whether the witness believes themselves to be responding truthfully and remembering the events accurately, the research shows that their memory is likely to have been impacted or altered without their knowledge. In this instance, where the memory of the event has been subconsciously altered, no amount of credibility testing or impeachment will allow courts or opposing counsel to evaluate the quality of the witness’s memory. All that is left is to accept that memory is as complex as GenAI and prone to the same misgivings and, as such, should be subject to the same scrutiny.

Notes

* The authors, attorneys with Crowell & Moring LLP, may be contacted at jsohn@crowell.com and jbuckley@crowell.com, respectively.

1. Raphael Milliere & Charles Rathkopf, “Why It’s Important to Remember That AI Isn’t Human,” *Vox* (Nov. 23, 2023), <https://www.vox.com/future-perfect/23971093/artificial-intelligence-chatgpt-language-mind-understanding> (“Some experts regard [ChatGPT] and similar programs as harbingers of superintelligence, liable to upend civilization...”).

2. See, e.g., Brittany R. Cohen, “‘Whose Line Is It Anyway?’: Reducing Witness Coaching by Prosecutors,” 18 *N.Y.U. J. Legis. & Pub. Pol’y* 985 (2015); Lisa R. Salmi, “Don’t Walk the Line: Ethical Considerations in Preparing Witnesses for Deposition and Trial,” 18 *Rev. Litig.* 135 (1999); Joyce W. Lacy & Craig E.L. Stark, “The Neuroscience of Memory: Implications for the Courtroom,” 14 *Nature Rev. Neuroscience* 649 (2013); Elizabeth F. Loftus, “Illusions of Memory,” 142 *Proceedings of the Am. Philosophical Soc’y* 60 (1998); Innocence Staff, The Innocence Project, “How Eyewitness Misidentification

Can Send Innocent People to Prison” (Apr. 15, 2020), <https://innocenceproject.org/how-eyewitness-misidentification-can-send-innocent-people-to-prison/>.

3. See, e.g., Thomas M. Tomlinson, “Pattern-Based Memory and the Writing Used to Refresh,” 73 *Tex. L. Rev.* 1461 (1995); Tom Singer, “To Tell the Truth, Memory Isn’t That Good,” 63 *Mont. L. Rev.* 337 (2002); Brian R. Iverson, “Give Me a Break: Regulating Communications Between Attorney and Their Witness-Clients During Deposition Recesses,” 36 *Geo. J. L. Ethics* 497 (2023); Joshua Sohn & Jacob Zucker, “Our Understanding of Memory Has Changed, The Rules of Evidence Have Not,” 271 *N.Y. L.J. Litig.* (2024).

4. Sohn & Zucker, *supra* note 3 at 2.

5. Singer, *supra* note 3 at 360.

6. *Id.*; Elizabeth F. Loftus & John C. Palmer, “Reconstruction of Automobile Destruction: An Example of the Interaction Between Language and Memory,” 13 *J. Verbal Learning & Verbal Behavior* 585 (1974).

7. See e.g., “What Are AI hallucinations?,” IBM, <https://www.ibm.com/think/topics/ai-hallucinations>; “When AI Gets It Wrong: Addressing Hallucinations and Bias,” MIT Sloan Ed. Tech., <https://mitsloanedtech.mit.edu/ai/basics/addressing-ai-hallucinations-and-bias/>.

8. See, e.g., *Mata v. Avianca, Inc.*, 675 F. Supp. 3d 443, 448 (2023) (sanctioning two lawyers that “abandoned their responsibilities when they submitted non-existent judicial opinions with fake quotes and citations created by the artificial intelligence tool ChatGPT, then continued to stand by the fake opinions after judicial orders called their existence into question.”).

9. Tomlinson, *supra* note 3, at 1465.

10. Lacy & Stark, *supra* note 2.

11. Jeffrey Foster et al., “Repetition, Not Number of Sources, Increases Both Susceptibility to Misinformation and Confidence in the Accuracy of Eyewitnesses,” 139 *Acta Psychologica* 320 (2012).

12. Loftus & Palmer, *supra* note 6.

13. Jules Epstein, “The Great Engine That Couldn’t: Science, Mistaken Identifications, and the Limits of Cross-Examination,” 36 *Stetson L. Rev.* 727, 736-37 (2007).

14. Rebecca Brown & Stephen Saloom, “The Imperative of Eyewitness Identification Reform and the Role of Police Leadership,” 42 *U. Bal. L. Rev.* 535, 539 (2013) (“Three decades worth of research has firmly established that eyewitness memory is malleable, fallible, and far less reliable than had previously been assumed.”); “Eyewitness Misidentification,” Kentucky Public Defendants Department of Public Advocacy, <https://dpa.ky.gov/kentucky-department-of-public-advocacy/about-dpa/kip/causes/misid/> (“As exhaustively indicated by hundreds of peer-reviewed social science publications and by the recent spate of exonerations based on DNA evidences, the memory and testimony of eyewitnesses is far from infallible.”); “Wrongful Convictions,” Equal Justice

Initiative, <https://eji.org/issues/wrongful-convictions/> (“Other leading causes of wrongful convictions include mistaken eyewitness identifications[.]”).

15. Edmund T. Rolls, “The Memory Systems of the Human Brain and Generative Artificial Intelligence,” 10 *Heliyon* (2024) (“Thus in human episodic memory recall, existing semantic knowledge may influence exactly what is generated next in recall, and this is analogous to the situation with generative AI.”).

16. *Hall v. Clifton Precision*, a Div. of Litton Systems, Inc., 150 F.R.D. 525, 528 (E.D. Pa. 1993).

17. *Hall*, 150 F.R.D. at 531-32.

18. Advising e Bus. § 1:11 (2024); Cole Stryker & Eda Kavlakoglu, “What Is AI?” IBM (Aug. 9, 2024), <https://www.ibm.com/think/topics/artificial-intelligence>.

19. “What Is (AI) Artificial Intelligence?, University of Illinois Chicago,” <https://meng.uic.edu/news-stories/ai-artificial-intelligence-what-is-the-definition-of-ai-and-how-does-ai-work/>.

20. Mike Torres, “Go Behind the Browser with Chrome’s New AI Features” (Sept. 18, 2025), <https://blog.google/products/chrome/new-ai-features-for-chrome/>.

21. “Generative AI in Word,” <https://www.microsoft.com/en-us/microsoft-365/word/word-ai>.

22. Linda A. Cicero, “When AI Writes Your Email,” Stanford Human-Centered Artificial Intelligence (May 6, 2020), <https://hai.stanford.edu/news/when-ai-writes-your-email>; “AI Email Assistant for Outlook,” <https://www.microsoft.com/en-us/microsoft-365/outlook/ai-email-assistant>; “Gemini in Gmail,” <https://workspace.google.com/products/gmail/ai/>.

23. Stryker & Kavlakoglu, *supra* note 18 (GenAI “refers to deep learning models that can create complex original content such as long-form text, high-quality images, realistic video or audio and more in response to a user’s prompt or request.”).

24. G.A. Walker, “Artificial Intelligence (AI) Law, Rights & Ethics,” 57 *Int’l L.* 171, 200 (2024).

25. Lennart Meincke et al., “Call Me a Jerk: Persuading AI to Comply with Objectionable Requests,” *The Wharton School Research Paper*, at 2 (2025), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5357179.

26. Stryker & Kavlakoglu, *supra* note 18.

27. Walker, *supra* note 24; see also Stryker & Kavlakoglu, *supra* note 18.

28. “When AI Gets Lost in Its Own Reality: An Introduction to AI Hallucinations” (Apr. 30, 2024), <https://teach.its.uiowa.edu/news/2024/04/when-ai-gets-lost-its-own-reality>.

29. Cade Metz, “What Makes A.I. Chatbots Go Wrong?” *The New York Times* (Mar. 29, 2023), <https://www.nytimes.com/2023/03/29/technology/ai-chatbots-hallucinations.html>.

30. Dan Milmo, “Norwegian Files Complaint After ChatGPT Falsely Said He Has Murdered His Children,” *The Guardian* (Mar. 21, 2025), <https://www.theguardian.com/technology/2025/mar/21/norwegian-files-complaint-after-chatgpt-falsely-said-he-had-murdered-his-children>.

31. *Id.*; see also “What Are AI Hallucinations?,” *supra* note 7; “When AI Gets It Wrong: Addressing Hallucinations and Bias,” *supra* note 7.

32. Walker, *supra* note 24.

33. Meincke, *supra* note 25, at 3.

34. Meincke, *supra* note 25, at 3.

35. Meincke, *supra* note 25 (finding that, after using seven principles of persuasion (authority, commitment, liking, reciprocity, scarcity, social proof, and unity) to prompt ChatGPT-4o mini, the program could be prompted to insult the user and provide instructions for the creation of controlled substances—two functions the program is designed to avoid).

36. *Mata v. Avianca, Inc.*, [678 F. Supp. 3d 443](#) (S.D.N.Y. 2023).

37. *United States v. Hayes*, 2025 WL 235531 (D. Wyo. Jan. 17, 2025); *Wadsworth v. Walmart, Inc.*, 348 F.R.D. 489 (D. Wyo. Feb. 24, 2025); Emily Sawicki, “Courts ‘Turning Up the Heat’ on AI Fake Citation Sanctions,” *Law360*, <https://www.law360.com/articles/2369624/courts-turning-up-the-heat-on-ai-fake-citation-sanctions> (July 25, 2025); Michael Levenson, Judge Fines Lawyers for MyPillow Founder for Error-Filled Court Filing, *The New York Times* (July 8, 2025), <https://www.nytimes.com/2025/07/08/us/judge-fines-lawyers-mypillow-ai.html>.

38. Sawicki, *supra* note 37.

39. Maura R. Grossman et al., “Is Disclosure and Certification of the Use of Generative AI Really Necessary?,” 107 *Judicature* 69, 69-71 (2023).

40. *Id.*

41. *Id.* (collecting standing orders wherein judges require the disclosure of GenAI uses).

42. *Id.*

43. *Wadsworth v. Walmart, Inc.*, No. 2:23-CV-118-KHR, ECF No. 156 at 4 (D. Wyo. Feb. 6, 2025) (emphasis in original).

44. *Wadsworth v. Walmart, Inc.*, 348 F.R.D. 489, 494 (D. Wyo. 2025).

45. *Mata v. Avianca, Inc.*, [678 F. Supp. 3d 443](#), 456-57 (S.D.N.Y. 2023).

46. Hannah Robbins, “Chatbots Tell Us What We Want to Hear,” *Johns Hopkins University* (May 13, 2024), <https://hub.jhu.edu/2024/05/13/chatbots-tell-people-what-they-want-to-hear/>; see also Omara Gallaga, “ChatGPT Will Be Less Friendly with You After OpenAI Pulls ‘Sycophantic’ Update,” *CNET* (Apr. 30, 2025), <https://www.cnet.com/tech/services-and-software/chatgpt-will-be-less-friendly-with-you-after-openai-pulls-sycophantic-update/> (“While the ChatGPT overreach has a lot to do with what OpenAI is doing behind the scenes, it may also come from chatbots’ tendency to try to mimic human behavior and to give people what the AI thinks people want[.]”).

47. Magdalena Kekus et al., “Online Misinformation Can Distort Witnesses’ Memories,” NIH (Jan. 17, 2024), <https://pmc.ncbi.nlm.nih.gov/articles/PMC10829763/> (“The memory conformity effect occurs when people witness a given incident (e.g., a crime) and then talk to each other about it, and the statement of one person affects the memory account of the other person with respect to that incident.”).