Journal of Technology Law & Policy

Volume 30 Spring 2026 Number 2

13TH ANNUAL UF LAW E-DISCOVERY CONFERENCE

ARTICLES

Journal of Technology Law & Policy

Volume 30 Spring 2026 Number 2

THE COMING USE AND MISUSE OF ARTIFICIAL INTELLIGENCE IN THE COURTROOM: A JUDICIAL PERSPECTIVE AND PROPOSAL

William Matthewman[*]

Abstract

While attorneys, litigants, and judges have always had to be vigilant for the attempted introduction of improperly enhanced or fabricated evidence in court proceedings, the rise of Artificial Intelligence (AI) and its subset, Generative Artificial Intelligence (GAI), has substantially increased the stakes. The task of identifying and exposing such fake or improperly enhanced AI evidence has become significantly more challenging. The current Federal Rules of Evidence and the Federal Rules of Civil and Criminal Procedure are insufficient to effectively address the coming advance of AI-generated or AI-enhanced evidence. The rules must be amended, in certain respects, to provide attorneys, litigants, and judges with the necessary tools and guardrails to handle issues arising during pretrial and trial proceedings concerning deepfakes and AI-generated or AI-enhanced evidence. This includes addressing pretrial discovery related to such AI evidence, notice requirements when a party seeks to rely on acknowledged AI-generated or AI-enhanced evidence, procedures for pretrial challenges to admission of such evidence, respective burdens of proof, trial procedures, and even jury instructions. Lawyers and litigants must become adept at discerning and challenging unacknowledged deepfakes and other improperly AI-generated or AI-enhanced evidence, while simultaneously learning to effectively utilize properly acknowledged AI evidence for potentially legitimate purposes. Judges must carefully take a gatekeeping role when it comes to AI-generated or enhanced evidence, ensuring that only legitimate, properly authenticated evidence is admitted. New rules and procedures are needed to guide judges in dealing with challenges to the admission of AI-

* J.D. 1983, University of Florida College of Law. Judge Matthewman is the Chief United States Magistrate Judge in the United States District Court for the Southern District of Florida, and he sits in the West Palm Beach Division. He has been on the federal bench for fourteen years and has presided over jury trials, bench trials, evidentiary hearings, and all types of hearings in civil and criminal cases. While in law school, Judge Matthewman served as a Research Editor on the *Florida Law Review.* Before his appointment to the federal bench in 2012, Judge Matthewman was a very active trial lawyer and was Board Certified by The Florida Bar as a criminal trial attorney. He was lead defense counsel on the first modern federal death penalty trial in the State of Florida, lead counsel on the first modern federal death penalty trial in the U.S. District Court for the District of Puerto Rico, and he tried numerous high-profile and significant cases. He started his legal career as a police officer and then was a Sergeant in the Miami Police Department. The author acknowledges with great appreciation and thanks the assistance of Natalie Rickards, Esq., law clerk to Judge Matthewman, 2024 University of Florida Levin College of Law graduate, and member of The Florida Bar.

generated or enhanced evidence during both pretrial and trial proceedings, and to ensure that jurors are not misled by fraudulent AI evidence such as deepfakes or by false claims of deepfakes. The rapid rise of AI and GAI creates unique issues and challenges in our courtrooms and legal proceedings, requiring the courts to take the necessary steps to be ready to meet these new challenges.

INTRODUCTION

The rapid rise and development of Artificial Intelligence (AI)[1] and its subset, Generative Artificial Intelligence (GAI),[2] will have profound and widespread effects in courtrooms and upon our entire legal system.[3] The use of this new technology by attorneys and parties has already created a plethora of hallucinated, false, and fabricated submissions in filings in state and federal courts, with judges and parties struggling to deal with these new issues.[4] These hallucinations, reflected in motions and other court filings containing fabricated cases, statutes, and other misrepresentations, are just the tip of the iceberg and foreshadow of some of the negative effects that AI evidence will have in our courtrooms and in legal proceedings. But there is also the potential for positive uses of AI evidence in the courtroom, so long as such uses are properly acknowledged and authenticated. This Article addresses the anticipated

---

1. AI "refers to computer systems and applications that are capable of performing functions normally associated with human intelligence, such as abstracting, reasoning, problem solving, learning, etc." Hon. Xavier Rodriguez, *Artificial Intelligence (AI) and the Practice of Law*, 24 SEDONA CONF. J. 783, 788 (2023) (quoting CYNTHIA CWIK, PAUL W. GRIMM, MAURA GROSSMAN & TOBY WALSH, ARTIFICIAL INTELLIGENCE AND THE COURTS: MATERIAL FOR JUDGES, ARTIFICIAL INTELLIGENCE, TRUSTWORTHINESS, AND LITIGATION 6 (2022)).

2. "Generative AI is a specific subset of AI used to create new content based on training on existing data taken from massive data sources in response to a user's prompt, or to replicate a style used as input. The prompt and the new content may consist of text, images, audio, or video." *Id.* at 789 (citing Maura Grossman et al., *The GPT Judge: Justice in a Generative AI World*, 23 DUKE L. & TECH. REV. 1, 9–10 (2023).

3. For ease of reference and purposes of clarity and convenience, the author will refer to both AI-generated evidence and GAI-generated evidence collectively as AI evidence.

4. *See* ByoPlanet Int'l, LLC v. Johansson, 792 F. Supp. 3d 1341, 1356–57 (S.D. Fla. 2025) ("We live in an age when two things are happening simultaneously: (1) institutions central to our constitutional republic are suffering from a loss of trust and confidence; and (2) technology has developed to a point that few could scarcely imagine even twenty years ago. At all times, attorneys must ensure that their conduct, including their use of technology, never contributes to any diminishment of trust and confidence held by the public for the practice of law and judicial proceedings. Here, [an attorney] fell far below the standard expected, and because he did so, numerous parties and court personnel expended substantial resources getting to the bottom of his AI-fueled hallucinations."); Versant Funding LLC v. Teras Breakbulk Ocean Navigation Enters., LLC, No. 17-CV-81140, 2025 WL 1440351, at *3 (S.D. Fla. May 20, 2025) ("[T]he Court has inherent authority to sanction the misuse of AI when it affects the Court's docket, case disposition, and ruling."); Mata v. Avianca, Inc., 678 F. Supp. 3d 443, 464 (S.D.N.Y. 2023) ("[An attorney] violated Rule 11 in not reading a single case cited in his [opposition] and taking no other steps on his own to check whether any aspect of the assertions of law were warranted by existing law."); Damien Charlotin, *AI Hallucination Cases*, https://www.damiencharlotin.com/hallucinations/ [https://perma.cc/EA94-69YK] (last visited May 26, 2026). Interestingly, in response to this epidemic, the Florida Supreme Court very recently amended Florida Rules of General Practice and Judicial Administration 2.515 to state that a signer of a court filing represents that "the legal authorities identified exist and are accurately cited." Rule 2.515(d)(2) also gives the court the authority to impose sanctions for filings containing hallucinated or inaccurately cited authorities—another example of how judges may combat this problem.

use and misuse of AI-generated or enhanced evidence in the courtroom and legal proceedings.

The author, a longtime judge and former active trial attorney, takes a practical approach intended to aid lawyers, litigants, and judges who are contending with evidence generated or enhanced by AI in the courtroom and legal proceedings. Part I discusses how deepfakes and fabricated AI evidence have already invaded some courtrooms and legal proceedings, as well as concerns regarding what the future may hold in this area. Part II discusses the fact that not all AI evidence is bad; that is, acknowledged AI evidence for certain purposes will most certainly have a place in court. The question is where and how courts decide to draw the line between admissible and inadmissible AI evidence. The proper role of AI evidence will depend upon, *inter alia*, authentication, relevance, provenance, reliability, and notice. Part III discusses the role of lawyers and litigants in guarding against the use of deepfakes and other fabricated evidence in court proceedings. This part also discusses how lawyers and litigants must seek creative and acceptable ways to utilize AI in the courtroom and legal proceedings. Part IV discusses the role of judges in addressing deepfakes and fabricated AI evidence in court and in legal proceedings. Judges must take an active, careful, gatekeeping role to prevent the unfair prejudice that can arise from the introduction or mention of fabricated or improperly enhanced AI evidence. Part V discusses the need for additional rules and procedures, including a discussion as to why the recently proposed Federal Rule of Evidence 707 is misguided, because it seemingly encourages and permits the introduction of AI evidence without a testifying witness. This part also discusses why new or amended rules of evidence and criminal and civil procedure are needed to manage and address AI evidence in the courtroom and in legal proceedings. Part VI addresses the "deepfake defense" or "liar's dividend" and its potential improper effect on jurors. Specifically, considering the public's growing knowledge of deepfakes and fabricated AI evidence reported in the media and elsewhere, procedures must be implemented to guard against a party's false or frivolous denial of legitimate evidence by cavalierly claiming, without any good faith basis and without any factual or evidentiary support, that the proffered or admitted evidence is an AI-fabricated deepfake. Part VII concludes with ten core components or proposals to permit courts and judges to properly and effectively deal with AI evidence in trials and legal proceedings over the coming years. A brave new world is upon us when it comes to AI evidence in our courtrooms and legal proceedings. Attorneys, litigants, and judges must be ready for this AI onslaught, and to be ready, they must be properly educated and informed to confront the unique and novel challenges created by this new technology.

I. IT'S ALREADY HERE—DEEPFAKES AND FABRICATED, UNACKNOWLEDGED AI EVIDENCE HAVE RECENTLY SURFACED IN THE COURTROOM AND LEGAL PROCEEDINGS

The discussion of deepfakes and unacknowledged[5] AI evidence in the courtroom and legal proceedings is not a hypothetical or abstract issue. The concern is real and currently confronting various judges, courts, and legal proceedings.[6] For example, in *Mendones v.Cushman and Wakefield, Inc.*,[7] a California Superior Court Judge imposed terminating sanctions against the plaintiffs in a housing dispute after the judge found that the plaintiffs intentionally submitted fabricated evidence in connection with their motion for summary judgment. Specifically, the court found that certain audio exhibits and graphic material were deepfakes or materially altered.[8] The court questioned the authenticity of other evidentiary submissions, such as voicemail messages and video testimony, but stated that "it does not have the time, funding, or technical expertise to determine the authenticity of Plaintiffs' statements or conduct a forensic analysis of the suspect evidentiary submissions."[9] The *Mendones* court struck the operative complaint and dismissed the case with prejudice.[10] The court stated that the terminating sanction imposed "serves the appropriately chilling message to litigants appearing before this Court: Use GenAI in court with great caution."[11] This case is just the tip of the coming iceberg. While the alleged deepfakes were discovered in this case, it begs the question as to whether unacknowledged AI evidence and deepfakes in any other cases have evaded scrutiny and have been introduced into evidence without knowledge of the opposing party or the court. As the saying goes, we don't know what we don't know. A sufficient level of technical and legal expertise regarding the use of AI in legal proceedings on behalf of attorneys, litigants, and judges is vital.

Issues surrounding unacknowledged AI evidence and deepfakes are beginning to infect legal proceedings, and others are taking notice. As noted by reporter Jared Perlo, "[w]ith the rise of powerful AI tools, AI-generated content is increasingly finding its way into courts, and some judges are worried that hyperrealistic fake evidence will soon flood their

5. For a discussion on the difference between acknowledged and unacknowledged AI evidence, see Part III, *infra*.

6. *See* Lars Daniel, *Deepfakes Are Entering U.S. Courtrooms—Judges Say They're 'Not Ready*', FORBES (Dec. 8, 2025), https://forbes.com/sites/larsdaniel/2025/12/08/deepfakes-are-entering-us-courtrooms-judges-say-theyre-not-ready/ [https://perma.cc/27Z7-LZ58].

7. No. 23CV028772, 2025 WL 2613764 (Cal. Super. Ct. Sep. 9, 2025).

8. *Id.* at *3–5.

9. *Id.* at *5.

10. *Id.* at *7.

11. *Id.*

courtrooms and threaten their fact-finding mission."[12] Court management consultant Michael Navin writes: "Attorneys are reporting increasing use of AI-faked evidence in divorce and custody cases—doctored images, fabricated texts, even audio clips. In one case, a judge discovered a fake audio recording submitted in a custody dispute and sanctioned the party who submitted it."[13] Another article reports that, "[a]cross industries and jurisdictions, synthetic media created by generative artificial intelligence is disrupting how we identify, interpret and admit evidence. Deepfakes—convincing, falsified digital content—are increasingly being used not only for political misinformation and social manipulation but now, directly in courtrooms."[14]

The issue is not limited to civil cases; criminal cases are also a prime target. A recent criminal case in Florida involving alleged domestic violence depicts an apparent, disturbing use of fake AI evidence. The defendant in the case, Melissa Sims, contended that her boyfriend "scratched and punched himself and used AI to create a video of her attacking him."[15] According to further news reports, Ms. Sims was initially arrested on a domestic violence charge, which she claimed was supported by the fabricated AI-generated video evidence.[16] After her initial arrest, she was released on bond. But while on bond awaiting trial, she was arrested a second time for an alleged violation of her bond conditions—contacting her ex-boyfriend—due to the alleged creation of fake text messages using a text-generating app.[17] Ms. Sims' claim was that the alleged victim, with the aid of AI, fabricated fake text messages to make it appear that she texted the victim, in violation of her bond conditions, even though she had never done so. The bond violation claim was eventually dropped in early 2025, and in December 2025, she was

12. Jared Perlo, *AI-Generated Evidence is Showing up in Court. Judges Say They're Not Ready.*, NBC NEWS (Nov. 18, 2025), https.www.nbcnews.com/tech/tech-news/ai-generated-evidence-deepfake-use-law-judges-object-rcna235976 [https://perma.cc/9RV8-B93V].

13. Michael Navin, *AI-Generated Evidence in the Courts: Are We Ready for What's Coming?*, LINKEDIN (July 14, 2025), https://www.linkedin.com/pulse/ai-generated-evidence-courts-we-ready-whats-coming-michael-navin-uog2c/ [https://perma.cc/U87J-XQXH].

14. Complex Discovery Staff, *Courts at the Crossroads: Confronting AI-Generated Evidence in the Age of Deepfakes*, COMPLEX DISCOVERY (May 29, 2025), https://complexdiscovery.com/courts-at-the-crossroads-confronting-ai-generated-evidence-in-the-age-of-deepfakes/ [https://perma.cc/3PKC-7Z24].

15. Aisling Swift, *Battery or Fake Charges? Woman Seeks New Law to Prevent Arrests Based on Artificial Intelligence*, THE NAPLES PRESS (Nov. 20, 2025), https://www.naplespress.com/local-news/battery-or-fake-charges-woman-seeks-new-law-to-prevent/article_01017a99-3a7a-44d6-8f34-325e99bd2211.html [https://perma.cc/PA3M-A34C].

16. Jensen Bird, *Innocent Woman Thrown into Hellhole Florida Jail After Twisted Ex-Boyfriend Used AI to Send Himself Fake Texts*, DAILY MAIL (Jan. 11, 2026), https://www.dailymail.co.uk/news/article-15449079/florida-woman-jail-ai-text-messages.html [https://perma.cc/T65X-DXLQ].

17. *Id.*

acquitted of the domestic violence battery charge.[18] Yet, serious and important questions remain as to whether evidence fabricated or improperly enhanced by AI could, or will, send an innocent person to prison—an entirely legitimate concern.[19] In fact, it has been reported that in a recent federal criminal case in the Western District of Texas, a confidential informant allegedly provided fake audiovisual evidence to the prosecution which led to the indictment of an innocent person.[20] The fake evidence apparently went undetected by the prosecution and the defense and was not discovered until the confidential informant entered a guilty plea in his case.[21] While the federal prosecutors in that case reportedly dismissed the case against the innocent person after learning of the confidential informant's alleged fabrication of the audiovisual evidence, the fact remains that an innocent person had to suffer an indictment and prosecution before the fraud was uncovered.[22] Frankly, if this happened in a federal criminal case, it can happen in any case. A truly scary thought!

Moreover, parties often seek to introduce evidence through expert witness affidavits and testimony, which can include AI generated materials. In *Kohls v. Ellison*,[23] the court was faced with an expert[24] declaration that allegedly contained AI-generated references to non-existent academic articles.[25] The court did not fault the expert for simply using AI, but warned that "when attorneys and experts abdicate their independent judgment and critical thinking skills in favor of ready-made, AI-generated answers, the quality of our legal profession and the Court's decisional process suffer."[26] I could not put it better. Moreover, the court suggested that "an 'inquiry reasonable under the circumstances,' Fed. R. Civ. P. 11(b), may now require attorneys to ask their witnesses whether they have used AI in drafting their declarations and what they have done

---

18. *Id.*

19. *See* Katie LaGrone, *AI-Generated Fake Evidence Is Landing People in Jail as Courts Struggle with New Technology*, TAMPA BAY 28 ABC (Jan. 21, 2026), https://www.tampabay28.com/news/local-news/i-team-investigates/ai-generated-fake-evidence-is-landing-people-in-jail-as-courts-struggle-with-new-technology [https://perma.cc/PF7K-3VMY].

20. Lars Daniel, *Federal Prosecutors Indicted an Innocent Person on a Deepfake*, FORBES (May 14, 2026), https://www.forbes.com/sites/larsdaniel/2026/05/14/federal-prosecutors-indicted-an-innocent-person-on-a-deepfake/ [https://perma.cc/PB7L-RJPW].

21. *Id.*

22. *Id.* ("Until the federal rules adopt a written threshold, the deepfakes that land hardest in federal court may keep being the ones nobody objected to.").

23. No. 24-CV-3754 (LMP/DLM), 2025 WL 66514 (D. Minn. Jan. 10, 2025).

24. The expert was a Stanford University professor and, ironically, a "credentialed expert on the dangers of AI and misinformation[.]" *Id.* at *3.

25. *Id.* at *1.

26. *Id.* at *4.

to verify any AI-generated content."[27] In *Kohls*, the expert's declaration was ultimately excluded as the fake citations undermined the expert's "competence and credibility."[28] In light of the growing use of AI, opposing parties and judges must be sensitive to the fact that an affidavit, which a party seeks to introduce into evidence, may include false facts, data, and citations. Lawyers should take note of this wise cautionary tale.

While unacknowledged AI evidence like deepfakes receive the most attention, courts and litigants will also have to deal with acknowledged AI evidence as well. This occurs where a party offers AI evidence with full disclosure to the court and opposing parties that the evidence was enhanced or created with AI.[29] Different and novel circumstances arise when AI-enhanced evidence is acknowledged.[30] In *Washington v. Puloka*,[31] the court confronted the issue of whether to admit video exhibits that were openly enhanced by AI. In advance of a criminal trial, a defense expert used AI to enhance seven video exhibits.[32] During trial, the defense intended to submit at least one AI-enhanced Snapchat video recording.[33] After hearing from competing prosecution and defense experts, the court found that the AI-enhanced video failed to meet the *Frye* standard because AI program tools used to enhance the video "have not been peer-reviewed by the forensic video analysis community" and are "not reproducible by that community, and are not accepted generally in that community."[34] Further, the defense did not offer any other case which "examined—let alone approved of—the introduction of AI-enhanced videos in a criminal or civil trial."[35] Lastly, the court concluded that the AI-enhanced video was not admissible under Washington evidence rules.[36] Ultimately, the court excluded the AI-enhanced evidence.[37] These types of admissibility issues surrounding acknowledged AI evidence will arise with regularity in courts and legal

27. *Id.*

28. *Id.* at *5.

29. An excellent law review article discussing the important distinction between acknowledged and unacknowledged evidence is Maura R. Grossman & Hon. Paul W. Grimm, *Judicial Approaches to Acknowledged and Unacknowledged AI-Generated Evidence*, 26 COLUMBIA SCI. & TECH. L. REV. 110 (2025), which appears later in footnote 54. This astute and informative article is highly recommended to the reader.

30. *See infra* Part III, for a discussion on the difference between acknowledged and unacknowledged AI evidence.

31. No. 21-1-04851-2 KNT, Findings of Fact & Conclusions of Law Re: *Frye* Hearing on Admissibility of Videos Enhanced by Artificial Intelligence (Wash. Super. Ct. Mar. 29, 2024).

32. *Id.* at 2 ¶ 2.

33. *Id.* 2 ¶ 3.

34. *Id.* at 5 ¶ 10.

35. *Id.* at 5 ¶ 11.

36. *Id.* at 6 ¶¶ 13–15.

37. *Id.* at 6 ¶ 18.

proceedings. Lawyers, litigants, and judges must be ready, willing, and able to effectively deal with these novel issues.

All of this leads to the inescapable conclusion that lawyers, litigants, and judges must be on alert for the possible presentation of unacknowledged and fabricated AI evidence in court and legal proceedings. Although we are not aware of a deluge of such cases in our courts at this time, the growing popularity and ease of use of AI programs, combined with the high stakes of criminal and civil litigation, strongly suggest that a wave of such fabricated evidence is soon coming to a court near you.

II. NOT ALL AI-GENERATED OR ENHANCED EVIDENCE IS BAD—THERE IS A PLACE FOR ACKNOWLEDGED AI-GENERATED OR PROPERLY ENHANCED EVIDENCE IN LEGAL PROCEEDINGS

Before proceeding further, it is important to understand that there will undoubtedly be several legitimate uses of AI evidence or documents in court proceedings. While this new technology brings many serious potential dangers, as discussed in the preceding section and elsewhere in this Article, it also creates an opportunity for the creative, legitimate, and ethical use of AI evidence. AI creates the opportunity to improve justice and assist the fact-finding process. In fact, prominent commentators in the AI space have raised the question as to whether attorneys failing to use AI could be committing legal malpractice.[38] For example, it was reported that U.S. District Judge Jesse Furman, during a New York State Bar Association panel discussion on artificial intelligence, stated that he "heard somebody say employers are risking malpractice by relying too much on AI." Judge Furman reportedly continued, "I think there may come a point where it's the opposite—where you're committing malpractice if you don't incorporate AI into your practice."[39] Claims for awards of attorneys' fees could certainly be opposed by arguments that the attorneys overbilled attorney time for matters that could have been done in far less time by AI applications. According to Judge Furman, "fees for 'thousands of hours' could be contested by someone arguing they are 'unreasonable, because all these tasks that were done by a lawyer could have been done in 30 seconds by an AI tool.'"[40]

38. Mike Vilensky, *NY Federal Judge Questions if Avoiding AI Could Be Malpractice*, BLOOMBERG LAW (Jan. 13, 2026), https://news.bloomberglaw.com/new-york-brief/ny-federal-judge-questions-if-avoiding-ai-could-be-malpractice [https://perma.cc/8DRZ-CLZQ].

39. *Id.*

40. *Id.*

A state court judge has also reflected that AI can assist courts in boosting access to justice.[41] According to Judge Samuel A. Thumma, "I look forward to seeing how courts use generative AI to enhance access to justice in the future, and am optimistic in those endeavors."[42] Certainly, if used properly, AI has the potential to assist *pro se* parties who are unable to afford legal representation. But the issue of *pro se* parties using AI during litigation is a double-edged sword. On one hand, the use of AI by *pro se* litigants, if used properly and if vetted and verified, could potentially assist such parties in presenting their complaints and arguments in a clear and concise manner. But the use of AI by *pro se* parties who do not check and verify their AI output, as unfortunately is often the case, will lead to serious problems, including the submission of false or inaccurate law, citations and statutes.

As an initial matter, of course, there must be clear and effective guardrails and rules that permit courts to competently and properly admit or exclude AI evidence. Looking at the issue globally, the first prerequisite to the proffer, use, or admission of any such evidence is that it must be acknowledged as AI. That is, no surreptitious or undisclosed use. Second, there must be a pretrial notice requirement. The party offering evidence created or enhanced by AI must promptly and fully notify all opposing parties and the court of the party's intent to introduce such evidence. Advance pretrial notice, at an early stage of the litigation, will allow opposing parties and counsel to properly investigate the evidence's origin, authenticity, provenance, reliability, and other issues of concern.

Moreover, as discussed in Part V, discovery must be allowed into the processes and programs that created or enhanced the AI evidence. It is important that parties be permitted to fully investigate and discover the method, process, provenance, authentication, and other factors that created the AI evidence. Early notice and full discovery will permit timely challenges to any proffered AI or GAI evidence so that the court can rule in a careful and reasoned manner. As discussed in Part V, certain rules of evidence and procedure will need to be amended, revised, or created to specifically address proffered AI evidence.

So then, what potentially legitimate purposes might there be for AI evidence in our courts? A few come to mind.

First, in a non-evidentiary capacity, parties could use AI to create illustrative aids under Federal Rule of Evidence 107. Although not admissible as evidence and not sent back to the jury room during deliberations, illustrative aids can be used to help jurors understand

41. Judge Samuel A. Thumma, *Judges on AI: How Courts Can Boost Access to Justice*, LAW360 (Jan. 12, 2026), https://www.law360.com/articles/2428128/judges-on-ai-how-courts-can-boost-access-to-justice [https://perma.cc/9CT8-RBMA].

42. *Id.*

arguments or evidence. Under Rule 107(a), "[t]he court may allow a party to present an illustrative aid to help the trier of fact understand the evidence or argument if the aid's utility in assisting comprehension is not substantially outweighed by the danger of unfair prejudice, confusing the issues, misleading the jury, undue delay, or wasting time." According to some commentators, "Rule 107 should allow for AI-created trial graphics and other illustrative aids to be used in courtroom presentations. But attorneys should be familiar with the limitations of the content and use of such graphics and aids, particularly as courts continue to grapple with the intersection of new Rule 107 and AI-generated graphics."[43] There most certainly is a place for AI-created illustrative aids in court proceedings, so long as they comply with Rule 107. The nature and types of such charts remain to be seen. Of course, judges will have to determine whether to allow such potentially legitimate AI creations.

Second, in an evidentiary capacity, the myriad uses of AI are only limited by a lawyer's or party's imagination, subject to applicable rules and court decisions. As discussed in a well-written and instructive law review article, AI may be used to discover, aggregate, and summarize evidence and data, create charts and timelines, produce and display forensic evidence, produce forensic sketches, create or enhance evidence such as photos, videos, and audios, and create simulations.[44] Potential attempted uses of AI evidence are, quite frankly, limitless. Lawyers and parties will undoubtedly attempt to use AI in the courtroom for a myriad of purposes. Predicting all the ways in which they will attempt to harness this technology for courtroom use is simply impossible.

Additionally, in a recent Federal Judicial Center publication, it was noted that litigants might seek to introduce "AI-derived evidence in a variety of tort law contexts, both as proof of probability and causation and as the underlying source of tort, as in cases alleging medical malpractice and involving AI-operated vehicle accidents."[45] That same article discussed the potential use or analysis of AI-derived evidence in First Amendment cases, Fourth Amendment cases, and Fifth and Fourteenth Amendment cases.[46] As noted in the publication, issues which will arise include, but are not limited to, the effect of surveillance on First

43. Diana Sterk & Sarah Prather, *AI Enters the Courtroom with New Rule Governing Illustrative Aids*, BLOOMBERG LAW (July 23, 2025), https://news.bloomberglaw.com/us-law-week/ai-enters-the-courtroom-with-new-rule-governing-illustrative-aids [https://perma.cc/22Q7-46SS].

44. Neal Feigenson & Brian Carney, *Generative AI as Courtroom Evidence: A Practical Guide*, 52 MITCHELL HAMLINE L. REV. 1, 28–65 (2025).

45. James E. Baker, Laurie N. Hobart & Matthew Mittelsteadt, *An Introduction to Artificial Intelligence for Federal Judges*, FED. JUD. CTR., at 59 (Feb. 13, 2023), https://www.fjc.gov/content/375968/introduction-artificial-intelligence-federal-judges [https://perma.cc/2G8E-KNZ6].

46. *Id.* at 60–81.

Amendment freedoms, law enforcement use of AI-enabled technology including the use of such technology to obtain search warrants and arrest warrants, liability for AI products, AI empowered contracts, criminal risk assessments, government watch lists, and intellectual property rights to AI.[47] The impact will affect not only civil litigation but also criminal litigation. Ultimately, judges will be called upon to decide the admissibility of such AI evidence in court proceedings—no easy task.

Recent examples of acknowledged AI use and attempted use in court proceedings are informative. To illustrate, in a recent Arizona criminal case, a defendant was convicted of the road rage killing of Christopher Pelkey. During the sentencing hearing, although Mr. Pelkey was deceased, an AI-generated video of him was offered as a victim impact statement. According to news reports, the victim's family "used voice recordings, videos and pictures of Mr. Pelkey . . . to recreate him in a video using AI."[48] The AI-generated video was admitted at the sentencing hearing over defense objection. The AI video appears to have had an impact on the sentencing judge, as he was quoted in the press as saying: "I loved that AI, thank you for that. As angry as you are, as justifiably angry as the family is, I heard the forgiveness . . . I feel that that was genuine."[49] Per a legal commentator, "[t]he video gave the impression that Pelkey himself had authored and delivered the statement. The defense attorney reportedly questioned whether the court placed undue reliance on the deepfake, and an appeal was filed citing potential judicial bias introduced by the emotionally resonant (but synthetic) testimony."[50] The defendant was sentenced to ten years and six months in prison on the manslaughter charge, and an appeal is pending.[51] Although the AI-generated video was admitted at the sentencing hearing in this particular criminal case, the issue is certainly controversial, and other courts may take a completely different view regarding the admissibility of such evidence.

AI animation or virtual reality in our courtrooms and legal proceedings is also already here, and more fulsome efforts to use it are likely on the horizon. For example, in the pending Florida criminal case, *State v. Rodriguez Albisu*,[52] the court was confronted with the issue of

47. *Id.* at 59–81.

48. Ana Faguy & Lily Jamali, *Arizona Man Shot Dead in Road Rage 'Returns' to Address his Killer*, BBC NEWS (May 7, 2025), https://www.bbc.com/news/articles/cq808px90wxo [https://perma.cc/6U6A-B8EX].

49. *Id.*

50. Joseline Jean-Louis Hardrick, *Deeply Moved by a Deepfake: AI's Expanding Role in the Courtroom*, 72 THE FED. LAW. 1, 47 (Fall 2025).

51. State v. Horcasitas, No. CR-2021-142720-001 (Maricopa Cnty. Super. Ct.), notice of appeal filed, No. 1 CA-CR 25-0191 (Ariz. Ct. App. May 7, 2025).

52. State v. Miguel Rodriguez Albisu, No. 23-002405-CF10A (Fla. Cir. Ct., Broward Cnty.).

whether acknowledged forensic animation and virtual reality evidence should be admitted in support of the defendant's self-defense "stand your ground" claim. Facing nine counts of aggravated assault with a deadly weapon, the defendant argued that he brandished the firearm in self-defense. In an effort to support his defense, the defendant filed a Motion to Allow Forensic Animation, which included a demonstrative aid at the Stand Your Ground hearing, and a forensic computer animation demonstrative exhibit at trial, to illustrate the material witnesses' testimony and expert opinion "as to psychological and physiological responses at particular relevant moments in time, and of how the incident occurred."[53] At the pretrial Stand Your Ground hearing, the trial court judge "donned Oculus headsets to watch a 3D recreation" of the events leading up to the time that the Defendant pulled out a gun.[54] While the court allowed the demonstrative aid at the pretrial hearing, it subsequently entered its Order Denying Defendant's Motion to Allow Forensic Examination.[55] Thus, the court ruled that it could not be used at trial. The court also denied the Defendant's Stand Your Ground Motion to Dismiss,[56] and the case is still pending trial. Potential appeals from this case could help to define the use of AI animation and virtual reality in Florida state courts.

As the above discussion shows, there will be permissible uses of AI-generated evidence in the courtroom. The question becomes how well attorneys and litigants will adapt, utilize, and confront such evidence, and how courts will determine the admissibility issues surrounding such evidence.

III. THE ROLE OF LAWYERS AND LITIGANTS IN GUARDING AGAINST DEEPFAKES AND FABRICATED AI EVIDENCE IN LEGAL PROCEEDINGS

Whether for better or worse, AI evidence, enhancements, information, and illustrative aids will become part and parcel of civil and criminal litigation over the coming years, and lawyers and litigants must understand their new role.

An insightful and enlightening law review article discusses the important distinction between acknowledged and unacknowledged AI

53. *Id.* at filing # 205261466, Defendant's Motion to Allow Forensic Animation ¶ 2 (Fla. Cir. Ct., Broward Cnty. Aug. 21, 2024).

54. Taylor Justice, *A New View: The Use and Consequences of Virtual Reality Evidence*, OHIO STATE J. OF CRIMI. L. ONLINE ARTICLES, https://moritzlaw.osu.edu/osjcl-ohio-state-journal-criminal-law-online [https://perma.cc/Y99Y-WMJR].

55. State v. Rodriguez Albisu, *supra*, Order Denying Defendant's Motion to Allow Forensic Animation (Apr. 29, 2025).

56. State v. Rodriguez Albisu, *supra*, Order Denying Defendant's Urgent Motion for Reconsideration on Denial of Defendant's Stand Your Ground Motion to Dismiss and Expedited Ruling in Chambers (May 8, 2025).

evidence.[57] Acknowledged AI evidence refers to situations where all the parties agree that the evidence or content was created, enhanced, or processed using AI.[58] Acknowledged AI evidence can include charts, timelines, enhanced audios or videos, forensic animations, virtual reality, and numerous other types of evidence. On the other hand, unacknowledged AI evidence involves efforts by attorneys and parties to surreptitiously (or perhaps unknowingly) introduce content which has been created, enhanced, or manipulated by AI without any disclosure to opposing parties or the court.[59] This evidence can include deepfakes, fabricated or altered audio or video recordings, phony documents, altered photographs, and numerous other types of AI evidence.

Lawyers and litigants must become sufficiently knowledgeable to both offer or oppose acknowledged AI evidence, while simultaneously being able to detect and challenge unacknowledged AI evidence. AI evidence, both acknowledged and unacknowledged, will arise in all types of local, state, and federal civil and criminal cases. As one attorney has observed, "the unique characteristics of AI demand a more sophisticated and cautious legal approach . . . [and] the need for legal professionals to be not just legal experts, but also technologically savvy, or at least capable of effectively collaborating with technical experts."[60] A good piece of advice is that "[l]awyers must double-check every citation, scrutinize every piece of digital evidence, and adopt a healthy skepticism about any content that appears too perfect to be real."[61]

Lawyers must attend training, CLEs, and other educational programs that will properly prepare them for dealing with AI evidence in legal proceedings. Technical proficiency, or at least access to and consultation with an AI expert with the necessary technological proficiency, is now required for competent counsel.

A recent California state court appellate case raises the issue as to whether lawyers have an ethical or legal responsibility to detect and report fabricated AI pleadings. In *Noland v. Land of the Free, L.P.*,[62] the plaintiffs' counsel allegedly used AI in preparing a legal brief, but unfortunately for counsel, most of the case citations were inaccurate or fabrications. The appellate court found these hallucinated case citations on its own without notification from opposing counsel. Subsequently, when opposing counsel sought an award of legal fees, the request was

57. Maura R. Grossman & Hon. Paul W. Grimm, *Judicial Approaches to Acknowledged and Unacknowledged AI-Generated Evidence*, 26 COLUMBIA SCI. & TECH. L. REV. 110 (2025).

58. *Id.* at 121 (finding that the issue before the court in these scenarios is the *accuracy* of the AI-generated evidence or the AI tool used to create the evidence).

59. *Id.* (explaining that the issue in these cases is the evidence's *authenticity*).

60. Briana C. Breault, *Artificial Intelligence and Fraud: A Bankruptcy Practitioner's Perspective*, 72 THE FED. LAW. 16, 17 (Fall 2025).

61. Hardrick, *supra* note 50, at 47.

62. Noland v. Land of the Free, L.P., 336 Cal. Rptr. 3d 897 (Cal. Ct. App. 2025).

denied due to opposing counsel's apparent failure to notice the hallucinations and report them to the court. Opposing counsel did not alert the court to the fabricated citations, "and appeared to have become aware of the issue only when the court issued its order to show cause" to the plaintiffs' counsel.[63] This asserted lack of diligence led to the court's denial of the opposing party's request for attorneys' fees. In light of this case, as stated by one legal commentator:

> [l]awyers must adopt a higher level of skepticism and diligence in response to growing dependence on AI. Recognizing proactive detection as a critical risk management and litigation recovery strategy is essential, as mandated by the Rules of Professional Conduct and supported by developing case law. Proactive detection goes beyond a mere courtesy to the court; it is fundamental to guaranteeing accountability and safeguarding client interests from an opponent's ethical or technological failures.[64]

While the *Noland* case arose in the context of hallucinated case citations in an appellate brief, the lesson learned is applicable to a lawyer's duty when it comes to faulty or fraudulent AI evidence and deepfakes. Lawyers must be sufficiently knowledgeable of the uses and misuses of AI if they are to fulfill their ethical and legal duties to their clients and to the court.

IV. THE ROLE OF THE JUDGE IN SERVING AS A CAREFUL GATEKEEPER WHEN IT COMES TO THE ATTEMPTED INTRODUCTION OF AI-GENERATED OR ENHANCED EVIDENCE IN LEGAL PROCEEDINGS

Although AI use in court is still in a relatively early stage, judges will continue to be confronted, on an escalating basis, with efforts by attorneys and litigants to introduce both acknowledged and unacknowledged AI evidence in court and legal proceedings. Judges must be prepared to deal with these efforts. The first key is education of the judiciary at both the state and federal levels. Seminars and training programs geared exclusively to teaching judges how to deal with AI evidence in the courtroom are paramount. While lawyers are slowly adapting and learning how to effectively use AI in the courtroom, judges must be ready to rule and intelligently distinguish the probative and admissible AI evidence from the AI junk and deepfakes.

---

63. *Id.* at 915.

64. Joe Stephens, *Detecting AI Misconduct by Opposing Counsel is a Lawyer's Duty*, BLOOMBERG LAW (Nov. 18, 2025), https://news.bloomberglaw.com/legal-exchange-insights-and-commentary/detecting-ai-misconduct-by-opposing-counsel-is-a-lawyers-duty [https://perma.cc/7S5Q-TNFH].

As stated by the Federal Judicial Center, "[j]udges must understand how AI works, its applications, its implications for the fact-finding process, and its risks."[65] This is no easy task, as AI is complicated and confusing to many. Training judges is important, and it must emphasize the potential use and misuse of AI evidence.

When AI meets the courtroom, judges must be ready to meet the rigors of the roles they will be called upon to play. According to the Federal Judicial Center, judges will play at least four roles when it comes to AI in the courtroom.[66]

> First, they will serve as evidentiary gatekeepers, applying the Federal Rules of Evidence (or state equivalents) to proffers of testimonial and documentary evidence, including and perhaps especially Rules 401, 402, and 403. Second, judges will serve as guardians of the law, specifically values embedded in the Bill of Rights as well as statutes and rules of procedure and evidence. Third, judges may serve as potential AI consumers who need to decide whether to receive or rely on AI-generated outputs to inform bail, probation, and sentencing decisions. Fourth, judges will serve as communicators, translating the sometimes-complex inputs behind AI into plain-language instructions for jurors and case law precedent for lawyers.[67]

That is a lot to put on the plate of a busy trial judge.

Arguably, the most prominent and important role that judges will play is that of the careful and cautious gatekeeper, ensuring that only properly authenticated and accurate evidence is presented in the courtroom. With the rise of AI-generated audio files, videos, documents, and deepfakes, this is no easy task. As discussed in the next section, the rules must be revised to provide judges with the necessary tools to perform this difficult task.

V. THE FEDERAL RULES OF EVIDENCE AND CIVIL AND CRIMINAL PROCEDURE NEED TO BE AMENDED TO DEAL WITH ACKNOWLEDGED AND UNACKNOWLEDGED AI EVIDENCE IN LEGAL PROCEEDINGS

The current Federal Rules of Evidence, Federal Rules of Civil Procedure, and Federal Rules of Criminal Procedure need to be amended to provide judges with the proper guidance for making the difficult admissibility decisions surrounding AI-generated or enhanced evidence. Current rules are inadequate. The current rules do not provide judges the necessary specific guidance to effectively deal with proffers or attempted

65. Baker, Hobart & Mittelsteadt, *supra* note 45, at 6.
66. *Id.* at 22.
67. *Id.*

admission of AI evidence in legal proceedings. Current rules must be amended, or new rules must be enacted, to meet the challenge of AI evidence in our courtrooms and legal proceedings.

First, there must be a rule that specifically mandates early pretrial notice from any party or attorney who intends to introduce AI-generated or AI-enhanced evidence. No such notice requirement directed to AI evidence of AI-enhanced evidence currently exists in our rules. Such a requirement is necessary. Currently, an attorney or party seeking to rely upon an audio recording in litigation could only disclose the audio recording without disclosing that it was created or enhanced with AI. This is a serious problem. Therefore, new or amended rules of civil and criminal procedure are required, mandating that all parties and counsel file a prompt pretrial "Notice of Intent to Use" AI-generated or enhanced evidence in a court proceeding. This notice should be required to be filed early in the litigation so that discovery and any necessary motion practice can be directed to the AI-generated or enhanced evidence.

In the civil context, the notice should be filed by all parties contemporaneously with, or as part of, the Initial Disclosures mandated by Federal Rule of Civil Procedure 26(a)(1). Rule 26(a)(1) should be amended to require parties and counsel to explicitly disclose any evidence or materials they intend to rely upon that were created or enhanced by AI. There should also be a continuing obligation to disclose AI evidence if it arises after the initial disclosures are served. In the criminal context, the notice should be required to be filed contemporaneously with, or as part of, the Federal Rule of Criminal Procedure Rule 16 discovery responses of the prosecution and defense. There should also be a continuing obligation to disclose AI evidence in the criminal context.

The Notice, whether filed in a civil or criminal case, should specifically state the type of AI evidence at issue, the name of the person or persons who created or enhanced the AI evidence, a short statement regarding the authenticity of the AI evidence, and the name of the witness or witnesses through whom the party will seek introduction of the AI evidence. Failure to comply with the notice requirement should result in exclusion of any AI-generated or enhanced evidence sought to be admitted by the offending party.

Second, the Federal Rules of Evidence must be amended to allow judges to effectively deal with proffers of AI-generated and enhanced evidence, as well as the challenges to the admission of such evidence. As it now stands, there is no rule that directly and clearly explains how the court should address these issues. When faced with challenges to such evidence, judges have no clear guidance on where the guardrails should be placed. How are judges supposed to determine reliability, authenticity, provenance, and other issues relating to evidence created or enhanced by AI? Which party has the burden of proof, and what precisely is that

burden of proof? Does the proffering party have a burden to show that the AI evidence is reliable and authentic? Does the party opposing admission have the burden to show otherwise, or is there a shifting burden?

One pending effort to regulate AI or machine-generated evidence is that from the Advisory Committee on Evidence Rules to the Committee on Rules and Procedure of the Judicial Conference of the United States. This came about due to concerns about "whether the existing Evidence Rules are sufficient to assure that evidence created by artificial intelligence . . . will be properly regulated for reliability and authenticity."[68] The proposal was previously released for public comment as proposed Federal Rule of Evidence 707, which states as follows:

> When machine-generated evidence is offered without an expert witness and would be subject to Rule 702 if testified to by a witness, the court may admit the evidence only if it satisfies the requirements of Rule 702(a)–(d). This rule does not apply to the output of simple scientific instruments.

The above proposal was published for public comment because "the Committee has determined that a new rule of evidence may be appropriate to regulate the admissibility of machine evidence that is introduced without the testimony of any expert."[69] The proposed rule is intended to apply to only those situations where AI evidence or machine learning is sought to be introduced without the testimony of any expert.[70] The Committee describes some examples of evidence intended to be covered by new Rule 707 as including "machine output analyzing stock trading patterns to establish causation; analysis of digital data to determine whether two works are substantially similar in copyright litigation; and machine learning that assesses the complexity of software programs to determine the likelihood that code was misappropriated."[71] The Committee suggests that "it is possible that the machine output may be offered through a lay witness, or directly with a certification of authenticity under Rule 902(13)," and that in such a situation "a showing of reliability must be made akin to that required under Rule 702."[72]

The fundamental problem with the above proposal is that it would authorize, and more importantly, encourage, attorneys and litigants to seek to introduce AI evidence without any witness whatsoever, expert or otherwise. This could be seen as an effort to sidestep the expert disclosure

---

68. *Preliminary Draft, Proposed Amendments to the Federal Rules of Appellate, Bankruptcy, Civil, and Criminal Procedure, and the Federal Rules of Evidence* (Aug. 2025).

69. *Id.*

70. *Id.*

71. *Id.*

72. *Id.*

rules or to preclude cross-examination—a dangerous road to go down. This is not prudent or wise.

How, under this proposed rule, does opposing counsel cross-examine AI evidence without a testifying witness? When it comes to AI evidence, the adage of "garbage in, garbage out" is applicable. There is simply too much opportunity for abuse to permit attorneys or litigants to introduce AI evidence or machine learning without an appropriate testifying witness. There is no way to cross-examine AI evidence that is introduced without an appropriate testifying witness, and this fact alone should cause great concern over the potential negative and unfairly prejudicial impact of this proposed rule on our legal system. Jurors may place too much emphasis or weight on AI evidence admitted without the rigors of cross-examination of a testifying witness. Ultimately, the potential for misuse is great. I believe that proposed Rule 707, while well-intentioned, is misguided and should not be enacted in its present proposed form.

Rather, I suggest a clear and concise brightline rule that prohibits the introduction of AI evidence without an appropriate testifying witness, unless all the parties stipulate and agree to its admission without a testifying witness. Whether the testifying witness will need to be an expert witness, or merely a lay witness with sufficient knowledge, will depend upon the facts of each case. In such a scenario, the existing Rules 701–705 will govern. And with a testifying witness, discovery, pretrial rulings, and vigorous cross-examination will ferret out the wheat from the chaff, or the legitimate AI evidence from the deepfakes.

Fundamentally, our legal system is headed in the wrong direction with a proposed rule that encourages and permits AI evidence or machine learning without any witness whatsoever. The downsides to proposed Rule 707, including the possibilities of abuse and misuse, are endless, while the upsides are marginal. Likely recognizing this, the Committee also stated that the public should not interpret the Committee's recommendation for Rule 707 to be open for public comment as an endorsement of the proposed rule.[73] The cautionary comment speaks volumes.

There is also a clear need to revise Federal Rule of Evidence 901, "Authenticating or Identifying Evidence," to specifically deal with AI evidence and deepfakes. The current Rule 901 merely requires "evidence sufficient to support a finding that the item is what the proponent claims it is," a very flexible standard that has worked well up until the advent of AI and deepfakes. Rule 901 is wholly unhelpful to judges who must now deal with AI, which may look and appear real but may nonetheless be a clever deepfake. Rule 901 was enacted at a time when AI and deepfakes were not even a thought in the drafters' minds. Current Rule 901 is simply

73. *Id.*

not up to the task of providing judges with the guidance necessary to determine whether to admit or exclude AI evidence. Rule 901 also risks the admission of deepfake evidence, resulting in potential injustice, on too low of an evidentiary standard. The issue of authentication of a handwritten letter, for example, is much simpler than the issue of authentication of AI evidence. As noted by one legal commentator in an astute opinion piece:

> Rule 901 was conceived in an era where forgery was typically crude, detectable, and difficult to scale. Its illustrative examples—testimony based on personal knowledge, distinctive characteristics, chain of custody—presume authenticity can be assessed through human perception and circumstantial context.
>
> Generative AI dismantles this logic. Modern models produce videos of individuals saying things they never uttered, audio clips indistinguishable from real recordings, and documents that mimic unique writing styles—often complete with consistent metadata and artifacts that evade casual scrutiny. In this context, a witness' assurance that evidence "looks real" offers scant probative value.[74]

AI evidence is not your grandparents' evidence. It is different in type, kind, and detectability. AI evidence can look very real but be fake, and deepfakes can be hard for most people to detect.[75]

Current Rule 901(b)(9), dealing with evidence about a process or system, is inadequate as the burden on the proponent—"evidence sufficient to support a finding that the item is what the proponent claims it is"—is too low to guard against deepfakes and other types of fraudulent or manipulated AI evidence. The burden shifting suggestions, described below and elsewhere in this Article, should be written into Rule 901, or into an entirely new rule, to address attempted admission of AI evidence.

Thus, what would an amended Rule 901, specifically addressing AI, look like in combination with other rules?

First, as discussed previously, there needs to be an early pretrial notice requirement in either Rule 901 or in another appropriate rule of procedure in both civil and criminal cases. No sandbagging or unjustified last-minute efforts to introduce AI-generated or enhanced evidence should be permitted. An early and fulsome pretrial notice requirement is fundamental.

---

74. Edward T. Kang, *Authenticity Under Pressure: Rethinking Rule 901 in the Age of AI*, LAW.COM (Jan. 7, 2026), https://www.law.com/thelegalintelligencer/2026/01/07/authenticity-under-pressure-rethinking-rule-901-in-the-age-of-ai/ [https://perma.cc/2JQF-TKFT].

75. Nils C. Köbis, Barbora Doležalová, & Ivan Soraperra, *Fooled Twice: People Cannot Detect Deepfakes but Think They Can*, 24(11) ISCIENCE (2021).

Second, discovery into, *inter alia*, authenticity and reliability—including what has aptly been described as "Provenance: Who created the evidence and using what tools/models? Process: What data trained the system? What safeguards were in place? Verification: What independent, technical verification exists (metadata, hash values, audit logs)?"[76]—must be permitted. Without such fulsome discovery into proffered AI evidence, attorneys and litigants opposing the introduction of deepfake evidence or improper AI evidence will face an unfair, uphill, and herculean task in seeking to prevent its admission.

Third, the respective burdens of proof must be clearly stated and delineated in the rules. I would respectfully suggest a burden-shifting structure, that is, after being provided with timely notice, an attorney or litigant opposing the introduction of AI evidence in a legal proceeding must first meet a burden of proof, akin to either reasonable suspicion or probable cause, that the AI evidence is a deepfake, fraudulent, inaccurate, or otherwise inadmissible. This will likely require expert or direct testimony that the proposed AI evidence is fraudulent, manipulated, or a deepfake. Such a burden on the party opposing introduction will hopefully prevent cavalier and frivolous boilerplate challenges to legitimate, authenticated, and accurate AI evidence. That is, a party or counsel opposing the introduction of AI evidence should not be permitted to make blanket, unsupported, meritless, boilerplate, and cavalier assertions challenging that the evidence is fraudulent, manipulated, or a deepfake without some sufficient level of support. And this level of support should be clearly delineated in our rules.

Next, the proponent of the AI evidence, when faced with a sufficient challenge, will have to meet a burden akin to a preponderance of the evidence standard to allow introduction of the AI evidence in a legal proceeding. Such a significant burden would require the proponent of AI evidence to go beyond the relatively low burden of plausibility, currently permitted by Rule 901, and will hopefully prevent the admission of false AI evidence in legal proceedings. Again, this will likely require expert or eyewitness testimony to allow the proponent to meet their burden.

Finally, if the case reaches the trial stage and the AI evidence is admitted, vigorous cross-examination and appropriate jury instructions will allow jurors to properly evaluate such evidence and reach a fair and just verdict. Certainly, cross-examination is key. Cross-examination will allow the opposing party to effectively address the AI evidence during the crucible of trial. And clear and concise jury instructions will allow jurors to follow the law when determining what weight, if any, to give the admitted AI evidence.

---

76. Kang, *supra* note 74.

These suggestions are just a start, but I think they are good ones. We simply cannot continue business as usual when AI evidence reaches our courtrooms and legal proceedings. Those who espouse the attitude that current rules are sufficient to deal with the coming deepfakes and fake evidence created by AI are, in my humble and respectful opinion, wrong. We need new and amended rules to assist attorneys, litigants, and trial judges in dealing with the coming wave of AI evidence.

VI. NEW JURY INSTRUCTIONS ARE NEEDED TO PREVENT A PARTY FROM GAINING A "LIAR'S DIVIDEND" BY FALSELY ASSERTING A "DEEPFAKE DEFENSE" AND CLAIMING LEGITIMATE EVIDENCE IS IMPROPER AI GENERATED OR ENHANCED EVIDENCE

There are two sides to this AI coin. While deepfakes and unauthenticated AI evidence must be dealt with, exposed, and excluded, as discussed above, frivolous and meritless arguments against legitimate AI evidence should not be countenanced by our courts. What do I mean by this? First, there is the so-called "deepfake defense," whereby attorneys or parties cavalierly argue that no video, audio, or technical evidence can be accepted by a judge or jury because one cannot trust any such evidence in this day and age of AI and deepfakes. This is a dangerous assertion when made without any factual support or good faith basis and can undermine the pursuit of justice in civil and criminal cases.

In a civil slip and fall case, for example, a plaintiff may attempt to argue to a jury, without any factual support or good faith basis whatsoever, that the jury simply cannot rely on the video showing the plaintiff faking a fall because the defendant could have possibly easily manipulated or modified the video. In a criminal case, the same could occur where a defendant, cavalierly and without any good faith basis, argues to the jury that the prosecution's legitimate Title III wiretap evidence, or undercover videos, cannot be believed or trusted because they too could be easily manipulated and modified by the government agents. There needs to be more than mere wild speculation for these types of arguments to be made to jurors.

Of course, if there are serious chain of custody issues or other legitimate issues, those certainly are fair game at trial, and the jury can assess the reliability and weight to give the evidence. For example, in *United States v. Whitehead*,[77] the defendant was convicted of wire fraud and other criminal charges, and subsequently filed a post-trial motion for judgment of acquittal and for a new trial, asserting, *inter alia*, that certain recordings of government-introduced phone calls between him and another individual were improperly authenticated as deepfakes. As stated by the trial judge, "[d]efendant challenged the reliability of the recordings

77. 2024 WL 3085019, *9 (S.D.N.Y June 21, 2024).

before the jury – both during a voir dire of Special Agent Loizias before the recordings were admitted, and during Special Agent Loizias's cross-examination regarding his inability to confirm whether the recordings had been altered before being turned over to the FBI or were 'deepfakes.'"[78] In denying the defendant's post-trial motion, the trial judge ruled that the jury "reasonably could choose to accept the reliability of the recordings despite Defendant's arguments at trial."[79] In this case, the trial judge dealt with the issue by allowing the matter to be presented to the jury, seemingly because there were some apparent chain of custody issues.[80]

But what about when there is no legitimate, good faith basis for a litigant to argue to a jury that certain evidence are deepfakes that simply cannot be trusted? In such a circumstance, what can be referred to as the "liar's dividend" comes into play, whereby "parties attempt to exploit public awareness of deepfakes to cast doubt on legitimate evidence."[81] Unscrupulous lawyers or parties could seek to build upon the public's understanding that audios and videos can, in fact, be fraudulently created deepfakes by baselessly arguing that jurors should not accept any audio, video, or other technological evidence admitted in court due to this general skepticism. In this age of conspiracy theories and deepfakes appearing on social media sites, blogs, in the political arena, and on the internet, such an argument could appeal to one or more skeptical jurors, causing an injustice. As noted by one legal commentator, "[j]uror perception of AI-related content is unpredictable. Will they trust AI? Will they be moved by a synthetic victim impact video? Judges must offer jury instructions on how to interpret such materials."[82]

This is a problem that will be increasingly rearing its head in civil and criminal cases over the coming years. Lawyers and parties, with a win-at-all-costs mentality, may aggressively attempt to divert the jury's attention away from legitimate audio, video, or technological evidence by asserting a meritless deepfake defense and attempting to earn a liar's dividend. Opposing parties and judges must be ready to deal with it.

So, what can prevent a party from gaining a liar's dividend? First, parties faced with such a meritless deepfake defense must bring it to the court's attention by appropriate motion at the earliest possible opportunity so that trial judges can effectively and timely rule on the issues presented. It may be that the trial judge, on a motion *in limine* or other pretrial motion, precludes the attorney or litigant from making such

78. *Id.*

79. *Id.*

80. *Id.*

81. Maryam Akhlaghi, *Deepfakes in Courtrooms: How Judges Are Adapting to a New Reality*, CYBERJUSTICE LAB'Y (Mar. 12, 2025), https://www.cyberjustice.ca/en/2025/05/12/deepfakes-in-courtrooms-how-judges-are-adapting-to-a-new-reality/ [https://perma.cc/57PE-RPAD].

82. Hardrick, *supra* note 50, at 48.

a spurious argument so that the issue does not improperly invade the province of the jury. It may also be that when there is a good faith basis for such an argument, such as authentication, chain of custody, or other legitimate issues, the trial judge allows it with an appropriate cautionary or limiting jury instruction. As noted by legal commentators discussing proposed Rule 707, the Judicial Conference's Advisory Committee on Evidence Rules draft committee note potentially recommends that judges "consider providing a limiting instruction that machine-generated evidence is subject to error and that evidence should not be presumed to be reliable—or unreliable—simply because it was produced by a machine."[83] Other limiting or cautionary jury instructions will likely also be needed. Further, there needs to be a discussion as to what standard of proof jurors need to apply when considering the reliability or unreliability of AI or machine generated evidence.

These are all issues that judges will be confronting over the coming years. The task will not be simple or easy. It will take the combined efforts of ethical counsel, learned judges, and discerning jurors to make our way through the haze of AI-generated and enhanced evidence coming to our courtrooms and legal proceedings.

VII. SUGGESTIONS AND PROPOSALS FOR THE FUTURE—10 CORE COMPONENTS

As someone who appreciates brevity in legal submissions, I propose the following ten core components for effectively dealing with the coming wave and onslaught of AI-generated evidence in our legal system:

1. Lawyers must become knowledgeable in dealing with AI by attending seminars, CLEs, and other training programs. They first must become adept at avoiding hallucinations in their filings, uncovering hallucinations in their opponents' filings, and promptly notifying the court when such issues arise. They must engage AI and ESI experts to assist them in this process.

2. Lawyers must become adept at using AI-generated evidence and illustrative aids properly. Attorneys must have the ability to detect false AI-generated evidence and deepfakes, while simultaneously having the knowledge and technical expertise to timely and effectively challenge improper AI-generated evidence or deepfakes. Lawyers should also engage technical AI and ESI experts as necessary.

3. Lawyers, as well as counseled parties and *pro se* litigants, who seek to knowingly, recklessly, or cavalierly introduce deepfakes or other

83. John Siffert, Jillian Berman & Cindy Kuang, *AI Evidence Rule Tweaks Encourage Judicial Guardrails*, LAW360 (Dec. 9, 2025), https://www.law360.com/insurance-authority/articles/2417855/ai-evidence-rule-tweaks-encourage-judicial-guardrails [https://perma.cc/VF7H-Q7EQ].

fraudulent AI-generated evidence in a legal proceeding or courtroom, must be seriously sanctioned and exposed to severe civil and even potential criminal penalties. There can be zero tolerance for the admission, attempted admission, or proffer of such fabricated or deepfake evidence.

4. Judges must be educated on the proper use, and the potential misuse, of AI-generated evidence. Effective judicial training programs must be developed and offered to our judges who will be increasingly faced with AI-generated evidence in local, state, and federal courtrooms across our country. Judges should also self-educate as to the benefits and perils of AI in our legal system.

5. The current Federal Rules of Evidence, Federal Rules of Civil Procedure, and Federal Rules of Criminal Procedure need to be amended and revised to provide attorneys, litigants, and judges with the tools necessary to deal with AI-generated and enhanced evidence, including deepfakes, in our legal proceedings. Serious consideration should be given to the adoption of a brightline rule precluding the admission of AI-generated evidence or AI-enhanced evidence without a testifying witness, absent the parties' agreement to such admissibility. Burdens of proof should be established for proponents and opponents of such evidence. Rule 901 also needs to be amended, or a new rule created, to provide a higher burden to deal with the potential admission of AI evidence.

6. Pretrial notice of AI-generated evidence or AI-enhanced evidence is essential and should be mandated by court rules at the earliest possible stage of litigation. This is a fundamental requirement that permits proper evaluation of the proffered AI evidence.

7. Discovery into AI-generated or enhanced evidence must be allowed to address issues such as authentication, provenance, reliability, bias, process, verification, methodology, and other relevant issues. This discovery must be robust, and the use of AI and ESI experts during the discovery process should be encouraged.

8. Challenges to AI-generated or enhanced evidence should be required to be made pretrial. Motions directed to AI evidence should be required to be made when dispositive motions, such as summary judgment motions or *Daubert* motions, are due under the court's pretrial scheduling order. This will allow the court to timely rule on such challenges and prevent a meritless deepfake defense from being presented to the jury and a party from gaining a liar's dividend.

9. Judges should endeavor to promptly rule on any pretrial challenges to AI-generated or enhanced evidence so that the attorneys and litigants have a clear understanding as to whether such evidence will be permitted at trial. No sandbagging at trial or at evidentiary hearings should be permitted.

10. Jury instructions need to be prepared and given so as to explain to jurors how to evaluate the reliability or unreliability of AI-generated or enhanced evidence. Current jury instructions are inadequate for this purpose. Clear and concise jury instructions should be developed to guide the jury's evaluation of such evidence.

CONCLUSION

Over the years, our courts and judges have been faced with—and have met—the challenges of new types of evidence in our courts and legal proceedings. Various types of evidence, such as fingerprints, bite mark evidence, firearms evidence, voice analysis, DNA, handwriting, blood and serum, and other kinds, have raised challenges regarding admissibility. But this new AI-generated evidence is hard to detect and different in numerous material respects. The challenges facing lawyers, litigants, and judges are similar to those challenges faced in past decades, yet markedly different. We must raise our collective knowledge and expertise to rise to the challenges of this new AI evidence. Only by training, education, and knowledge can we meet these challenges to ensure that our courts dispense justice fairly and impartially in this age of AI evidence.

As lawyers, litigants, and judges confront the growing wave of AI in our legal system, the sage words of renowned Vietnamese Zen Buddhist monk, poet, and scholar, Thich Nhat Hanh, seem appropriate here:

"Smile, breathe and go slowly."

WHO HAS "POSSESSION, CUSTODY, OR CONTROL" OF THE EMPLOYEE'S PERSONAL MOBILE DEVICE? TIME FOR AMENDMENTS TO THE FEDERAL RULES

Judge Xavier Rodriguez[*]

Abstract

As mobile devices have become indispensable tools for professional communication, the distinction between personal and business data has blurred—creating critical challenges for litigators and organizations navigating discovery obligations. This Article explores the evolving jurisprudence surrounding whether and when a corporate defendant has "possession, custody, or control" over data stored on an employee's personal mobile device. Drawing from a recorded podcast conversation[1] and expanding on the legal foundations, I scrutinize the two primary frameworks courts use to analyze control—the "legal right" test and the "practical ability" test—and question their fitness in the context of modern mobile usage patterns.

This Article evaluates competing conceptions of "control" against three core criteria: (1) doctrinal coherence with the text and structure of the Federal Rules; (2) ex ante predictability for litigants and courts; and (3) the ability to reduce socially costly errors—particularly spoliation risk and unnecessary invasions of employee privacy. A test is "adequate," in this sense, if it supplies administrable guidance ex ante, produces reasonably consistent outcomes across jurisdictions, and does not systematically externalize preservation burdens onto individual employees or shield parties who exploit technical arrangements to avoid discovery. Given this measurement, both the legal right and practical ability tests are inadequate: the former is underinclusive and encourages evasion, while the latter is overinclusive, unevenly applied, and insensitive to privacy and power imbalances.

This Article also considers the growing disconnect between technological reality and static discovery rules, emphasizing the need for clarity in preservation and production expectations. Finally, this Article addresses the implications of emerging state statutory and case law developments regarding privacy, the General Data Protection

---

* Judge Rodriguez extends his thanks to Christian Krueger, a third-year law student at the University of Texas School of Law for his invaluable research, proofing and cite checking assistance. My thanks also to Kelly Twigger for prompting the idea of writing on this topic and providing me access to Minerva for research purposes.

1. To hear a discussion with the author on this topic, see Meet and Confer with Kelly Twigger, *Mobile Minutes: Judge Xavier Rodriguez on Possession, Custody, or Control*, SPOTIFY (July 24, 2025), https://open.spotify.com/episode/30ygUsDGeFHRAYgRhIvpu3?si=ff19b619907f4e60 [https://perma.cc/9APB-VRW5].

Regulation,[2] and the Hague Evidence Convention[3] (where an employer's limited control over personal devices may be further complicated by these privacy laws and jurisdictional boundaries). In response, this Article proposes a more functional approach to assessing control and offers practical recommendations for litigators, courts, and policymakers to modernize discovery frameworks in a Bring Your Own Device (BYOD) world.

This Article proposes an agency-anchored conception of "control" as one promising candidate for a more uniform national standard. Under this approach, a rebuttable presumption of control would attach to electronically stored information (ESI) on the personal devices of officers, directors, and individuals with meaningful managerial authority acting within the scope of their agency. In contrast, other employees' devices would ordinarily be treated as outside the responding party's control and only reachable, if at all, through Rule 45.

2. Regulation 2016/679, General Data Protection Regulation (GDPR), 2016 O.J. (L 119) 1.

3. Convention on the Taking of Evidence Abroad in Civil or Commercial Matters, Mar. 18, 1970, 847 U.N.T.S. 231.

I. INTRODUCTION

A. *Mobile Devices Meet 1938 Rules*

Few developments in modern civil litigation have unsettled the discovery obligations imposed by the Federal Rules of Civil Procedure more than the explosion of personal mobile devices. Where smartphones are ubiquitous, always on, and deeply entwined with both personal and professional life, courts increasingly face the question: when must a corporate defendant collect and produce information stored on employees' personally owned devices? The issue implicates a central phrase in the Federal Rules of Civil Procedure—"possession, custody, or control"—that has guided document production obligations since the Rules' inception in 1938.

This question is far from academic. In contemporary workplaces, employees routinely use personal phones to send emails, exchange text messages, access corporate databases, and communicate via messaging platforms. Many corporations encourage or even mandate this integration through BYOD policies. Others tolerate the practice informally. In either event, the blending of personal and professional data on personal mobile devices complicates discovery obligations and raises profound questions about privacy, proportionality, employee morale, and the extent of corporate responsibility for evidence spoliation.

Federal Rule of Civil Procedure 34(a)(1) provides that a party may request certain items in the responding party's possession, custody, or control.[4] The phrase seems straightforward, but when applied to personal

4. FED. R. CIV. P. 34(a) states:

> (a) In General. A party may serve on any other party a request within the scope of Rule 26(b):
>
> (1) to produce and permit the requesting party or its representative to inspect, copy, test, or sample the following items in the responding party's possession,

devices, it quickly becomes elusive. Does "control" mean a corporation must collect text messages from employees' personal phones? Or does it mean the employer has no obligation unless it can legally compel access? Or does it occupy some middle ground, such as when the employer has the "practical ability" to obtain the information?

B. *Stakes for Litigants*

The question of whether an employer has control over ESI on employees' personal mobile devices has become common, especially in the wake of the BYOD movement,[5] thus deserving careful attention.

When the duty to preserve is triggered, litigants must take measures to preserve relevant material or potentially face later sanctions. The application of this duty becomes more complex when relevant data may reside on employees' personal mobile devices, raising a number of practical and legal questions. Does a company take measures to preserve mobile devices owned by an employee? What form should those preservation measures take? How are they enforceable? Should employers discipline employees who refuse to cooperate? How do requesting parties know whether the corporation will assist in the discovery process or later argue that they do not possess or control any data that may be resident in an employee mobile device? If these conversations will not take place until the meet and confer sessions, will relevant and unique data be lost? If the parties agree to certain ESI protocols that contemplate corporation control over employee devices, can a corporation later be sanctioned when employees refuse to cooperate? At what point in these discussions should a requesting party issue Rule 45 subpoenas to individuals? At what point do courts become involved in this issue—before the issuance of any requests for production (RFP)? Or does court involvement take place after responses and objections to the RFPs are served asserting a lack of control over employee devices? If Rule 45 subpoenas are issued to individual

---

> custody, or control:
>
> (A) any designated documents or electronically stored information—including writings, drawings, graphs, charts, photographs, sound recordings, images, and other data or data compilations—stored in any medium from which information can be obtained either directly or, if necessary, after translation by the responding party into a reasonably usable form . . . .

5. *See generally* Garry G. Mathiason & Michael J. McGuire, *The 'Bring Your Own Device' to Work Movement: Engineering Practical Employment and Labor Law Compliance Solutions*, THE LITTLER REP. (2012), https://www.littler.com/sites/default/files/press/pdf/WP_12EE_BringYourOwnDevice_7-23-13.pdf [https://perma.cc/NZE3-G8WU] (discussing BYOD issues).

employees, does the corporation retain counsel for these employees to navigate the process of responding or objecting?

Courts are split and litigants face uncertainty. In some jurisdictions, requesting parties (usually plaintiffs) can compel corporations to collect and produce employee text messages, especially when corporate policies suggest some measure of control.[6] In other jurisdictions, courts shield employers, leaving plaintiffs to pursue non-party subpoenas under Rule 45.[7] The difference is consequential. Rule 34 production requests impose obligations directly on the corporate defendant, whereas Rule 45 subpoenas[8] place the burden on employees themselves to respond, object, or resist production.[9] In addition, while the issue of possession, custody, or control is disputed, the passage of time may cause relevant data on mobile devices to be lost through failures to disable automatic deletion settings in text settings, the overwriting of data, physical damage to the device, or replacement of the device.

For requesting parties (usually plaintiffs), access to data on an employee's mobile device may be critical to proving claims. In employment discrimination suits,[10] wage-and-hour cases,[11] antitrust

---

6. O'Bryan v. U.S. Bank Nat'l Assoc., No. 3:20-CV-00153, 2022 WL 22736591, at *4 (M.D. Tenn. Mar. 10, 2022) (Court finds that the Bank's Employee Mobility Program purports to give U.S. Bank ownership of specific information on its employees' personal devices under certain circumstances, but it was not clear that those circumstances existed.).

7. Control "for purposes of Rule 34 means demonstrating that the party served with the document request 'has the legal right to obtain the documents on demand' from someone else. This is understood to include 'the legal right to command release from the party with actual possession.'" Halabu Holdings, LLC v. Old Nat'l Bancorp, No. 20-10427, 2020 WL 12676263, at *3 (E.D. Mich. June 9, 2020) (internal citations omitted).

8. FED. R. CIV. P. 45(a)(1)(C) provides:

> A command to produce documents, electronically stored information, or tangible things or to permit the inspection of premises may be included in a subpoena commanding attendance at a deposition, hearing, or trial, or may be set out in a separate subpoena. A subpoena may specify the form or forms in which electronically stored information is to be produced.

9. FED. R. CIV. P. 45(d)(2)(B) states:

> Objections. A person commanded to produce documents or tangible things or to permit inspection may serve on the party or attorney designated in the subpoena a written objection to inspecting, copying, testing or sampling any or all of the materials or to inspecting the premises—or to producing electronically stored information in the form or forms requested.

10. Memorandum and Order, Shim-Larkin v. City of N.Y., No. 1:16-CV-6099-AJN-KNF, 2019 WL 5198792, at *1 (S.D.N.Y. Sep. 16, 2019) (Defendant sanctioned for loss of text messages stored on the personal cellular telephone of the defendant's assistant lifeguard coordinator.).

11. Small v. Univ. Med. Ctr., No. 2:13-CV-0298-APG-PAL, 2018 WL 3795238, at *71 (D. Nev. Aug. 9, 2018) (sanctions assessed against employer).

conspiracies,[12] and trade secret disputes,[13] text messages or encrypted chats may contain decisive evidence. From the perspective of plaintiffs and other requesting parties, personal devices may be the only realistic source of evidence of harassment, off-the-clock work, or informal communications that never pass through corporate systems. A standard that presumptively insulates most personal devices from Rule 34 may therefore raise serious access-to-justice concerns in employment, wage-and-hour, and civil rights litigation. For corporate defendants, insisting on identifying, reviewing, and potentially producing data from individual employees' mobile devices raises burdens and risks: alienating employees, potentially invading personal privacy, and creating opportunities for spoliation if employees refuse cooperation.

It is unclear what constitutes possession, custody, or control. Though the phrase "possession, custody, or control" is used in Rules 26, 34, and 45, the Federal Rules do not define the phrase.[14] Consequently, courts have had to reason with the words' meaning. Typically, physical "possession" and "custody" are straightforwardly determined.[15] Much of the real debate focuses on "control," specifically, with different courts adopting different standards.[16] Courts have broadly coalesced around three tests, respectively: legal right, practical ability, and legal right plus notification.[17] In addition to there being multiple tests, there is also variation in the application of each test, resulting in an even messier amalgamation.[18] The consequent inconsistency causes uncertainty among litigants on an important matter. Pursuant to the Federal Rules, a party is

12. *In re* Pork Antitrust Litig., No. 18-CV-1776 (JRT/HB), 2022 WL 972401, at *15 (D. Minn. Mar. 31, 2022).

13. Allergan, Inc. v. Revance Therapeutics, Inc., No. 3:23-CV-00431, 2025 WL 2182324, at *1 (M.D. Tenn. June 20, 2025).

14. FED. R. CIV. P. 26(a)(1)(A)(ii) ("a copy—or a description by category and location—of all documents, electronically stored information, and tangible things that the disclosing party has in its possession, custody, or control and may use to support its claims or defenses"); FED. R. CIV. P. 34(a)(1) ("to produce and permit the requesting party or its representative to inspect, copy, test, or sample . . . items in the responding party's possession, custody, or control"); FED. R. CIV. P. 45(a)(1)(A)(iii) ("produce designated documents, electronically stored information, or tangible things in that person's possession, custody, or control").

15. The Sedona Conference, *Commentary on Rule 34 and Rule 45 "Possession, Custody, or Control"*, 25 SEDONA CONF. J. 1, 14 (2024) [hereinafter *Rule 34 and 45 Commentary*].

16. *See id.* at 17–34 (outlining the different tests and which courts use each of them).

17. *See id.* at 18 (explaining the legal right test as "the legal right to obtain the Documents and ESI," the practical ability test as "the 'practical ability' to [obtain the Documents and ESI]" without a legal right to do so, and the legal right plus notification test as "the legal right to obtain the Documents and ESI" plus the duty to "notify . . . [the] adversary" if "the party does not have the legal right to obtain the Documents and ESI").

18. *See id.* at 28, 34 ("To further complicate matters, even within these general categories there are differences in the ways in which federal courts within the circuits define and apply the standards," leaving "parties and courts with conflicting guidance to consider when making defensible discovery decisions.").

obligated to preserve relevant ESI under penalty of sanction.[19] In addition to raising spoliation concerns, irregularity in the standard for "control" over devices may cause lengthy discovery disputes that delay expeditious lawsuit resolution. Finally, interjurisdictional discrepancies contravene the Federal Rules' interest in uniformity.[20]

There is a need for greater predictability regarding employers' control of employees' personal mobile devices, but none of the established tests are without fault. While The Sedona Conference has advocated for the universal adoption of the legal right test,[21] this approach is underinclusive, leading to spoliation, evasion, and omission. The practical ability test has its own shortcomings, most notably in implementation (what if an employee simply says, "No"?). A test based on the principles of agency law would be superior. Under it, a requesting party would be able to gather data from the responding party's officers, directors, supervisors, and managers. In other words, ESI on the personal mobile devices of high-level agents would presumptively be under the company's control, whether by right or ability. Agents authorized to act on behalf of the company would have to produce relevant information. To avoid overbreadth, all other employees would have to be individually subpoenaed by the requesting party under Rule 45. Such an approach to a uniform national standard merits further consideration.

II. THE HISTORICAL ORIGINS OF RULES 34 AND 45

As Wright and Miller explain, Rule 34 was patterned after the English Judicature Act and state statutes that allowed inspection of documents in a party's "possession or control."[22] The Advisory Committee intended to abolish the restrictive bill in equity practice[23] and create a streamlined

19. FED. R. CIV. P. 37(e).

20. Rhys Dipshan, *Diverging 'Possession, Custody, or Control' Tests Impact E-Discovery Outcomes. But is a Uniform Standard Feasible?*, LEGALTECH NEWS (May 5, 2023), https://monitor.lawnext.com/article/diverging--possession--custody-or-control--tests-impact-e-discovery-outcomes.-but-is-a-uniform-standard-feasible- [https://perma.cc/96SJ-9F5W].

21. *Rule 34 and 45 Commentary*, *supra* note 15, at 63–64.

22. 8B WRIGHT & MILLER'S FEDERAL PRACTICE & PROCEDURE § 2201 (3d ed. 2025).

23. Equity Rule 58, promulgated in 1912, permitted interrogatories and the inspection of documents in the possession or control of the opposing party. The language of "possession or control" thus had deep roots in equity practice, reflecting a recognition that litigants sometimes required access to adversaries' records to prove their cases. But the remedy was limited, hedged with procedural technicalities, and available only in the context of equitable claims. As Sunderland's influential treatise *Discovery Before Trial* (1939) explained, American practice before the Federal Rules of Civil Procedure was a patchwork of state statutes, local rules, and equity precedents. Sunderland catalogued statutes from across the states, most of which allowed some form of pretrial inspection of documents "in the possession or under the control of" an opposing party. The aim was functional—to provide litigants a fair opportunity to access evidence—but the statutes varied widely, and technical objections often frustrated their purpose.

mechanism for production.[24] Rule 45, in turn, provided a means of reaching non-parties, ensuring that litigants could access information beyond their adversaries' files.[25]

Yet the 1938 drafters could not have imagined a world where terabytes of personal and professional data would coexist on a single handheld device, carried in a pocket. Their concern was with paper ledgers, contracts, correspondence, and land records—not encrypted instant messages or ephemeral chat applications. Even so, the doctrinal vocabulary they crafted continues to govern twenty-first century disputes. Complicating the matter, the 1938 drafters (and the Rules Advisory Committee to this day) fail to define the phrase "possession, custody, or control."

If a party has actual physical possession of a document or data, those documents or data must be produced if relevant, proportional, and non-privileged.[26] Likewise, if a party is the custodian of the documents or data, it must comply with Rule 34.[27] Why the framers of the 1938 Rule included the term "control" remains elusive. Cases and secondary sources from that time merely parrot the phrase "possession, custody, or control" without any insight as to what the term "control" should mean.

As noted above, prior to 1938 most state rules of discovery contained a reference to some variation of possession, custody, or control.[28] Alabama, for example, required a responding party "to attach to his answers copies of letter and documents, the originals of which are in his possession or under his custody or control."[29] Florida allowed courts to "require the party notified . . . to produce books and other writings in his possession, power, or custody."[30] Connecticut permitted a party to "file a motion praying for a disclosure of facts or production of papers, books, or documents . . . within the knowledge, possession or power of the

---

Edson R. Sunderland, *Discovery Before Trial Under the New Federal Rules*, 15 TENN. L. REV. 737 (1939).

24. WRIGHT & MILLER, *supra* note 22, § 2201.

25. FED. R. CIV. P. 45; *see also* WRIGHT & MILLER, *supra* note 22, § 2451 Purpose and History of the Rule ("A 'subpoena' is a mandate lawfully issued in the name of the court, usually by the clerk thereof, but under current practice, by attorneys. Its function is to compel the attendance of witnesses and the production of documents so that the court may have access to all of the available information for the determination of controversies before it.").

26. FED. R. CIV. P. 26(b)(1).

27. *See* United Mercantile Agencies v. Silver Fleet Motor Exp., 1 F.R.D. 709, 712 (W.D. Ky. 1941) ("In the usual course of events the defendant would either have actual physical possession of its own records or would be in a position to obtain them from someone who has temporary custody of them with the defendant's permission. Such control, without actual possession, is sufficient.").

28. *See* GEORGE RAGLAND, JR., DISCOVERY BEFORE TRIAL 267–391 (1932) (detailing statutory provisions on discovery in many jurisdictions).

29. *Id.* at 274.

30. *Id.* at 292–93.

adverse party."[31] The courts of England, too, stipulated that a judge may "order the production by any party . . . of such of the documents in his possession or power, relating to any matter in question."[32] As the representative examples discussed above demonstrate, the concept of "possession, custody, or control" preceded the Federal Rules and was already widespread among the state courts.[33]

Still earlier, courts here and abroad recognized an analog of "possession, custody, or control" in discovery. Characterizing the state of the law in 1873, the scholar John Adams, Jr. explained that discovery "must be to the best of the defendant's information and belief," and "the information meant is not only that which he actually possesses," but rather "that also which, either by inspecting his books, or by making inquiries of his solicitors or agents, or of others from whom he has a right of information, is fairly within his reach."[34] Furthermore, a responding party's "mere allegation" that "such parties will not give him the information, or even that they have refused to do so, will not be sufficient to excuse its want."[35] The standard extended beyond physical possession or actual custody, incorporating "whatever means of information he has a right to possess" as "look[ed] upon as being in his possession."[36] While the requested documents had to "be in the defendant's possession or power," it was "sufficient that they [we]re admitted to belong to him, although . . . out of his actual custody."[37] The possession of an "agent, or of any other person whose possession he c[ould] control, [wa]s equivalent to his own."[38] At the outer limit, though, a party was "not bound to

31. *Id.* at 281.

32. *Id.* at 287.

33. The remaining states and the terminology referenced in their statutes are as follows: Arizona (possession or control), California (possession or control), Colorado (possession or control), Delaware (possession or control), Georgia ("within his knowledge"), Idaho (possession or control), Illinois (possession or power), Indiana (possession or power), Kansas (possession or control), Kentucky (control), Maine (possession), Maryland (possession or power), Massachusetts (possession or control), Michigan (possession or control), Minnesota (possession or control), Mississippi (possession or control), Missouri (possession or control), Montana (possession or control), Nebraska (possession or control), Nevada (possession or control), New Jersey (possession or control), New Mexico (possession or control), New York (possession or control), North Carolina (possession or control), North Dakota (possession or control), Ohio (possession or control), Oklahoma (possession or control), Oregon (possession or control), Pennsylvania (possession or power), Rhode Island (possession or control), South Carolina (possession or control), South Dakota (possession or control), Utah (possession or control), Virginia (possession), West Virginia (control), Wisconsin (possession or control), Wyoming (possession or power). *Id.* at 275–391.

34. JOHN ADAMS, THE DOCTRINE OF EQUITY: A COMMENTARY ON THE LAW AS ADMINISTERED BY THE COURT OF CHANCERY 75 (6th Am. ed. 1873).

35. *Id.*

36. *Id.*

37. *Id.* at 79.

38. *Id.*

produce documents for which he merely h[eld] a covenant for production against a third party."[39] Like courts today, centuries-old courts grappled with possession, custody, or control.

In England, too, possession, custody, or control boasts an extended history. As former Harvard Law School Dean Christopher C. Langdell chronicled, discovery itself "was borrowed by the Court of Chancery, directly from the English ecclesiastical courts."[40] In the Ecclesiastical Courts, in matters of probate, parties were required to file an affidavit "specifying all writings of a testamentary nature in his possession, or within his knowledge."[41] Writing in 1897, Dean Langdell defined discovery as "the compulsory disclosure by a litigant of such facts within his knowledge, and of the contents of such documents in his possession."[42] Despite parallels between American and English law, Dean Langdell noted that "discovery is not an original product of English soil" and is "entirely foreign."[43] It may be traced back to Roman law.[44] "Possession, custody, or control," whatever its source, predates the Federal Rules.

Before the enactment of the Federal Rules, federal law—statutory and common—incorporated the premise of "possession, custody, or control." At that time, there were three means by which a party could compel the production of documents.[45] One, a deposition statute,[46] was construed by courts to "include such a power" to "compel the production of papers or documents under the control of such a witness."[47] Another, a subpoena duces tecum statute,[48] allowed parties to "apply for an order directing a person to produce documents supposed to be in his possession."[49] Lastly, Equity Rule 58 enabled courts to direct parties to "effect the inspection or production of documents in the possession of either party."[50] To obtain a court order, the requesting party had to show materiality "together with an admission by his adversary that the documents sought were actually

---

39. *Id.* at 75 n.1.

40. Christopher C. Langdell, *Discovery Under the Judicature Acts, 1873, 1875*, 11 HARV. L. REV. 137, 138 (1897).

41. *Id.* at 147.

42. *Id.* at 137.

43. *Id.* at 138.

44. Robert W. Millar, *The Mechanism of Fact-Discovery: A Study in Comparative Civil Procedure*, 32 ILL. L. REV. 261, 263 (1937).

45. H.S.T., *Notes and Legislation*, 26 VA. L. REV. 769, 797 (1940).

46. 28 U.S.C. § 639 (1934).

47. H.S.T., *supra* note 45, at 797–98.

48. 28 U.S.C. § 647 (1934).

49. H.S.T., *supra* note 45, at 798; *see also* Smith v. Nat'l Bank, 193 F. 255 (C.C.D. Nev. 1910) (allowing a requesting party to make a copy of the responding party's document without the responding party surrendering it).

50. H.S.T., *supra* note 45, at 800 (quoting 28 U.S.C. § 723 (1928)).

under the latter's power or control."[51] The "many technicalities" of Equity Rule 58 "limited [its] utility" as a discovery tool.[52] Against that backdrop, early caselaw interpreted the new Federal Rules as marking a departure from the Equity Rules.[53]

Rule 45 did not originally refer to possession, custody, or control. The first iteration stated, "A subpoena may command the person to whom it is directed to produce the books, papers, or documents designated therein," omitting the qualifier of possession, custody, or control.[54] It was added via the 1970 amendment to the Federal Rules, however. Including possession, custody, or control in this subsequent edition, the Federal Rules clarified that "the scope of discovery through a subpoena is the same as that applicable to Rule 34 and the other discovery rules."[55] Though Rule 45 incorporated possession, custody, or control language later than Rule 34 did, the Advisory Committee underscored that the interpretative exercise should be the same in each case. In any case, courts have had over a half century of possession, custody, or control under Rule 45. The longevity of possession, custody, or control offers useful context, but is no panacea for litigants searching for definitive answers to new-age problems.

III. THE PUSH FOR UNIFORM FEDERAL PROCEDURE

By the 1920s and 1930s, the inadequacy of federal discovery practice had become clear. Under the Conformity Act of 1872, federal courts were required to follow state procedure "as near as may be." [56] This created a bewildering array of systems: as many distinct regimes of federal practice as there were states, and sometimes more, given district-level variations.[57]

Reformers, including Roscoe Pound and Charles E. Clark, argued that federal practice required unification and simplification. The American Bar Association, prodded by critiques of delay and uncertainty, formed

---

51. *Id.* (citing United States v. Interstate Cir., 20 F. Supp. 463 (N.D. Tex. 1937); Dixie Drinking Cup Co. v. Paper Utils. Co., 5 F.2d 322 (E.D.N.Y. 1925); Hartford v. Cleveland Auto Club, 275 Fed. 590, 591 (D. Del. 1921)).

52. H.S.T., *supra* note 45, at 801 (citing Moore v. Backus, 78 F.2d 571 (7th Cir. 1935) (denying discovery request despite the witness being a party's agent)). By citing *Moore* for the proposition that Equity Rule 58 was restrictive, that piece's author suggests that defendants' agents should fall within the purview of discovery.

53. *See, e.g.*, Nichols v. Sanborn Co., 24 F. Supp. 908, 911 (D. Mass. 1938) ("[I]t is better that liberality rather than restriction of interpretation [of the new *Federal Rules*] be the guiding principle. This will avoid the confusion and complexities which have resulted as a result of the diversified interpretations of Equity Rule 58 by the Courts.").

54. FED. R. CIV. P. 45(d) (1938).

55. FED. R. CIV. P. 45 advisory committee's note to 1970 amendment.

56. Judicial Conformity Act of 1872, ch. 255, 17 Stat. 197 (repealed 1938).

57. Elmo Hunter, *One Year of Our Federal Rules*, 5 MO. L. REV. 1, 3–4 (1940).

committees to study reform. When Congress enacted the Rules Enabling Act in 1934,[58] it authorized the Supreme Court to promulgate rules of civil procedure for the federal courts, abolishing the law–equity divide and creating an opportunity for comprehensive discovery reform.[59]

The Advisory Committee debated extensively how to integrate discovery into the new rules. They drew from English practice under the Judicature Acts, American equity, and progressive state codes. The result was a robust discovery system, including Rules 26–37, that placed disclosure at the center of federal litigation.

While Rule 34 governed discovery between parties, Rule 45 governed subpoenas directed to non-parties. Adopted in 1938 alongside Rule 34, Rule 45 authorized subpoenas duces tecum to compel production of documents or things by non-parties.

The history of Rule 45 reveals the balance between fairness and burden. Initially, a subpoena duces tecum could not issue without a court order. The 1948 amendments eliminated this requirement, allowing attorneys to issue subpoenas directly, while preserving judicial authority to quash or modify subpoenas that imposed undue burden. Over time, Rule 45 evolved to include explicit protections for non-parties, including cost-shifting and privilege safeguards.

The relationship between Rule 34 and Rule 45 is crucial to the mobile device conundrum. If an employee's personal phone is deemed outside the employer's possession, custody, or control, then Rule 45 provides the mechanism for accessing its contents.

IV. BEFORE THE MODERN MOBILE-DEVICE ERA

Immediately after the Federal Rules' adoption, some courts treated the new Federal Rules as a marked departure from the previous Equity Rules. In *C.F. Simonin's Sons v. American Can Co.*,[60] decided just one month after the Federal Rules took effect, the court declared that "they afford the plaintiff facilities for discovery . . . which are more liberal than it could obtain in a suit in equity."[61] Courts considering discovery matters struck a similar tone. Two months after Rule 34's arrival, the court in *Gielow v. Warner Bros. Pictures*[62] described taking a liberal approach to discovery.[63] It remarked that Rule 34 was intended to "simplify[] the

58. Rules Enabling Act, ch. 651, 48 Stat. 1064 (1934) (codified as amended at 28 U.S.C. §§ 2071–77).

59. Hunter, *supra* note 57.

60. 24 F. Supp. 765 (E.D. Pa. 1938).

61. *Id.*

62. 26 F. Supp. 425 (S.D.N.Y. 1938).

63. *Id.*

issues," and that it would account for the liberal spirit animating the Federal Rules.[64]

The first courts to encounter the Federal Rules for discovery did not engage in intensive analysis of them. Instead, they were taken as self-explanatory, and courts' decisions relied on conclusory statements without detailed reasoning. In *Orange County Theatres v. Levy*,[65] the court invoked possession, custody, or control but did not scrutinize it.[66] The requesting party sought the Strand Theatre's accounting records, and the court simply posited "the plaintiff is only required to produce the [records] if 'they are in its possession, custody, or control,'" without more.[67] In so doing, the court treated possession, custody, or control as apparent. In *Flynn v. Magraw*,[68] a plaintiff requested various business documents from Thomas Manor, Inc., by way of John McGinty, a director of the corporation.[69] Mr. McGinty countered that he was a director but not an officer and therefore did not "have the custody or possession of any of the . . . [documents] which he ha[d] been called upon to produce for inspection."[70] Accepting that and without evaluating the distinction between a director and officer for the purposes of control, the court denied the request and advised, "[i]t may be a hardship for the plaintiff but an application to examine the known officers, and for the production and inspection of the . . . [records] can be made in Florida."[71] While courts framed possession, custody, or control expansively, they applied the standard formalistically.

Mid-twentieth century courts transitioned toward performing careful individualized inquiry and reaching fact-based decisions. In *Tobin v. WKRZ, Inc.*,[72] the Secretary of Labor sought employment-related documents from the defendant-company.[73] The court identified the sole

64. *Id.*

65. 26 F. Supp. 416 (S.D.N.Y. 1938).

66. *Id.* at 417–18.

67. *Id.*

68. 27 F. Supp. 936 (S.D.N.Y. 1939).

69. *Id.* at 936.

70. *Id.*

71. *Id.*; *see also* H.S.T., *supra* note 45, at 804–05 ("[I]n an action against a corporation an order to produce its books and records will issue only against an officer who has access to these documents."); *cf. In re Harris*, 27 F. Supp. 480, 481–82 (S.D.N.Y. 1939) (denying discovery even though the requested bank records were within the responding party's physical possession because the responding party was found not to have control of documents maintained by a foreign affiliate under the National Banking Act). *But see* H.S.T., *supra* note 45, at 805 ("[T]he defense of non-possession is not open to a person when it appears that he has the ultimate control over, if not the physical possession of, the things sought."); *Bough v. Lee*, 29 F. Supp. 498, 501 (S.D.N.Y. 1939) ("The excuse of non-possession was not sufficient, even if true" where the responding party did not deny it was in possession of relevant documents.).

72. Tobin v. WKRZ, Inc., 12 F.R.D. 200 (W.D. Pa. 1952).

73. *Id.* at 201.

question as whether the requested items were within the responding party's possession, custody, or control.[74] Analyzing the issue, the court observed that "the cases have not indicated precisely what constitutes a sufficient showing."[75] It reasoned, though, that "any showing is sufficient which, on the record, satisfies the Court."[76] Looking at the specific circumstances of the discovery request, the court concluded that records maintained for ordinary business under government regulation would be presumed within possession, custody, or control.[77] In *Biferato v. States Marine Corp. of Delaware*,[78] the responding party sought to resist production of vessel logs and reports because its attorney possessed them, but the court declined to allow "[p]ossession by an attorney or a third party . . . [to] be used as a means of avoiding compliance with a direction . . . . The true test is control and not possession."[79] As courts engaged more directly with possession, custody, or control analysis, the facts and circumstances became significant.[80]

Courts in the latter portion of the twentieth century opened the aperture of possession, custody, or control and imposed more affirmative obligations on responding parties. In *Herbst v. Able*,[81] the plaintiffs requested that some of the defendant's current and former employees procure from the Securities and Exchange Commission (SEC) and make available to the plaintiffs transcripts of testimony offered at private

74. *Id.*

75. *Id.*

76. *Id.*

77. *Id.* at 201–02.

78. 11 F.R.D. 44 (S.D.N.Y. 1951). *But cf.* Hickman v. Taylor, 329 U.S. 495, 504–08 (1947) (denying a request for documents in the possession of an attorney because the original sources were otherwise available to the requesting party and stating "Rule 34 . . . is limited to parties to the proceeding, thereby excluding their counsel or agents."); *see also* H.S.T., *supra* note 45, at 805 ("It should be noted that Rule 34 does not require inspection of writings when it appears that the information contained therein is equally available to the moving party from an independent source.").

79. *Biferato*, 11 F.R.D. at 46; *see also* Reeves v. Pa. R.R. Co., 80 F. Supp. 107, 109 (D. Del. 1948) ("Clearly there is no showing that the photographs are in [plaintiff's] possession. The defendant suggests that a request of the plaintiff for the custody of the photographs would probably be complied with. No right is shown to exist in the plaintiff to enforce a demand for such X-ray photographs from the person having the possession, custody or control and the lack of a right to enforce the demand negatives the legal right of custody or control.").

80. *See Enforcing Discovery of Documents Under Federal Rule 34: The Effect of Foreign Law on the Concept of Control*, 62 YALE L.J. 1248, 1248–49 (1953) ("Whether a party controls designated documents is a question of fact."). *But cf.* Humphries v. Pa. R.R. Co., 14 F.R.D. 177, 178 (N.D. Ohio 1953) (failing to discuss possession, custody, or control relating to a request "seek[ing] production of 'all' statements of defendant's employees, agents, crew members and witnesses, 'all' photographs, sketches and reports submitted by Republic Steel to defendant and 'all' photographs and sketches taken or drawn by defendant's agents").

81. 63 F.R.D. 135 (S.D.N.Y. 1972).

hearings.[82] Agency policy dictated that the SEC would not produce witness testimony for private litigation but that witnesses were permitted to acquire copies of their own testimony.[83] Here, one employee had resisted doing so in response to plaintiffs' request.[84] Reasoning that the company "plainly" had Rule 34 control over its employees, the court ruled that the plaintiffs could "request Douglas [Aircraft Co.] to have its non-defendant employees procure copies of their private testimony before the SEC so that Douglas may give same to plaintiffs."[85] For a former employee, the court dismissively observed that "[t]here is no assertion . . . [he] is unwilling to cooperate," so the company was "required to produce a copy" of his testimony.[86] In so deciding, the court treated current and former employees as subject to company control. Additionally, in *Gray v. Faulkner*,[87] a suit by an inmate against prison officials, the court placed on the responding party "an affirmative duty to seek that information reasonably available to him from his employees, agents, or others subject to his control."[88] In comparison to the earliest cases on possession, custody, or control, those in the late-twentieth century expanded the definition of control through more deliberate analysis, imposing greater burdens on responding parties to bring information within the scope of Rule 34.

Courts encountered technology that was the precursor to the modern mobile device in the early aughts. For example, in *Hatfill v. New York Times, Co.*,[89] a plaintiff sued the *New York Times* for defamation and requested access to a reporter's interview notes kept on a flash drive.[90] The court assessed the case's unique circumstances and written agreements between the reporter and newspaper.[91] It referenced the "legal

82. *Id.* at 136.

83. *Id.* at 137.

84. *Id.*

85. *Id.* at 138; *see also In re* Domestic Air Transp. Antitrust Litig., 142 F.R.D. 354, 356 (N.D. Ga.1992) (requiring companies to get consent from employees to procure transcripts of deposition testimony given in a Department of Justice investigation).

86. *Herbst*, 63 F.R.D. at 138.

87. 148 F.R.D. 220 (N.D. Ind. 1992); *cf.* Cooper Indus., Inc. v. British Aerospace, Inc., 102 F.R.D. 918, 919–20 (S.D.N.Y. 1984) (relaxing the standard for control in the direction of practical ability, saying "The documents plaintiff seeks all relate to the planes that defendant works with every day; it is inconceivable that defendant would not have access to these documents and the ability to obtain them for its usual business."); Golden Trade, S.r.L. v. Lee Apparel Co., 143 F.R.D. 514, 525 (S.D.N.Y. 1992) (using the term "practical ability" for the first time, writing, "[I]n practice the courts have sometimes interpreted Rule 34 to require production if the party has the practical ability to obtain the documents from another, irrespective of his legal entitlement to the documents.").

88. *Gray*, 148 F.R.D. at 223.

89. 242 F.R.D. 353 (E.D. Va. 2006).

90. *Id.* at 354.

91. *Id.* at 355.

right" to obtain information as supporting Rule 34 control, but acknowledged that the Fourth Circuit had not interpreted the meaning of possession, custody, or control.[92] It emphasized that the interview notes were not stored on *Times* computers, the *Times* had a policy in its collective bargaining agreement with its reporters' union to allow journalists to retain their unpublished materials, and the *Times* routinely permitted departing reporters to keep their documents.[93] Nor had the reporter shared his notes with anyone, and he had kept the flash drive in his personal possession despite attaching it to company computers for work purposes.[94] Based on these indicia of the reporter's independence, the court deemed the *Times* to have "ceded any legal rights" to the notes.[95] Finally, the court pointed to "the clear substantive purpose" of the *Times*' relevant policies.[96] Considering this request for technologically stored information, the court examined the responding party's behavior and company guidelines to find no Rule 34 control.

V. MOBILE DEVICES AT WORK

Many employees use cell phones in their work, whether employers like it or not. In a 2011 study, carried out only a few years after the iPhone was launched, seventy-five percent of companies allowed employees to use their personal devices for work.[97] And in the years since, that figure has likely increased. People and their phones are practically inseparable, so it is unsurprising that fifty percent of people use their personal devices for work while on vacation, twenty-nine percent while in bed, and five percent while in a place of worship.[98]

One response was straightforward: corporations began issuing company-owned phones to key employees. By doing so, they secured clear control for discovery purposes and avoided disputes. Courts often

92. *Id.*

93. *Id.*

94. *Id.* at 354.

95. *Id.* at 355.

96. *Id.* Additionally, in *Han v. Futurewei Technologies, Inc.*, which dealt with personal computer hard drives, the court primarily discussed Rule 26 but quoted the Advisory Committee Notes to Rule 34 for the proposition that "testing and sampling" of ESI "*is not meant to create a routine right of direct access to a party's electronic information system*." No. 11-CV-831-JM (JMA), 2011 WL 4344301, at *3 (S.D. Cal. Sep. 15, 2011) (emphasis in original) (citing FED. R. CIV. P. 34 advisory committee's note to 2006 amendment). The court questioned the request's relevance, intrusiveness, and burdensomeness to deny the requesting party access to the responding party's personal computer hard drives. *Id.* at *4–6.

97. Mathiason & McGuire, *supra* note 5, at 4.

98. *Id.* at 4–5; *see also* The Sedona Conference, *Commentary on BYOD: Principles and Guidance for Developing Policies and Meeting Discovery Obligations*, 19 SEDONA CONF. J. 495, 502 (2018) [hereinafter *Commentary on BYOD*] ("organizations are permitting or encouraging workers to use their own personal devices to access, create, and manage their information—often after hours and outside the office").

held that documents and communications on company-issued devices were within the employer's possession or control.[99] This approach, however, was costly and often impractical. As smartphones became ubiquitous, many companies adopted "Bring Your Own Device" policies, which allowed or required employees to use personal devices for work. These policies often included provisions reserving rights for the employer to access, inspect, or remotely wipe data. BYOD adoption is estimated at over eighty percent across industries.[100] BYOD allowed employees to carry out personal and business functions, including accessing proprietary mobile applications.[101]

Many employers established policies governing BYOD use.[102] BYOD programs rose in popularity during the 2000s and 2010s.[103] BYOD policies gained favor first in IT circles and subsequently in a myriad of professions.[104] In practice, though, "the idea behind BYOD has existed ever since employees began bringing their own flash drives."[105] Employees using their own devices for business is exceedingly common, and the line between the personal and professional has blurred.

BYOD programs present numerous benefits on which companies want to capitalize. BYOD offers cost savings, reducing employers' expense on technology by leveraging employees' own investment in devices.[106] It boosts morale and engagement, as employees can use their preferred, familiar devices.[107] It solves the "two pocket problem," condensing the personal and professional into a single device.[108] Furthermore, BYOD has the effect of modernizing the workplace, as employees buy the latest devices that increase productivity and

99. *See, e.g.*, Lalumiere v. Willow Springs Care, Inc., No. 1:16-CV-3133-RMP, 2017 WL 6943148, at *2 (E.D. Wash. Sep. 18, 2017) ("A company controls the text messages of its employees on work phones, as well as the emails of its employees via company email accounts.").

100. Curt Steinhorst, *The Company Cell Phone Comeback: A Win-Win for Businesses & Employees*, FORBES (May 8, 2024), https://www.forbes.com/sites/curtsteinhorst/2024/05/08/the-company-cell-phone-comeback-a-win-win-for-businesses--employees/ [https://perma.cc/9684-G8EW].

101. Lindsey Blair, *Contextualizing Bring Your Own Device Policies*, 44 J. CORP. L. 151, 151 (2018).

102. *Id.*

103. *Id.* at 151–52.

104. *Id.* at 152.

105. *Id.*

106. Mathiason & McGuire, *supra* note 5, at 6; *Bring Your Own Device (BYOD) vs. Corporate-Owned: Which Policy Is Better for Your Needs?,* MODEONE (Mar. 21, 2023), https://modeone.io/blogs/bring-your-own-device-byod-vs-corporate-owned-which-policy-is-better-for-your-needs/ [https://perma.cc/L4L3-HFJV] [hereinafter *Bring Your Own Device*].

107. Mathiason & McGuire, *supra* note 5, at 6; *Bring Your Own Device*, *supra* note 106.

108. Mathiason & McGuire, *supra* note 5, at 6.

connectivity.[109] BYOD presents intangible social benefits and practical efficiency advantages.

That said, BYOD has several deficiencies. Of primary importance here are the legal ones. "[T]ime has exposed legal issues for both employers and employees."[110] Many discovery matters turn "on whether an employer is in 'control'" of BYOD devices because they are not physically possessed.[111] When faced with litigation or the likelihood thereof, employers with BYOD programs must inform their employees of the duty to preserve any relevant data on their dual-use devices.[112] The ramifications for not placing a timely hold on relevant information are significant—"harsh sanctions and penalties."[113] Companies with less control over employees' mobile devices may also lose the opportunity in eDiscovery to scan ESI for relevance and privilege.[114] In the absence of a formal BYOD policy, companies may still have a personal-device workplace.[115] Companies with BYOD workers risk being held responsible and incurring penalties for employee evidence spoliation.[116] Companies have thus been required to implement and enforce their BYOD guidelines and litigation hold practices over employee devices.[117]

---

109. *Id.*

110. Blair, *supra* note 101, at 151.

111. *Id.* at 153; *see also Commentary on BYOD*, *supra* note 98, at 531 ("In the BYOD context, the concept of "control" can be particularly murky and ripe for disputes due to the overlap of personal and business information on the device, as well as the physical possession and ownership of the device by the employee.").

112. Mathiason & McGuire, *supra* note 5, at 17–18, 21.

113. Blair, *supra* note 101, at 161; *see also Commentary on BYOD*, *supra* note 98, at 528 ("Organizations cannot ignore their discovery obligations merely because a device containing unique, relevant ESI is also used for personal purposes.").

114. Blair, *supra* note 101, at 161.

115. *Id.*

116. Mathiason & McGuire, *supra* note 5, at 23.

117. The Sedona Conference, *Commentary on Discovery of Mobile Device Data*, 26 SEDONA CONF. J. 735, 752–53 (2025) [hereinafter *Commentary on Mobile Device Discovery*]. Among BYOD policy factors that may influence a judicial finding of Rule 34 control are:

> (i) require[ments for] employees to cooperate with company requests for information on or access to mobile devices in their possession; (ii) specifi[cation] that the organization retains ownership of or control over any business information on an employee's personal device at all times; (iii) permi[ssion for] employees to use personal devices for company business and to access company systems, and whether such permissions are granted in exchange for the organization's right to obtain the device or access or collect data on the device on demand; (iv) state[ments] that an employee waives any rights or expectations of privacy with respect to their personal devices or data on those devices; (v) require[ments for] employees to waive any rights or expectations of privacy as a condition of using the device to communicate about company business or access

Furthermore, uncontrolled data on employee devices is an information security threat. "[T]he biggest risk to companies implementing a BYOD program is corporate data loss."[118] Company data on personal devices poses challenges for eDiscovery and business integrity.

To address BYOD weaknesses, companies have altered their practices. To that end, mobile device management (MDM) software has grown. MDM is "a very popular safety and administrative solution" enabling IT departments to "control and secure" devices connected to a company's network.[119] With MDM services, companies can encrypt, lock up or wipe clean, track, and access employees' personal mobile devices.[120] If a company could "containerize" its data on an employee's personal device, it would gain control of and insight into the information.[121] Setting up an MDM system in place is useful for companies, because they "can't just expect to have every employee somehow bring in their device when an e-discovery event happens and try to cultivate data at that time."[122] Regulators are also urging companies to exercise greater control over employee devices. The Financial Industry

---

> company systems; or (vi) . . . silen[ce] on giving an organization access to the employees' mobile device, and to personal content on the mobile device.

Id.; *see also Commentary on BYOD*, *supra* note 98, at 531 ("A 'consent' or 'acknowledgement' or other agreement that the employee signs and that recognizes that the organization owns or controls the ESI would likewise give the organization possession, custody, or control of the ESI, and the resulting obligation to consider the device when meeting its discovery obligation.").

118. Blair, *supra* note 101, at 166; *see also* Jacob A. Young, Tyler J. Smith, & Shawn H. Zheng, *Call Me BIG PAPA: An Extension of Mason's Information Ethics Framework to Big Data*, 2020 J. THE MIDWEST ASS'N FOR INFO. SYS. 30–31 (2020), https://www.metrostate.edu/sites/default/files/2020-09/volume-july-2020-oneclickprintable.pdf [https://perma.cc/ZQ2B-KJAH] (declaring the BYOD phenomenon "leads to several privacy and security implications, especially as it pertains to data breaches stemming from lost and stolen devices.").

119. *Bring Your Own Device*, *supra* note 106.

120. Blair, *supra* note 101, at 166. For a discussion of BYOD policies and MDM systems in relation to employee privacy, *see generally* Lauren E. Douglas, *Is Workplace Privacy Dead? The Effects of "Bring Your Own Device" Policies on Employee Privacy*, WAKE FOREST L. REV. CURRENT ISSUES BLOG (Mar. 11, 2022), https://www.wakeforestlawreview.com/2021/03/is-workplace-privacy-dead-the-effects-of-bring-your-own-device-policies-on-employee-privacy/ [https://perma.cc/JJ6A-9W92].

121. *Proliferation of BYOD Leads to E-Discovery Headaches*, 51 INFO. MGMT. 11, 11 (2017). *But cf. Commentary on BYOD*, *supra* note 98, at 555 (raising privacy concerns and saying, "There are unique legal challenges to the successful implementation of a BYOD program, particularly in the international context and due mainly to data privacy and data protection laws. In the European Union and many other jurisdictions, data privacy is considered a human right.").

122. Jaclyn Jaeger, *e-Discovery Challenges in Mobile Computing Era*, COMPLIANCE WK. (Feb. 2012) (internal quotations omitted); *see also Commentary on Discovery of Mobile Device Data*, *supra* note 117, at 754 ("Given the potential complexity of identifying, preserving, and collecting mobile device data, it is often advisable to meet and confer with opposing parties early in a matter and attempt to reach consensus on mobile device data that will be considered in-scope for the matter.").

Regulatory Authority (FINRA), for instance, directed financial firms to supervise, retrieve, and retain mobile-device communications.[123] Some companies are now returning to employer-provided phones.[124] MDM and "work phones" may afford companies security and control that was weakened by BYOD.

VI. JUDICIAL TESTS FOR "POSSESSION, CUSTODY, OR CONTROL"

There are various standards that courts have developed for determining possession, custody, or control.[125] Some courts, alternatively, combine these tests into a unique variation or apply an altogether different analytical framework. The three main tests are (1) legal right; (2) legal right plus notification;[126] and (3) practical ability. Insofar as one of these particular tests is used, legal right and practical ability are the most common.

A. *"Legal Right" Standard*

According to the legal right test, a party has possession, custody, or control when it has the legal right to obtain the documents on demand.[127] Some courts unhelpfully state the test as "possession, custody, or control, or has the legal right to obtain the documents on demand."[128] The legal right test is grounded in legal authority as the basis for control—a contract, statute, or policy, for example. A party's actual possession of information is not necessary, and the practical ability to procure information is not sufficient. Rather, some formal legal mechanism must exist conferring upon a party the right to such information.

---

123. Jaeger, *supra* note 122.

124. *See* Scott Moritz, *Work Phones Make a Comeback as Offices Ban WhatsApp, TikTok*, L.A. TIMES (Mar. 3, 2023), https://www.latimes.com/business/story/2023-03-03/work-phones-make-a-comeback-as-offices-ban-whatsapptiktok [https://perma.cc/Z7NZ-W3GU] ("Businesses, especially those in finance, have grown concerned about the security of their data, and the Securities and Exchange Commission and the Commodity Futures Trading Commission have stepped up their scrutiny over unauthorized private communication on applications such as WhatsApp and through personal email.").

125. *See Rule 34 and 45 Commentary*, *supra* note 15, at 2 (There is "a split of authority between circuits that apply a 'practical ability' standard, circuits that apply a 'legal right' standard, [and] those that have applied a 'legal right plus notification' standard.").

126. Few cases have mentioned this "Legal Right Plus Notification Standard" except by way of a passing reference. *See, e.g.*, St. Clair Cnty. Emp.'s Ret. Sys. v. Acadia Healthcare Co., No. 3:18-CV-00988, 2022 WL 4095387, at *9 (M.D. Tenn. Sep. 7, 2022), *aff'd*, No. 3:18-CV-00988, 2023 WL 3659734 (M.D. Tenn. May 25, 2023).

127. "Control is defined not only as possession, but as the legal right to obtain the documents requested upon demand." Searock v. Stripling, 736 F.2d 650, 653 (11th Cir. 1984).

128. *In re* Bankers Tr. Co., 61 F.3d 465, 469 (6th Cir. 1995); *see also Rule 34 and 45 Commentary*, *supra* note 15, at 19 n.10 (collecting cases). The legal right test has been used, though non-exclusively, in the Third, Fifth, Sixth, Seventh, Eighth, Ninth, Tenth, and Eleventh Circuits. *Id.* at 19.

B. *"Legal Right Plus Notification" Standard*

The legal right plus notification test is a close corollary of the legal right test. Like the latter, it requires a legal right to obtain relevant documents and ESI.[129] It also requires, as the name suggests, a notification. In the event that a party "cannot fulfill this duty to preserve because he does not own or control the evidence," then "he still has an obligation to give the opposing party notice of access to the evidence or of the possible destruction of the evidence if the party anticipates litigation."[130] This test applies when its condition precedent—inability to preserve or lack of control over evidence—is triggered. Thus, for determining control as under the Federal Rules, the analysis is substantially similar to that undertaken for the legal right test.

C. *"Practical Ability" Standard*

The practical ability test is widely adopted and more expansive than the legal right test. It requires production of documents and ESI "not in a party's possession . . . if a party has the *practical ability* to obtain the documents from another, irrespective of legal entitlements to [or actual possession of] the documents."[131] In other words, if a person or entity other than the responding party has actual possession of documents or ESI, the responding party may still have control under the Federal Rules if it may direct or influence the actual possessor. Under the practical ability standard, if one can easily obtain documents or ESI from another, then it has control for the purposes of discovery preservation and production.

---

129. *Rule 34 and 45 Commentary*, *supra* note 15, at 20.

130. Silvestri v. General Motors Corp., 271 F.3d 583, 591 (4th Cir. 2001); *see also Rule 34 and 45 Commentary*, *supra* note 15, at 24–27 nn.17–21 (collecting cases). The practical ability test has been used, though non-exclusively, in the Second, Fourth, Eighth, Tenth, Eleventh, and District of Columbia Circuits. *Id.* at 24.

131. Ice Corp. v. Hamilton Sundstrand Corp., 245 F.R.D. 513, 517 (D. Kan. 2007) (citation omitted); *see also Rule 34 and 45 Commentary*, *supra* note 15, at 22 n.15 (collecting cases).

D. *"Hybrid" Standards*

Some courts blend the two major tests,[132] others evaluate the nature of the employment relationship and the terms of any BYOD policies,[133] and some courts have considered prior practices of information sharing. For example, in *Stanz v. Brown*,[134] the Court recognized that control was defined as the legal right to obtain documents upon demand, recognized that a party has a duty to conduct a reasonable inquiry into the factual basis of discovery responses, and thereafter extrapolated that "a party responding to a Rule 34 production request is under an affirmative duty to seek that information reasonably available to it from its employees, agents, or others subject to its control."[135] The caselaw is filled with the blending of tests that suggest legal right is the standard for control and then extend beyond—whether intentionally or inadvertently—into a practical ability application.

VII. APPLICATION OF THESE TESTS TO PERSONAL EMPLOYEE MOBILE DEVICES

Because of variations in tests and their applications, and the circumstances of each case, outcomes in discovery disputes have likewise varied.

---

132. *See, e.g.*, *In re* Bankers Tr. Co., 61 F.3d 465, 469 (6th Cir. 1995) ("[D]ocuments are deemed to be within the 'possession, custody or control' for purposes of Rule 34 if the party has actual possession, custody or control, or has the legal right to obtain the documents on demand."); Prokosch v. Catalina Lighting, Inc., 193 F.R.D. 633, 636 (D. Minn. 2000) ("Specifically, control is defined as 'the legal right, authority, or ability to obtain upon demand documents in the possession of another.'") (quoting Florentia Cont. Corp. v. RTC, No. 92 Civ. 1188, 1993 WL 127187 at *3 (S.D.N.Y. Apr. 22, 1993)); Goodman v. Praxair Servs., Inc., 632 F. Supp. 2d 494, 515 (D. Md. 2009) ("Rule 34 'control' would not require a party to have legal ownership or actual physical possession of any documents at issue. Instead, documents are considered to be under a party's control when that party has 'the right, authority, or practical ability to obtain the documents from a non-party to the action.'") (quoting *In re* NTL, Inc. Secs. Litig., 244 F.R.D. 179, 195 (S.D.N.Y.2007)).

133. The "BYOD policy does not assert Hormel's ownership over any data other than data 'sourced from Hormel systems and synced between the mobile device and its servers'—which does not include text messages (except, perhaps, if the employee copied data sourced from a Hormel system and embedded it in a text)—nor does it assert Hormel's right to demand that its employees allow it to access or inspect any other data. Hormel's right and ability to remotely wipe an entire phone is for the sole and express purpose of removing company data—such as in response to the phone being lost or stolen. The company's ability to wipe personal data from a personally-owned device by resetting the device to a factory floor state in order to purge company data does not give the company control—legal or practical—over that personal data." *In re* Pork Antitrust Litig., No. 18-CV-1776 (JRT/HB), 2022 WL 972401, at *5 (D. Minn. Mar. 31, 2022).

134. No. 22-CV-01164-GPC-JLB, 2025 WL 1617102 (S.D. Cal. June 6, 2025).

135. *Id.* at *6 (citing A. Farber & Partners, Inc. v. Garber, 234 F.R.D. 186, 189 (C.D. Cal. 2006)).

A. *The "Legal Right"*

Contemporary cases employing the legal right test with respect to employee ESI look for clear signs of an agreement or arrangement ratifying control. In the context of employer control over employees' personal devices, the legal right test has generally been difficult to pass.

In *Cotton v. Costco Wholesale Corp.*,[136] the court assessed the nexus between an employee's personal phone and the contents thereon to determine whether text messages were discoverable under Rule 34.[137] A former Costco employee suing the company for discrimination requested the text messages of other Costco employees that were housed on these individuals' personal mobile devices.[138] In considering whether the messages fell within the parameters of Rule 34, the court invoked the legal right test.[139] The court reasoned that the cell phones were not issued by Costco to its employees and the employees represented that they did not use the cell phones for work purposes.[140] On that basis, the court determined that Costco did not have possession, custody, or control of the text messages and denied the plaintiff's motion to compel.[141] The court's analysis demonstrates that an employer does not have control over an employee's personal mobile device under the legal right test absent a mechanism of legal authority. Costco had neither provided employees with phones nor outlined an official use policy conferring power over personal mobile devices.

B. *Employee Handbooks and Agreements and "Legal Right"*

The court in *Krishnan v. Cambia Health Solutions, Inc.*,[142] identified several factors to consider in determining whether a company has control over personal employee devices and evaluated the terms of a company handbook governing electronic messaging systems.[143] The court stated that to determine whether an employer has legal right control over workers' personal devices, it evaluates whether the employer issued the devices, the devices were used for work, and the company had a right to obtain communications from the devices.[144] With respect to those factors, the court found employees had used personal devices to discuss work, but the cell phones were not company-owned and were not routinely used for

136. No. 12-2731-JW, 2013 WL 3819974 (D. Kan. July 24, 2013).
137. *Id.* at *6.
138. *Id.*
139. *Id.* (citing Noaimi v. Zaid, 283 F.R.D. 639, 641 (D. Kan. 2012)).
140. *Id.*
141. *Id.*
142. No. 2:20-cv-574-RAJ, 2021 WL 3129940 (W.D. Wash. 2021).
143. *Id.* at *2.
144. *Id.*

business.[145] The employer also had a company policy that the court weighed.[146] The applicable handbook provided, "employees 'do not have an expectation of privacy in the electronic messaging systems used during . . . employment.'"[147] It added that company management could access such data at its discretion.[148] Such access included reviewing, investigating, and disclosing "any use" of the "computer, email, instant messaging, telephone and internet systems."[149] The court, however, interpreted the handbook "language [to] clearly refer to employees' use of the *employer*'s devices and electronic systems, not employees' use" of personal mobile devices and text messaging abilities.[150] Thus, the court ruled that defendants were not in control of the requested messages.[151] Despite the employer maintaining a company policy on electronic messaging and employees exchanging messages relating to work on their personal cell phones, the plaintiff could not establish control under the legal right standard.

On the other hand, in *Weston v. Docusign, Inc.*,[152] the court found an explicit employment agreement pertaining to personal mobile device use sufficient to establish employer control by legal right.[153] In this securities fraud litigation, plaintiffs sought text messages from non-defendant custodians.[154] Plaintiffs had demonstrated that employees did carry out their work by text at times.[155] Most important for the court's analysis, Docusign's employment agreements expressly stipulated, "if employees used any personal device to transmit any Company information, they agree[d] . . . [the company] may have access to such personal devices to retrieve it."[156] Thus, despite not having actual possession of employees' phones, the court concluded the employer had the right to obtain the desired text messages.[157] When the employer has an agreement directly

145. *Id.*

146. *Id.*

147. *Id.*

148. *Id.*

149. *Id.* (internal quotations omitted).

150. *Id.* (emphasis in original).

151. *Id.*; *cf. In re* Citric Acid Litig., 191 F.3d 1090, 1108 (9th Cir. 1999) ("Ordering a party to produce documents that it does not have the legal right to obtain will oftentimes be futile, precisely because the party has no certain way of getting those documents."); Robinson v. Moskus, 491 F. Supp. 3d 359, 364 (C.D. Ill. 2020) (concluding "a party's 'practical ability' to obtain the documents is irrelevant absent a legal right to do so").

152. No. 22-cv-00824-WHO, 2024 WL 3446924 (N.D. Cal. July 15, 2024).

153. *Id.* at *7.

154. *Id.* at *3.

155. *Id.* at *6.

156. *Id.* at *7 (citation modified).

157. *Id.*; *cf.* Order RE: Player Agent Discovery Dispute, *Jones v. PGA Tour, Inc.*, No. 5:22-CV-04486 at *2 (N.D. Cal. Nov. 17, 2022) (applying the legal right test in response to a request for custodial ESI from agents representing professional golfers, finding "It is not credible that the

establishing its authorization to access employees' personal devices, that establishes control under the legal right test.

But, in *Halabu Holdings, LLC v. Old National Bancorp*,[158] the court found that "without further indicia of control" or employment policies, the requesting party failed to establish the employer's legal right to control employees' personal cell phones.[159] The plaintiff had requested records generated in the scope or course of employees' "performance of their responsibilities for the [b]ank."[160] The court, though, found the plaintiff's argument, "that information on an employee's personal device may be compelled by way of a document request directed to the employer without any further indicia of control over the device by the employer," to be without "support in any case . . . for good reason."[161] The court rejected such a principle as seriously infringing employees' rights to privacy and property.[162] Rule 34 control, the court explained, "is context-specific."[163] As the court described, "[s]ome workplaces have, by agreement and practice, defined rights and responsibilities regarding personal cell phone and computer device use in connection with the employer's business."[164] Company contracts that provide a right of access to employees' personal devices and agreements that manage corporate information thereon "are entitled to judicial respect."[165] In this case, there were no writings or standards of that sort to which the plaintiff could point.[166] Moreover, the plaintiff "made no effort to establish . . . control over the personal devices."[167] Absent proof of a "set of circumstances

---

Players cannot obtain responsive documents from their own agents," and thus concluding "custodial ESI (email and device-level data) in the actual possession of the Players' agents is within their 'control' within the meaning of Rules 34 and 45."); Miniace v. Pac. Mar. Ass'n, No. C 04-3506 SI, 2006 WL 335389, at *7–8 (N.D. Cal. Feb. 13, 2006) (directing that documents in possession of current directors be produced).

158. No. 20-10427, 2020 WL 12676263 (E.D. Mich. June 9, 2020).

159. *Id.* at *2–3; *see also* Rattie v. Balfour Beatty Infrastructure, Inc., No. 22-cv-05061-RS (LJC), 2023 WL 5507174, at *4 (N.D. Cal. Aug. 25, 2023) (finding that the plaintiff "failed to establish a sufficient factual record" to request personal text messages from management witnesses. The court also highlighted that the responding party provides work-issued phones to management and expects them to use such phones for work. On these grounds, the court deemed the defendant–employer not to have control under the legal right test.).

160. 2020 WL 12676263, at *1.

161. *Id.* at *2.

162. *Id.* The court also explained that "workplaces are ever evolving" through remote work and smartphones, thus confusing the distinction between personal and work. *Id.* at *3. "But an employee's sense of privacy and ownership of information should not be forfeited without an adequate discovery process to address those interests." *Id.*

163. *Id.*

164. *Id.*

165. *Id.*

166. *Id.*

167. *Id.*

demonstrating control," the requesting party did not satisfy its burden.[168] Finally, the court advised that Rule 45 subpoenas were available and proper.[169] As *Halabu* reflects, the legal right test prioritizes formal agreements in demonstrating Rule 34 control. In practice, courts have used the legal right test to good effect where clear contractual or regulatory rights exist—for example, in enforcing BYOD acknowledgments that expressly reserve employer access to work content or in regulated industries with detailed record-keeping rules. Yet those cases highlight the test's limits: actors who wish to avoid discovery can structure relationships to leave critical communications outside any formal entitlement, even as they remain central to corporate decision-making.

Defenders of the legal right test argue that its under-inclusiveness is a virtue, not a defect, because it respects personal-property boundaries and encourages careful drafting of policies and agreements. That virtue, however, depends on an unrealistic assumption that employees freely bargain over device-access terms and that regulators and courts can keep pace with new technologies and practices that make work communications appear "personal" by design.

C. *"Practical Ability"*

Courts applying the practical ability test take a permissive tack. The test is less formalistic and more pragmatic than the legal right test. Courts applying it may impose expectations and obligations on responding parties beyond what would be required under the legal right test. If the circumstances evince the ability to obtain requested information, courts are likely to compel production.

In *Ronnie Van Zant, Inc. v. Pyle*,[170] the court analyzed whether a movie production company had control over a director it had hired to determine whether to impose Rule 37[171] spoliation sanctions in light of unpreserved text messages.[172] After a dispute arose in the planning of a movie about Lynyrd Skynyrd, the director switched cell phone providers, acquired a new cell phone, and did not back up existing text messages that included correspondence with the plaintiff.[173] The production company contested the motion for sanctions, arguing that it had not

168. *Id.*

169. *Id.*

170. 270 F. Supp. 3d 656 (S.D.N.Y. 2017), *vacated*, No. 17-cv-3360, 2019 WL 13550673 (S.D.N.Y. Jan. 7, 2019). While the 2017 decision was vacated by order of the Second Circuit, the analysis therein remains instructive.

171. *See* FED. R. CIV. P. 37(e) (imposing sanctions for failure to preserve ESI).

172. *Ronnie Van Zant*, 270 F. Supp. 3d at 669.

173. *Id.* at 667.

controlled the director's phone and the information that was lost.[174] The court, reasoning through the practical ability test, underscored that the director had contracted, worked closely, and shared litigation documents with the production company.[175] Additionally, the director had a financial interest in the litigation's outcome because he was owed a percentage of the proposed film's earnings.[176] Though "determining practical control is not an exact science," the court decided that the aforementioned factors and "common sense" established that the production company, "practically speaking," controlled the director's text messages.[177] A professional relationship and record of close collaboration indicate that a responding party has the practical ability to control information in the hands of such non-party.

Similarly, if a corporate officer possesses information as part of her company duties, the employer is likely to control it under the practical ability standard. In *CA, Inc. v. AppDynamics, Inc.*,[178] corporate officers used personal email accounts for work purposes.[179] The court presented the issue as "whether the documents were created in connection with the officer's functions as a corporate employee."[180] And if they were, then the corporation retains a proprietary interest in the documents and the employee has to turn them over.[181] Put simply, "corporations have control over their officers and employees sufficient to require production of the documents in their possession."[182] Thus, if an employee produces information in furtherance of her job as an officer, the information is within the company's control and must be disclosed.[183] Here, company founders and officers had personal email accounts which the plaintiff suspected with "some justification" to house relevant emails.[184] Testimony showed that on "at least one occasion," a corporate officer forwarded one email from his work account to his personal account.[185] Another officer testified, "albeit somewhat ambiguously," that he had

174. *Id.* at 669.

175. *Id.* ("Cohn was contracted by Cleopatra to work on the Film, and the evidence has establishes [sic] that he worked closely with Cleopatra for over the past year. Over the course of the instant litigation, Cohn has participated by providing documents and took a deposition sought by Plaintiffs during discovery.").

176. *Id.*

177. *Id.*

178. CV 13-2111 (WFK) (SIL), 2014 WL 12860591 (E.D.N.Y. Sep. 8, 2014).

179. *Id.* at *4; *but see* Owen v. Elastos Found., 19-CV-5462 (GHW) (BHM), 2023 WL 2537287, at *2 (S.D.N.Y. Mar. 16, 2023) (declining to find control over a corporate director's personal Gmail account).

180. *CA*, 2014 WL 12860591, at *3.

181. *Id.*

182. *Id.*

183. *Id.*

184. *Id.* at *4.

185. *Id.*

used his personal account in creating the company.[186] The court found "these examples sufficient to direct that Defendant produce responsive documents maintained on these accounts."[187] Doing so, the court opined it "worth recognizing" that concluding otherwise "would not make sense."[188] To rule that a company did not have Rule 34 control, the court asserted, "would be tantamount to an invitation for executives" to maintain two email accounts—"one discoverable and the other not."[189] Thus, the court deemed the practical ability test satisfied on a modest evidentiary showing and without imposing a strict burden on the requesting party.

The practical ability test responds to precisely this evasion concern by instructing courts to look beyond formal entitlements to the realities of how information flows within an organization. Properly applied, it can prevent parties from defeating discovery by offloading critical communications to loosely affiliated insiders or personal devices.

At the same time, practical ability is vulnerable on the very criteria that matter most in the mobile-device context: it provides little ex ante guidance to parties designing compliance programs, invites highly fact-bound mini-trials about "ability" in every case, and can be read to support de facto mandates that employers pressure employees to surrender access to deeply personal devices.

D. *Hybrid Applications*

Often, courts enunciate that they are using one test but then start to blend other concepts into their analysis. Some courts modify, reinterpret, distinguish, or decline to follow the prevailing possession, custody, or control tests. That contributes to inconsistency in the law and unpredictability for litigants. When courts lean into judicial discretion, case outcomes can vary widely.

For example, the court in *Goolsby v. County of San Diego*,[190] stated it was applying a legal right test, but then started to assess whether employees' personal devices had been used for work purposes to thereby establish employer control.[191] A prisoner had filed a civil rights action against the county arising out of his stay in the San Diego County Jail.[192] He sought messages from deputies' personal phones, alleging that employees carried and used such devices at work.[193] The county denied

186. *Id.*
187. *Id.*
188. *Id.*
189. *Id.*
190. No. 3:17-cv-564-WQH-NLS, 2019 WL 3891128 (S.D. Cal. Aug. 19, 2019).
191. *Id.* at *4.
192. *Id.* at *1.
193. *Id.* at *4.

possession, custody, or control of the personal mobile devices but had instructed the employees to search their work-issued phones and computers for responsive information.[194] Though the court acknowledged that federal courts are divided on what constitutes Rule 34 control, it relied on the legal right test and cited to *Cotton*.[195] "Generally," the court said, "the plaintiff must show that personal devices were used for business purposes."[196] The burden of showing responding-party control over the requested ESI lies with the requesting party.[197] Here, though, the plaintiff "only offer[ed] his own speculation that deputies use[d] their personal devices while at work."[198] Additionally, no evidence supported that deputies' use of personal devices was for their work.[199] Finally, the court explained that the defendant had already searched government-owned devices and attested to producing all relevant information.[200] "Under these circumstances," the court determined, "there [wa]s nothing left to compel."[201] Because the plaintiff could not show that the county had any legal basis for claiming a right to its employees' personal phones on these facts, the employer did not have control.

In *Servicios Funerarios GG, S.A. de C.V. v. Advent International Corp*.,[202] the court analyzed the issue of control even though it recognized that the employee was using a company-issued cell phone. Hence, it should have just found that the company possessed the device.

A court may decline to follow a test but still reason on the basis of both. In *In re Pork Antitrust Litigation*,[203] the court explained the legal right and practical ability tests but declined to find company control over text messages on employees' personal phones "regardless of the standard applied."[204] The court stated that the Eighth Circuit had not defined "control" but that different courts employ the legal right and practical ability tests, respectively.[205] It presented each test, methodically detailing

194. *Id.*
195. *Id.*
196. *Id.*
197. *Id.*
198. *Id.*
199. *Id.*
200. *Id.*
201. *Id.*
202. Servicios Funerarios GG, S.A. de C.V. v. Advent Int'l Corp., No. 23-10684-IT, 2023 WL 8544790, at *1 (D. Mass. Dec. 11, 2023), *aff'd*, No. 1:23-cv-10684-IT, 2024 WL 2748348, at *2 (D. Mass. May 29, 2024).
203. *In re* Pork Antitrust Litig., No. 18-cv-1776 (JRT/HB), 2022 WL 972401, at *3 (D. Minn. Mar. 31, 2022).
204. *Id.* at *4. The plaintiffs also complained, reminiscent of the legal right plus notification test, that the responding party had not notified them that it disclaimed control of the desired cell phones. *Id.* at *1.
205. *Id.* at *3.

the strengths and weaknesses.[206] It decided, though, that "[i]n this case," it "need not choose between" them.[207] In its factual examination, the court described the responding party's BYOD policy as providing ownership of company data but disclaiming ownership of text messages on personal devices.[208] The court rejected the BYOD justification for control.[209] The court evaluated a practical-ability-grounded argument premised on the employer-employee relationship.[210] While the court accepted that "custodians may feel a sense of company loyalty," or "have an interest in the company's financial health," the court was sensitive to the power dynamics of an employer-employee relationship and thought it went "too far" to conclude that an employer would have a practical ability to control texts on personal phones.[211] That reasoning conflicts with *CA, Inc.*, however, which proclaimed the opposite while also reasoning through a practical ability analysis: "[C]orporations have control over their officers and employees sufficient to require production of the documents in their possession."[212] While the court did not declare its adherence to either predominant test, it considered both and vacillated between them, which does not provide litigants a clear sense of direction.

In the absence of a definitive test for control, courts may apply the law inconsistently. In *Allergan, Inc. v. Revance Therapeutics, Inc.*,[213] the court discussed both tests, selected one, and ultimately waffled between them.[214] Much like in *In re Pork*, the court outlined both tests and posited,

206. *See id.* at *4 (noting that it may be futile to command a party to produce documents it does not have the right to obtain but pointing out "a strong argument" that a party's relationship with a non-party may establish that "fairness would require" producing documents).

207. *Id.* at *4.

208. *Id.* at *5.

209. *Id.* at *6.

210. *Id.*

211. *Id.* In *Sports Rehab Consulting LLC v. Vail Clinic, Inc.*, No. 1:19-cv-02075-WJM-SBP, 2025 WL 1144559, at *8 (D. Colo. Apr. 18, 2025), the court similarly wrote, "Surely, a typical employee does not expect to be compelled to turn over their personal cell phone simply because their employer has become embroiled in litigation . . . . And such an employer-directed document request plainly implicates both the privacy and property rights of the targeted employee." As in *In re Pork*, the *Sports Rehab* court referenced both the legal right and practical ability tests, finding no control over personal device text messages under either. *Id.* at *6. It concluded, however, by combining the logic underlying each test, stating that the requesting party "failed to meet their obligation to show that Vail Health had the *legal right—coupled with the practical ability*—to obtain and produce text messages on the personal devices of either current or former board members or employees." *Id.* at *9 (emphasis added).

212. *CA*, 2014 WL 12860591, at *3.

213. No. 3:23-cv-00431, 2025 WL 984792 (M.D. Tenn. Mar. 17, 2025).

214. *Id.* at *9; *see also* Kelly Twigger, *Episode 171: Is a BYOD Policy the Key to Determining Employer Control Over Employees' Personal Phones?*, MINERVA 26 (Apr. 20, 2025), https://minerva26.com/episode-171-is-a-byod-policy-the-key-to-determining-employer-control-over-employees-personal-phones/ [https://perma.cc/EA2L-NW5B] (discussing *Allergan*); *cf.* Selectica, Inc. v. Novatus, Inc., No. 6:13-cv-1708-Orl-40TBS, 2015 WL 1125051,

"[w]ithin these varying applications of Rule 34(a) 'control' standards, relatively few courts have addressed . . . the context of an employer's Rule 34(a) control over its employees' personal devices."[215] The court, then, cited with approval *In re Pork*.[216] Ultimately, though, it decided to "apply the 'legal right' standard."[217] In response to the requesting party's argument that the employer had control over employees' personal devices because they were used for work purposes, the court deemed that suggestion lacking.[218] In doing so, it overlooked and implicitly rejected reasoning employed in *Goolsby* and *Krishnan*, other cases which considered Rule 34 control under the legal right test.[219] After "adopt[ing] and appl[ying]" the legal right test, the court "further rule[d]" that the responding party did not have control based on the practical ability test, either.[220] The court concluded the practical ability test failed to establish control "for the same reasons" that the legal right test fell short.[221] By giving that assessment, the court collapsed the distinction between the two tests, even as the legal right test is tailored to requiring a formal mechanism of legal control, whereas the practical ability test is rooted in the functional aspects of employer-employee relations.

Courts may perform an individualized inquiry separate from either test based on the facts of a particular case. In *Miramontes v. Peraton, Inc.*,[222] the court acknowledged the existence of the legal right test but rejected any one-size-fits-all standard in favor of a fact-based analysis.[223]

---

at *4–5 (M.D. Fla. Mar. 12, 2015) (citing the practical ability test but subsequently invoking the employer–employee relationship as one "'which evidence[s] a *legal right* . . . to obtain a document.'") (citation omitted) (emphasis added).

215. *Allergan*, 2025 WL 984792, at *6.

216. *Id.* at *7.

217. *Id.*

218. *Id.*

219. *See* Goolsby v. Cnty. of San Diego, No.:3:17-cv-564-WQH-NLS, 2019 WL 3891128, at *4 (S.D. Cal. Aug. 19, 2019) (describing the standard for the legal right test: "[G]enerally the plaintiff must show that personal devices were used for business purposes."); Krishnan v. Cambia Health Sols., Inc., No. 2:20-cv-574-RAJ, 2021 WL 3129940, at *2 (W.D. Wash. 2021) (identifying whether a personal device was used for work as a factor to consider in evaluating a party's legal right to control a non-party's personal device).

220. *Allergan*, 2025 WL 984792, at *9.

221. *Id.*

222. No. 3:21-CV-3019-B, 2023 WL 3855603 (N.D. Tex. June 6, 2023).

223. *Id.* at *5; *see also* H.J. Heinz Co. v. Starr Surplus Lines Ins. Co., No. 2:15-cv-00631-AJS, 2015 WL 12791338 (W.D. Pa. July 28, 2015), *report and recommendation adopted*, 2015 WL 12792025, at *4 (W.D. Pa. July 31, 2015) (finding employer control over employees' personal phones under a BYOD policy without discussing any possession, custody, or control test). *But see* Ewald v. Royal Norwegian Embassy, No. 11-CV-2116 SRN/SER, 2013 WL 6094600, at *10 (D. Minn. Nov. 20, 2013) (declining to order production of personal mobile devices without discussing either possession, custody, or control test where the requesting party did not "demonstrate[] her entitlement to such devices"); *but cf.* Bratcher v. Navient Sols., Inc., 249 F. Supp. 3d 1283, 1286 (M.D. Fla. 2017) (declining to allow the requesting party to inspect

The responding party claimed it had no control over employee communications without a "legal right to obtain them on demand."[224] The court appreciated the "intuitive appeal" of such a "bright-line test," but declined to follow one.[225] Rather, it explained that "the realities of modern business require a fact-specific approach."[226] The court found persuasive that the responding company did not provide company-owned devices to employees and that employees often used their phones for work.[227] As the court laid out, most employees use their cell phones for work, many companies do not provide work-issued devices, and courts following the legal right test would permit employers to shield their information on employees' personal phones from discovery.[228] In recognition of those pragmatic considerations, the court "agree[d] with other courts" that ESI on "employees' personal devices may be under the control of their employer in certain circumstances."[229]

VIII. EMPLOYER RESPONSIBILITY FOR EMPLOYEE SPOLIATION

Employers may be held responsible for employee spoliation. Spoliation arises when evidence is destroyed, materially altered, or inadequately preserved for another party's use in pending or reasonably foreseeable litigation.[230] Courts' authority to administer sanctions for spoliation derives from their inherent power to manage judicial processes and Rule 37.[231] The integrity of the justice system is the guiding

---

a cell phone "akin to a computer" in a Rule 26 request because "'[T]here is no routine right of direct access to a party's electronic information system.'") (citation omitted); Weisman v. Barnes Jewish Hosp., No. 4:19-CV-00075 JAR, 2022 WL 850772, at *3 (E.D. Mo. Mar. 22, 2022) (declining to require production of Telegram messages without discussing either test because of the requesting party's failure to show control by "not provid[ing] the court with information sufficient to . . . question the truthfulness of Defendants' representation that no responsive document exist."). For additional analysis of the *Miramontes* decision, *see* Kelly Twigger, Episode 117: *Will an Employer's Failure to Preserve Text Messages on an Employee's Personal Mobile Devices Result in Sanctions?*, MINERVA 26 (Sep. 7, 2023), https://minerva26.com/episode-117-will-an-employers-failure-to-preserve-text-messages-on-an-employees-personal-device-result-in-sanctions/ [https://perma.cc/9ABV-4Q8H].

224. *Miramontes*, 2023 WL 3855603, at *5.

225. *Id.* In *Lalumiere v. Willow Springs Care, Inc.*, NOL 1:16-CV-3133-RMP, 2017 WL 6943148, at *2 (E.D. Wash. Sep. 18, 2017), the court followed a bright-line test of a different sort. Without mentioning either the legal right or practical ability tests, the court simply stated as a rule that "a company does not possess or control the text messages from the personal phones of its employees and may not be compelled to disclose text messages from employees' personal phones."

226. *Miramontes*, 2023 WL 3855603, at *5.

227. *Id.*

228. *Id.*

229. *Id.*

230. Pension Comm. of Univ. of Montreal Pension Plan v. Banc of Am. Sec., LLC, 685 F. Supp. 2d 456, 465 (S.D.N.Y. 2010).

231. *Id.*; *see also* FED. R. CIV. P. 37(e) (authorizing sanctions).

principle.[232] The duty to preserve evidence is well established at common law and is triggered once a party reasonably anticipates litigation.[233] It is "crystal clear" that any "breach of the duty to preserve, and the resulting spoliation of evidence, may result in the imposition of sanctions."[234] Thus, a party must lift information destruction procedures and implement a litigation hold to preserve documents when litigation is reasonably anticipated.[235] If information is not adequately preserved and produced on employees' personal devices deemed under the employer's control, liability may still be imputed to the company and sanctions imposed.

It is important to distinguish three related but conceptually distinct questions: when a party has "control" for purposes of Rule 34 production; when a duty to preserve arises; and when sanctions are appropriate for loss of information. Existing cases sometimes elide these inquiries, using broad language about an employer's responsibility for its agents' acts of spoliation to justify equally broad production obligations. An agency-anchored test can clarify the relationship: it ties presumptive control—and thus the core preservation duty—to a defined subset of agents, while leaving sanctions doctrine to address egregious failures to supervise or deliberate efforts to circumvent discovery, even when control in the narrow sense is lacking.

A company is likely to be liable where its employee–agent spoliated evidence. In *Microvention, Inc. v. Balt USA, LLC*,[236] the plaintiffs in a trade secret misappropriation case alleged that the defendant's employees spoliated ESI on laptops and tablets by deleting thousands of files after litigation had commenced.[237] The defendant–employer argued that it did not control employees' personal devices and that it had fulfilled its obligations by issuing a hold with respect to information on work-issued laptops.[238] Reasoning that "[t]he current trend . . . [is] to impute liability for an agent's spoliation to the principal based on 'traditional notions of

232. *Pension Comm.*, 685 F. Supp. 2d at 465.

233. *Id.* at 466.

234. *Id.*

235. *Id.*

236. No. 8:20-cv-02400-JLS-KES, 2023 WL 7476998 (C.D. Cal. Oct. 5, 2023).

237. *Id.* at *1, *9.

238. *Id.* at *26. Additionally, *La Belle v. Barclays Cap., Inc.*, 340 F.R.D. 74, 84 (S.D.N.Y. 2022) is instructive. There, the court found that the "duty to search for messages on . . . [employees'] personal cellphones did not arise until there was some indication that . . . [relevant evidence] was contained on the personal devices." *Id.* Barclays had a company policy prohibiting employees from carrying out work on personal devices. *Id.* Barclays had preserved and produced ESI from company-owned devices already. *Id.* at 83. Reasoning that "it [wa]s a close question," the court did not consider it unreasonable for Barclays to assume employees honored the company policy. *Id.* at 84. The court conceded, "it is a better practice for a company to make a searching inquiry of all relevant employees to determine whether they violated a company policy regarding use of devices," but "in light of the enormous demands that discovery places on any party" it did not "find that Barclays acted unreasonably." *Id.*

agency law,'" the court decided it was "appropriate to hold Balt responsible" for employee spoliation of evidence on personal devices, lest an employer merely ask an employee to store information on its behalf until after discovery.[239] Even if the employee takes individual action that spoliates evidence, the employer may be liable if it had sufficient control over the information held by an employee–agent.

Even if an employee acted independently and against or in disregard of company interests, the employer may still be liable. In *E.I. du Pont de Nemours and Co. v. Kolon Industries, Inc.*,[240] the court imputed employees' spoliation to the company despite the company's argument that the employees' actions were unauthorized, outside the scope of employment, not taken in furtherance of the company's interests, and in contravention of corporate directives.[241] Rather, the court stated that "[s]tandard principles of agency law govern the attribution of employees' spoliation to the company."[242] Accordingly, the company "breached its duty to preserve when key employees . . . deleted files and email items from their personal computers in the days after DuPont filed the action and after being apprised of their duty to preserve relevant information."[243] Regardless of the employer's "attempt to disassociate itself from the acts of its executives and employees," the court emphasized that employees deleted files from computers used for work.[244] The court declared, "[n]o other conclusion" could be reached, and it would not "entertain suggestions that short of a corporate policy or directive encouraging spoliation . . . , then employees' spoliation of relevant evidence should not be imputed."[245]

---

239. *Microvention*, 2023 WL 7476998, at *26, *28; *see also* Victor Stanley, Inc. v. Creative Pipe, Inc., 269 F.R.D. 497, 515 n.23 (D. Md. 2010) ("[A]gency law is directly applicable to a spoliation motion, and the level of culpability of the agent can be imputed to the master."); Goodman v. Praxair Servs., Inc., 632 F. Supp. 2d 494, 522 n.16 (D. Md. 2009) ("A party may be held responsible for the spoliation of relevant evidence done by its agents."); Valentine v. Mercedes-Benz Credit Corp., No. 98 CIV. 1815 (MBM), 1999 WL 787657, at *5 (S.D.N.Y. 1999) (imputing liability through agents for failure to preserve evidence).

240. 803 F. Supp. 2d 469 (E.D. Va. 2011).

241. *Id.* at 506.

242. *Id.*

243. *Id.* at 500.

244. *Id.* at 506.

245. *Id.* at 507. Liability may attach to individual employees, too. *See In re* Pradaxa (Dabigatran Etexilate) Prod. Liab. Litig., No. 312MD0285DRHSCW, 2013 WL 6486921, at *18 (S.D. Ill. Dec. 9, 2013), *rescinded on other grounds*, 745 F.3d 216, 218–20 (7th Cir. 2014) (warning, "[a]ny employee who refuses to allow the auto delete feature for text messages turned off or to turn over his or her phone for the examination of the relevant space on that phone will be subject to a show cause order of this Court to appear personally in order to demonstrate why he or she should not be held in contempt of Court"); Paisley Park Enterprises, Inc. v. Boxill, 330 F.R.D. 226, 235 (D. Minn. 2019) (". . . the duty to preserve still attached upon both individuals").

Ephemeral messaging presents a new issue for evidence preservation and may be viewed skeptically by a court.[246] In *Fed. Trade Commission v. Noland*,[247] after parties inadvertently learned about an FTC subpoena, "Noland sent an invitation to Harris to install 'Signal,' a mobile messaging application that emphasizes user privacy, and separately sent a message to the 'SBH Leadership Council' . . . stating that he had '[j]ust sent y'all an important invite to an app you need to install.'"[248] Messages were deleted, and Rule 37(e) sanctions followed.

IX. THE SEDONA CONFERENCE POSITION

The Sedona Conference,[249] a research and educational organization, convened a working group to consider possession, custody, or control under Rules 34 and 45 and adopted a recommendation of the legal right test.[250] The Sedona Conference supports the legal right test as "reliable" and "objective."[251] It argues that "control" under Rules 34 and 45 "should be defined as the legal right to obtain Documents and ESI and ability to produce them on demand."[252] It reasoned that the legal right standard is based in clear factors and established legal authorities, affords notice of

246. *See* Jennifer Joyce & Tracey Tran, *What Ephemeral and Third-Party Messaging Apps Mean for eDiscovery*, EY (Oct. 1, 2024), https://www.ey.com/en_us/insights/forensic-integrity-services/third-party-and-ephemeral-messaging-updated-guidelines [https://perma.cc/7C43-D5YN] ("The rise of third-party and ephemeral messaging applications presents significant eDiscovery challenges for companies," including missing evidence, spoliation risk, and comingled data); *see also* Tanya Ganguli & Nilesh Sanwalka, *Disappearing Messages: Updating Best Practices*, INT'L BAR ASS'N (Sep. 13, 2022), https://www.ibanet.org/disappearing-messages-updating-best-practices#_edn4 [https://perma.cc/3L94-VPJF] ("instant messaging technologies which provide an option only to retain data for a short period of time, (usually 'self-destructing' on review of the receiver or within up to 24 hours of message being delivered and referred to as 'ephemeral messaging platforms'), means that managing communication records is becoming increasing [sic] difficult"). Federal regulators are also reining in ephemeral messaging. Press Release, FTC and DOJ, FTC and DOJ Update Guidance That Reinforces Parties' Preservation Obligations for Collaboration Tools and Ephemeral Messaging (https://www.ftc.gov/news-events/news/press-releases/2024/01/ftc-doj-update-guidance-reinforces-parties-preservation-obligations-collaboration-tools-ephemeral [https://perma.cc/N39X-C3NY]) (announcing "updat[ed] language in their standard preservation letters and specifications for all second requests, voluntary access letters, and compulsory legal process, including grand jury subpoenas, to address the increased use of collaboration tools and ephemeral messaging platforms in the modern workplace."). *See generally* The Sedona Conference, *Commentary on Ephemeral Messaging*, 22 SEDONA CONF. J. 435 (2021) (outlining guidelines for regulators, courts, and organizations).

247. No. CV-20-00047-PHX-DWL, 2021 WL 3857413 (D. Ariz. Aug. 30, 2021).

248. *Id.* at *2.

249. THE SEDONA CONF., https://www.thesedonaconference.org/ [https://perma.cc/WU5H-WTFE] (last visited Nov. 27, 2025).

250. *See Rule 34 and 45 Commentary*, *supra* note 15, at 1–3.

251. *Id.* at 64.

252. *Id.*

applicable standards to parties, brings consistency, and takes into account any competing legal consideration (e.g., foreign data privacy laws and corporate structuring).[253] Additionally, the Sedona Conference contends that the legal right test is supported by *Restatement* principles.[254] The Sedona Conference argues that the legal right test "leads to fairer results."[255] By endorsing the legal right test, The Sedona Conference sought to realize a uniform national standard establishing "clear guidance resulting in consistent application."[256] Separately, The Sedona Conference explained that a requesting party can use a Rule 45 subpoena if a Rule 34 responding party does not possess the legal right to obtain information.[257] The Sedona Conference criticizes the practical ability test as "inherently vague" and "unevenly applied," leading to inconsistent results.[258] Yet, as noted above, some plaintiffs and courts note that the legal right test creates perverse incentives for employers to route communications through personal devices to avoid discovery obligations.

Furthermore, The Sedona Conference urges the incorporation of a notification principle, as under the legal right plus notification test. The Sedona Conference recommends that responding parties who lack possession, custody, or control of requested information "should, in a reasonably timely manner," notify the requesting party.[259] The Sedona Conference proposes such notification as a defense against spoliation sanctions or other liability arising out of a third party's failure to preserve information.[260] Notification aims to allow the requesting party the early opportunity to "take steps to attempt to preserve or obtain" documents from third parties.[261]

The Sedona Conference found the practical ability test produces "unfair results."[262] The Sedona Conference similarly emphasizes that organizations should not be compelled to terminate or discipline

253. *Id.* at 77–78.

254. *Id.* at 67.

255. *Id.* at 78.

256. *Id.*

257. *Id.* at 79. Notably, The Sedona Conference highlighted as an "illustrative situation[]/example[]" where there is no Rule 34 control "employer/employee relationships, e.g., employer does not have the legal right to obtain personal Documents and ESI from a director, officer, or employee's personal cell phone, personal email account, or personal social networking sites; employee does not have the legal right to demand or remove data from his/her employer" and "former directors, officers, and employee relationships where no legal right to demand data exists." *Id.* at 80–81.

258. *Rule 34 and 45 Commentary*, *supra* note 15, at 75.

259. *Id.* at 105–06.

260. *Id.*

261. *Id.* at 106.

262. *Id.* at 35.

employees for refusing to surrender devices.[263] It adds that the practical ability test could compromise compliance with foreign privacy laws and give "short shrift" to legally distinct corporate entities.[264] Additionally, The Sedona Conference shared that the practical ability test might force parties to violate their employment contracts and compel document production from current or former employees.[265] The Sedona Conference acknowledges that the practical ability test may limit gamesmanship and allows for judicial discretion.[266] But, despite any potential benefits, The Sedona Conference has concluded that the practical ability test is inherently vague and unevenly applied.[267]

X. PRIVACY CONCERNS ENTER THE FRAY

Various states in the United States, countries in the European Union and around the world, and American courts have recognized that individuals have a substantial privacy interest in their personal cell phones and communications.[268] Some courts have questioned whether the "practical ability" test adequately recognizes that the employment relationship is not level, and some employees may feel compelled to surrender their personal mobile devices when asked to do so.[269] The same observation could be made regarding the application of a "legal right" standard based on company handbooks or BYOD policies, where no bargaining was present and continued employment was conditioned on the acknowledgment or signing of such policies.

Any reconceptualization of "control" also operates against a backdrop of constitutional and statutory limits. Public-sector employers face Fourth Amendment constraints on searches of employees' personal devices, and several states have enacted statutes that restrict employer access to personal social media and electronic accounts. An expansive practical ability regime that effectively compels employers to demand device-level access risks colliding with these protections. An agency-anchored rule

263. *See id.* at 46–47 ("The Practical Ability Standard could arguably put employers in the awkward position of asking for the personal documents and ESI of their employees (and former employee) which may be deemed improper or 'coercive.'").

264. *Id.* at 35.

265. *Id.* at 41, 45.

266. *Id.* at 72.

267. *Id.* at 75; *see also Commentary on BYOD*, *supra* note 98, at 532 ("consider the practical implications of commanding employees to turn over devices . . . . [T]he organization should not be required to use a threat of termination to force the employee to turn over the device. Such a rule would impose too heavily on the relationship between employees and their employer.").

268. *See, e.g.*, Riley v. California, 573 U.S. 373, 394–98 (2014); *In re* Rule 45 Subpoena issued to Nancy Lucas, No. 23-MC-00040, 2023 WL 4362822, at *2 (D. Minn. July 6, 2023 (cell phones hold the privacies of life).

269. *See, e.g.*, *In re* Pork Antitrust Lit., No. 18-cv-1776(JRT/HB), 2022 WL 972401, at *6 (D. Minn. Mar. 31, 2022).

mitigates those conflicts by narrowing the circumstances under which employers are expected to seek access and by encouraging courts to channel most disputes over non-agent devices into the Rule 45 framework, where individualized objections can be heard and resolved.

As employers seek greater insight into and control over employees' devices, employee privacy suffers. While novel technologies may ease employers' information security or eDiscovery difficulties, they expose employees to invasive oversight and surveillance. For example, some new systems allow companies to "observe any imaginable use of a workplace computer."[270] With these mechanisms, employers can record which information is accessed, observe employees' digital activities, and capture employees' screens.[271] As in the rest of society, the "explosion" in technology and "rise of the internet, social media, and connected devices" has changed the discovery system.[272] Meanwhile, "little attention has been paid to privacy as opposed to time and expense" in

270. Andrew M. Jones, *Employee Monitoring: An Overview of Technologies, Treatment, and Best Practices*, STATE BAR OF TEX., https://www.texasbar.com/AM/Template.cfm?Section=articles&Template=/CM/HTMLDisplay.cfm&ContentID=48657 [https://perma.cc/5FWQ-HARF] (last visited Mar. 6, 2026); *see also* David P. Biros, *The Challenges of New Information Technology on Security, Privacy and Ethics*, 2020 J. THE MIDWEST ASS'N FOR INFO. SYS. 2 (2020) (arguing the greatest source of information security vulnerabilities and privacy infringement is the internet of things (IoT) and wearables); *see generally* Madigan Wolford, Note, *Is Your Employer Watching You?: Invasive Employee Surveillance in the Modern Era*, 26 N.C. J.L. & TECH. 617 (2025) (discussing intrusive technologies that allow for companies to monitor employees digitally and which raise privacy concerns); Rafi Bortnick, Note, *Restoring Reasonable Expectations to Privacy at Work in the Face of Modern Electronic Monitoring Practices*, 75 UC L. SF L.J. 1479 (2024) (discussing data collection in the workplace and arguing for greater employee privacy protection); Lothar Determann & Robert Sprague, *Intrusive Monitoring: Employee Privacy Expectations are Reasonable in Europe, Destroyed in the United States*, 26 BERKELEY TECH. L.J. 979 (2011) (setting forth a comparative perspective on employer monitoring and employee privacy between the United States and Europe); Isabela Possino, Note, *Employee Monitoring: As Technology Advances Yet the Electronic Communications Privacy Act Stays in the Past*, 26 SMU SCI. & TECH. L. REV. 135 (2023) (explaining the Electronic Communications Privacy Act of 1986 and arguing it is outdated in light of modern employee monitoring practices).

271. Jones, *supra* note 270.

272. Allyson Haynes Stuart, *A Right to Privacy for Modern Discovery*, 29 GEO. MASON L. REV. 675, 676 (2022); *see also* Gordon v. T.G.R. Logistics, Inc., 321 F.R.D. 401, 403–04 (D. Wyo. 2017) (noting intrusiveness of request for extensive social media information, and the "substantial risk that the fear of humiliation and embarrassment will dissuade injured plaintiffs from seeking recovery for legitimate damages or abandon legitimate claims"); Donald H. Nichols, *Window Peeping in the Workplace: A Look into Employee Privacy in a Technological Era*, 27 WILLIAM MITCHELL L. REV. 1587, 1588 (2001) (providing an outdated but apt illustration: "Most employees who have computers at work often check their e-mail first thing in the morning, about the same time they sort through their postal mail. In both cases, they likely assume that they are the first, and only, person to be reading the contents.").

discovery.[273] Even as the Federal Rules were revised to account for ESI, "they did not give specific attention to issues of privacy."[274] That said, courts have begun to recognize privacy as a factor in the proportionality consideration of discovery.[275] Many courts, especially, recognize the privacy concerns of third parties.[276] In the context of personal mobile devices and ESI, courts "are reluctant to order litigants to submit their cell phones to their opponent for . . . forensic examination absent necessity."[277] In settings where data is being sought from employees residing internationally, the General Data Protection Regulation (GDPR) applies to any processing of personal data carried out by or on behalf of a controller (e.g., the employer), regardless of whether the data is stored on company-owned or personal devices. The parameters of when and how such data can be processed without violating the GDPR remain unclear. For multinational firms, a presumption that all employee devices are within a corporation's "control" can force untenable choices between violating foreign data-protection law and risking sanctions at home. Limiting presumptive control to a discrete group of agents and relying on case-specific analysis or Rule 45 mechanisms for others reduces the frequency and breadth of such conflicts, even if it cannot eliminate them altogether. Sometimes, foreign employers will insist upon seeking data by way of the Hague Convention on the Taking of Evidence, rather than compliance with the Federal Rules.[278] In the United States it has been reported that twenty-eight states have enacted some form of restrictions on the ability of employers to access employee social media and mobile device contents.[279]

These privacy concerns are not incidental. In an era when a single device may hold banking information, health data, intimate communications, and location history alongside work emails, any test for "control" that routinely pushes employers to demand device-level access

---

273. Stuart, *supra* note 272, at 677. While "There has long been the idea" of privacy from government intrusion, "there is no such concept in civil discovery." *Id.*; *cf.* Matthew T. Bodie, *The Law of Employee Data: Privacy, Property, Governance*, 97 IND. L.J. 707, 754 (2022) ("The existing patchwork of privacy laws provides at best uncertain relief . . . . The law has encased employer power over employee data.").

274. Stuart, *supra* note 272, at 705.

275. *Id.*; *see generally* Robert D. Keeling & Ray Mangum, *The Burden of Privacy in Discovery*, 105 JUDICATURE 66 (2021) (arguing that privacy can be part of proportionality analysis for discovery).

276. Stuart, *supra* note 272, at 701.

277. *Id.* at 706–07.

278. *See, e.g.*, Société Nationale Industrielle Aérospatiale v. United States Dist. Court for S. Dist. of Iowa, 482 U.S. 522, 540 (1987).

279. *See* Douglas M. Oldham, *The State of Employment Law - 28 States Limit Employer Access to Private Social Media Accounts*, NAT'L L. REV. (Oct. 16, 2025), https://natlawreview.com/article/state-employment-law-28-states-limit-employer-access-private-social-media-accounts [https://perma.cc/3CUW-KQZN].

risks sweeping far beyond what proportionality can justify. The question is not whether courts care about privacy—they increasingly do—but whether the prevailing tests give them usable tools to calibrate discovery obligations in light of these interests.

XI. AGENCY PRINCIPLES: A WAY FORWARD

None of the existing possession, custody, or control tests balance the important interests just right. This author agrees with the Sedona Conference, in part, wherein it properly acknowledges that the various prevailing tests have resulted in inconsistent application, uncertainty, and inequitable or even futile results.

The current tests invite confusion from the start of the discovery spectrum. Corporations must place a litigation hold on relevant data that it knows or reasonably should know will be relevant. Without a clearer standard, companies are left to speculate as to whether their identification and preservation efforts have been sufficient. Likewise, requesting parties are left to guess whether to issue Rule 45 subpoenas or Rule 34 requests for production. In the interim, relevant and unique data could be lost because of the passage of time, replacement of devices, or deletion settings.

Secondly, examination of policies and contracts for the establishment of the legal right standard invites clever drafting (either to mollify employee concerns of privacy invasion[280] or outright abuse). An agency-anchored standard offers one way to internalize privacy costs. By limiting presumptive control to a defined subset of agents and directing parties toward Rule 45 for other employees, it reduces the universe of individuals whose personal devices may be drawn into party-to-party discovery while still preserving a path to obtain critical information where truly necessary. Courts retain tools—protective orders, tailoring of search protocols, and cost-shifting—to further mitigate privacy intrusions even within the presumptive category.

Third, the burden of establishing control is on the requesting party.[281] Often, a requesting party is not in possession of enough information to establish this showing. Rather than effectively encouraging litigation disputes on the issue of control, courts should adopt a new presumptive test based on principles of agency law. Under it, employers would control relevant data on the personal mobile devices of company officers,

280. *See, e.g.*, *In re* Pork Antitrust Litig., No. 18-CV-1776 (JRT/HB), 2022 WL 972401, at *5 (D. Minn. Mar. 31, 2022); Allergan, Inc. v. Revance Therapeutics, Inc., No. 3:23-CV-00431, 2025 WL 984792, at *1 (M.D. Tenn. Mar. 17, 2025), report and recommendation adopted, No. 3:23-CV-00431, 2025 WL 1020485 (M.D. Tenn. Apr. 4, 2025).

281. *See, e.g.*, Mamani v. De Lozada Sanchez Bustamante, No. 07-22459-CIV, 2017 WL 3456327, at *2 (S.D. Fla. Aug. 11, 2017) ("Defendants bear the burden of establishing that the documents are under [a party's] possession, custody, or control.").

directors, supervisors, and managers, and relevant data would be subject to preservation and potentially production.[282] This also comports with requirements of Rule 34 requests sent to a corporate party and Rule 30(b)(6) notices sent to a corporate representative. Under those rules a party cannot simply respond based on personal, immediate knowledge, but is under an affirmative duty to seek information reasonably available to it.[283] Under the agency test, non-managerial employees' personal mobile devices would not be subject to preservation and production obligations under Rule 34. Requesting parties would need to issue Rule 45 subpoenas for them. That allocation is not costless; Rule 45 subpoenas impose additional burdens on individual plaintiffs and can be prohibitively expensive in small-value cases. To the extent those burdens fall disproportionately on low-income plaintiffs, courts should use existing tools—such as staged discovery, targeted subpoenas, and fee-shifting—to ensure the agency-anchored standard does not become a de facto bar to proving claims. The agency test would appropriately balance modern realities, legal principles, trial efficiency, data preservation, and employee privacy.

Rejecting the practical ability test would also recognize that ordering a party to produce documents from regular employees or former employees could prove futile, as there is no enforceable means of obtaining them.[284]

While a formal agency test for control would be somewhat novel, the law underlying it would not be. Rather, agency law is well established, and courts have consistently invoked its principles across vast areas of the law. At bottom, an agent is one who acts on behalf of a principal and

---

282. Some courts have already applied this test. "If a party has the right, and the ready ability, to obtain copies of documents gathered or created by its agents pursuant to work done for the party, such documents are clearly within the [party's] control." McAllister-Lewis v. Goodyear Dunlop Tires N. Am., Ltd., No. 4:14-CV-04103-LLP, 2015 WL 5794697, at *4 (D.S.D. Oct. 1, 2015); Waymo LLC v. Uber Techs., Inc., No. 17-CV-00939-WHA(JSC), 2017 WL 2972806, at *2 (N.D. Cal. July 12, 2017) ("The cases Otto Trucking does cite are all inapposite; none involve a corporation refusing to produce documents involving corporation business in the possession, custody or control of the corporation's officers. To accept Otto Trucking's argument would mean that it could not compel its current officers to produce Otto Trucking's own trade secrets to Otto Trucking merely because the officers conducted Otto Trucking business with their personal email accounts. Nonsense. Otto Trucking must produce responsive documents in the custody, control or possession of its officers and agents."); NDN Collective v. Retsel Corp., No. 5:22-CV-05027-LLP, 2024 WL 3401052, at *7 (D.S.D. July 12, 2024) (applying an agency test without acknowledging the term agent—"Defendants have the legal right to obtain from its board members documents regarding Defendants' business.").

283. *See, e.g.*, Hayse v. City of Melvindale, No. CV 17-13294, 2018 WL 3655138, at *6 (E.D. Mich. Aug. 2, 2018).

284. *See In re* Citric Acid Litig., 191 F.3d 1090, 1108 (9th Cir. 1999) (but also including supervisor's cell phone as not within the company's control).

is subject to the principal's control.[285] The principal–agent relationship is mutual and bidirectional, as the agent is a representative of the principal with power to act on its behalf, while the principal may control the agent's actions.[286] In *Proofpoint, Inc. v. Vade Secure, Inc.*,[287] the court found that a company had control over the laptop of the Chief Technology Officer, who was the employer's agent.[288] Considering whether to impose sanctions for spoliation on the company, the court rejected the employer's argument that it lacked control over the personal laptop from which the employee had deleted data.[289] The court found "this distinction—between Vade, the employer, and Lemarié, Vade's Chief Technology Officer, to be immaterial in the present context."[290] The court stated, "Contrary to Vade's self-serving contentions that it had no control over its *own* source code simply because it was reposed on the personal computers of its *own* agent," it did not "believe the suggestion that Vade was unable to direct its agent."[291] If the employer could not direct its agent, "then any employer or principal in any case could always avoid producing" any evidence "by simply asking an employee or agent to hold it for them until after the discovery cut-off date."[292] As the court acknowledged in

285. RESTATEMENT (THIRD) OF AGENCY § 1.01 (2006).

286. *Id.* at § 1.01(c); *see also* Deborah A. DeMott, *The Contours and Composition of Agency Doctrine: Perspectives from History and Theory on Inherent Agency Power*, 2014 U. ILL. L. REV. 1813, 1833 (2014) (An agent effectuates a "legally-salient extension of a principal's personality.").

287. No. 19-CV-04238, 2021 WL 11108111 (N.D. Cal. June 1, 2021).

288. *Id.* at *5; *see also* DR Distribs., LLC v. 21 Century Smoking, Inc., 513 F. Supp. 3d 839, 893 (N.D. Ill. 2021) ("At the very least, Edmiston was a volunteer agent for Defendants. So, Edmiston's documents [we]re in the possession, custody, or control of Defendants.") (citation omitted); Allen v. Woodford, No. CV-F-05-1104 OWW LJO, 2007 WL 309945, at *2 (E.D. Cal. Jan. 30, 2007) ("'Control' may be established by the existence of a principal–agent relationship."); Shim-Larkin v. City of New York, 16-CV-6099 (AJN)(KNF), 2019 WL 5198792, at *10 (S.D.N.Y. Sep. 16, 2019) ("[T]he defendant, as Kravitz's employer, had control over the contents of Kravitz's text messages . . . that Kravitz stored on his personal cellular telephone and which are pertinent to this action.").

289. *Proofpoint*, 2021 WL 11108111, at *5.

290. *Id.* Under Texas law, there is a designation of "vice principal." It encompasses:
(a) corporate officers; (b) those who have authority to employ, direct, and discharge servants of the master; (c) those engaged in the performance of nondelegable or absolute duties of the master; and (d) those to whom the master has confided the management of the whole or a department or a division of the business. Mobil Oil Corp. v. Ellender, 968 S.W.2d 917, 922 (Tex. 1998).

291. *Proofpoint*, 2021 WL 11108111, at *5 (emphasis in original).

292. *Id.*; *cf.* Schnatter v. Papa John's Int'l, Inc., No. 2018-0542-AGB, 2019 WL 194634, at *16 (Del. Ch. Jan. 15, 2019), *abrogated on other grounds by* Tiger v. Boast Apparel, Inc., 214 A.3d 933 (Del. 2019) ("The reality of today's world is that people communicate in many more ways than ever before . . . . In my view . . . if the custodians identified here—the Company's other directors, CEO, and General Counsel—used personal accounts and devices to communicate . . . they should expect to provide that information to the Company.").

McBryar v. International Union,[293] "every organization . . . can operate only through the actions of its officers and agents;" therefore, "it can possess, control, or have custody of documents and things only if its officers and agents possess, control, or have custody of them." [294]

The choice to focus on officers, directors, and managers is not arbitrary. As cases like *Proofpoint* and *McBryar* illustrate, courts already rely on agency concepts to attribute knowledge and conduct of high-level actors to the organization for purposes of sanctions and liability. Those actors are entrusted with discretionary authority to bind the firm and routinely act as its voice in internal and external communications; their devices are more likely to contain uniquely probative work-related ESI that cannot practically be replicated elsewhere. By contrast, rank-and-file employees may possess relevant information, but their communications are less systematically tied to corporate decision-making and more likely to be duplicative of other sources.

Of course, titles do not perfectly track authority. Any workable rule must leave room for courts to look past nominal labels and consider whether a given employee actually exercises supervisory or decision-making functions that justify treating their device as within the firm's control. The agency-anchored test therefore operates as a presumption tied to function, not a rigid taxonomy of job descriptions.

CONCLUSION

The meaning of "possession, custody, or control" has become increasingly unstable as technology outpaces the assumptions that underwrote the 1938 Rules. This Article has argued that the prevailing legal right and practical ability tests are ill-suited to the hybrid, privacy-sensitive world of employee-owned mobile devices, and that an agency-anchored conception of control offers a more coherent and administrable starting point. Whether implemented through formal amendment of Rules 34 and 45, through Advisory Committee guidance, or through gradual doctrinal convergence, moving toward such a standard would advance the Rules' core goals of fairness, predictability, and proportionality in discovery.

The agency test, according to which employers have control over the personal devices of company officers, directors, managers, and supervisors who act as the corporation's agents, may provide a way

---

293. McBryar v. Int'l Union of Auto. Aerospace & Agric. Implement Workers, 160 F.R.D. 691 (S.D. Ind. 1993).

294. *Id.* at 695; *cf.* Riddell Sports, Inc. v. Brooks, 158 F.R.D. 555, 558 (S.D.N.Y. Nov. 8, 1994) ("'[A]n individual party to a lawsuit can be compelled to produce relevant information . . . relating to a non-party corporation of which it is an officer . . . .' This case presents the mirror-image issue.").

forward to achieving an appropriate balance.[295] The proposed approach lessens the rigidity and under-inclusivity of the legal right test and the malleability and over-intrusiveness of the practical ability test. It is time for a uniform national standard.

295. The author concedes that an agency test may not resolve all disputes. Controversy has always existed as to whether some employees were managers, supervisors, or "rank and file" workers. Job titles and structures vary widely. It can also be asked why do first-line supervisors fall on the "control" side but highly autonomous non-managerial specialists (e.g., lead engineers) do not? Further, whether a corporation can be required to acquire data in the possession of a Software as a Service (SaaS) vendor remains unaddressed by this agency test. *See* Brown v. Tellermate Holdings Ltd., No. 2:11-CV-1122, 2014 WL 2987051, at *4 (S.D. Ohio July 1, 2014), *adopted as modified*, No. 2:11-CV-1122, 2015 WL 4742686 (S.D. Ohio Aug. 11, 2015) (salesforce.com information).

FORCE MULTIPLIER?: ARTIFICIAL INTELLIGENCE, UNEVEN COMPETENCE, AND THE INTEGRITY OF THE ADVERSARIAL SYSTEM

Ralph Artigliere,[*] *David Horrigan*,[**] & *Rose Hunter Jones*[***]

Abstract

Generative artificial intelligence has entered legal practice with unprecedented speed, transforming how lawyers and judges synthesize information, develop arguments, and evaluate evidence. While early attention has focused on visible failures such as hallucinated citations and fabricated authority, these incidents do not capture the deeper structural risk AI poses to the legal system. The authors argue that generative AI is a stress test for legal institutions: it offers substantial opportunities to improve legal practice and access to justice, but it also operates as a force multiplier that can deepen structural asymmetries when professional competence and oversight are uneven. Unlike prior legal technologies that primarily assisted retrieval or organization, generative systems participate directly in synthesis, framing, and analysis, altering how legal knowledge is formed and contested. Where AI competence is uneven, disparities in analytical capacity, litigation leverage, and epistemic control emerge in ways existing doctrine was not designed to absorb.

Drawing on the historical evolution of electronic discovery, Technology-Assisted Review, and digital evidence, this Article demonstrates why prior models of incremental adaptation are no longer sufficient. Ethical doctrine and sanctions regimes can identify individual failures, but they operate reactively and cannot correct structural asymmetries. Empirical evidence suggests, however, that the competence divide is not fixed: with appropriate education, supervision, and institutional design, generative AI can narrow disparities rather than entrench them. This Article concludes by proposing a calibrated institutional roadmap focused on formation, oversight, incentives, and coordination to ensure that AI advances the core commitments of fairness, accuracy, and legitimacy in an adversarial system increasingly shaped by probabilistic tools.

---

* Circuit Judge, Florida Tenth Judicial Circuit (Ret.).

** Discovery Counsel and Legal Education Director, Relativity, and Adjunct Professor of Law, Thomas R. Kline School of Law of Duquesne University.

*** Partner, Hilgers PLLC.

The authors acknowledge the encouragement and assistance of UF Law Professor William F. Hamilton in the preparation of this Article and dedicate this Article to his own selfless and productive contributions to law and technology education. All views and any errors remain the authors' alone.

INTRODUCTION

The justice system depends on the ability of lawyers and judges to understand, evaluate, and responsibly use information, and on the assumption that this ability is distributed evenly enough to sustain adversarial testing. When that ability is uneven across the profession, fairness and accuracy suffer. Over the past quarter-century, legal practice has undergone successive technological transformations, moving from digitized research and electronic document management to algorithmic tools in discovery, and now to an era increasingly defined by generative artificial intelligence. Each wave required adaptation by the profession. What distinguishes the current moment is not simply the pace of innovation, but the consequences of uneven adaptation.

For decades, technological competence in law functioned primarily as an advantage rather than a prerequisite. Lawyers who mastered electronic research tools or Technology-Assisted Review gained efficiency and cost benefits, and courts gradually incorporated those technologies into accepted practice. Yet the profession has long experienced a widening divide between those who develop technological competence and those who do not.[1]

Generative AI has transformed that divide from a professional inconvenience into a systemic risk. For purposes of this Article, "artificial intelligence" or "AI" refers to generative artificial intelligence systems,[2] including large language models,[3] that produce text, analysis, or other substantive outputs through probabilistic generation rather than through independent verification of information.[4] Its accessibility, opacity, and persuasive fluency compress the timeline for adaptation while

---

1. This divide did not go unnoticed and prompted ethical rules and guidance addressing lagging technological competence. *See, e.g.*, MODEL RULES OF PRO. CONDUCT r. 1.1 cmt. 8 (A.B.A.) (2012) (requiring lawyers to keep abreast of relevant technology); R. Regulating Fla. Bar r. 4-1.1 cmt. (2016) (describing the requirements of attorney competency); Cal. State Bar Formal Op. No. 2015-193 (explaining that a lack of technological knowledge may render a lawyer incompetent to handle matters involving e-discovery); Fla. Ethics Op. 10-2 (addressing lawyer competence and supervision in the use of cloud-based and electronic information systems).

2. *See generally* Nat'l Inst. of Standards & Tech., *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (Jan. 2023), https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf [https://perma.cc/C5WH-HZ96]; Nat'l Inst. of Standards & Tech., *CSRC Glossary*, https://csrc.nist.rip/glossary/term/generative_artificial_intelligence [https://perma.cc/9W63-A3ZC] (last visited Feb. 17, 2026) (defining "generative artificial intelligence" as "[t]he class of AI models that emulate the structure and characteristics of input data in order to generate derived synthetic content[,]" including "text, and other digital content").

3. A large language model (LLM) is "a narrow artificial intelligence (AI) system that has been trained on a massive amount of text data to interpret natural language and generate human-like responses to text-based prompts or questions." Bassel Almarie et al., *The Use of Large Language Models in Science: Opportunities and Challenges*, 9 PRINC. PRACT. CLIN. RES. 1, 1 (2023).

4. *See id.*

magnifying the consequences of misuse. Unlike prior technologies, AI does not merely accelerate work; it produces plausible but potentially false authority, compresses verification time, and obscures provenance. Its most significant risks extend beyond fabricated citations and synthetic evidence to subtle distortions in how facts are summarized, evidence is synthesized, and expert judgment is framed. AI is neither an apocalypse nor a miracle. It functions as a multiplier that amplifies skill where skill exists and magnifies error where oversight fails or verification is absent.

This Article argues that artificial intelligence functions as a stress test for the legal profession because it magnifies an existing competence divide in technological knowledge and discipline.[5] That divide separates legal professionals who understand the capabilities and limits of AI well enough to use it responsibly from those who do not. Its effects extend beyond lawyers to judges, clients, and pro se litigants, manifesting in structural unfairness in how evidence is developed, arguments are framed, and outcomes are reached. Unless lawyers and judges develop the ability to use these tools with appropriate verification and oversight, misuse, uneven access, and institutional inertia will undermine procedural fairness and public confidence in the justice system.[6] The competence divide exposed by generative AI also carries direct implications for access to justice, affecting not only sophisticated litigants and well-resourced institutions, but also the courts that must manage them, the small firms that cannot match their resources, and the self-represented parties navigating a system increasingly shaped by tools they cannot access or employ in practice.

This Article first traces how earlier waves of digital technology exposed competence divides in discovery, legal workflow, and evidence. It then explains why AI amplifies those divides. Finally, it proposes a practical roadmap for courts, law schools, firms, and individual lawyers to close the gap. The profession can meet this moment, but only by acting with purpose and urgency. The pace of technological change is not merely relentless; it is accelerating.

Drawing on the authors' combined experience as a trial judge and judicial educator, a discovery counsel and legal technology commentator, and a litigator with an engineering background, this Article examines how the legal profession has historically adapted unevenly to

5. *See generally* Harry Surden, *Artificial Intelligence and Law: An Overview*, 35 GA. ST. U. L. REV. 1305, 1312–15 (2019) (explaining that contemporary AI systems function by detecting and operationalizing patterns at scale, producing intelligent-seeming outputs that reshape how decision-making tasks are performed even where humans remain formally responsible).

6. *See* Hon. John M. Facciola, *Discovery: Faster and Shorter*, 2005 FED. CTS. L. REV. 7, 9–10 (reviewing Hon. Paul W. Grimm, Charles S. Fax & Paul Mark Sandler, *Discovery Problems and Their Solutions* (A.B.A. 2005)) (endorsing the view that discovery competence is central to effective advocacy and warning that lawyers who do not understand how to conduct or manage discovery pose systemic risks comparable to lawyers who lack core litigation skills).

technological change.[7] Across judicial service, complex litigation, systems-oriented practice, and exposure to crisis-driven institutional disruption, the authors observed a consistent pattern: technological competence in law has rarely been anticipatory. Instead, it has developed reactively, under pressure, through specialization, delegation, and institutional lag. The profession ultimately adapted to each prior technological shift, but only after friction, uneven learning, and avoidable cost. AI compresses that cycle. Tools that now participate directly in synthesis, reasoning, and judgment leave far less margin for uneven understanding or delayed adaptation. The question is no longer whether the profession will adjust to technological change, but whether it can do so deliberately—before error, asymmetry, and inequity harden into structural features of the legal system.

I. THE EARLY DIGITAL ERA: INNOVATION AHEAD OF DOCTRINE

A. *Digital Evidence Before Digital Doctrine*

Generative AI is not the first technological disruption to strain procedural doctrine. Technology has repeatedly entered legal practice before doctrine was ready, requiring courts to adapt reactively through litigation, expert testimony, and case-by-case decision-making. Rulemakers followed practice rather than anticipating it. Amendments to Rule 34, the formal recognition of electronically stored information, and later reforms to Rule 37(e) all occurred after digital evidence was already embedded in litigation.[8]

Before electronically stored information was recognized as a discrete category within the Federal Rules, courts confronting technologically complex evidence operated without a framework tailored to digital data. Judges and lawyers often lacked technical understanding and frequently depended on expert explanations to interpret complex digital evidence. Uneven competence was tolerated as a practical necessity because prior

---

7. Some observations and evaluative judgments in this Article draw on the authors' professional experience in technology, electronic discovery, and litigation practice and support. That experience includes sustained involvement in eDiscovery and technology education, research, writing, and rulemaking from the early 2000s to the present, including participation in The Sedona Conference Working Groups 1 and 13, the University of Florida eDiscovery Conference, and continuing legal and judicial education programs nationwide.

8. *See* FED. R. CIV. P. 34 advisory committee's note to 1970 amendment (clarifying that "data compilations" are discoverable and providing early doctrinal accommodation for electronically stored information before its formal recognition); FED. R. CIV. P. 34 advisory committee's note to 2006 amendment (acknowledging the dramatic growth in electronically stored information and the diversity of systems used to create and store it); FED. R. CIV. P. 37(e) advisory committee's note to 2015 amendment (responding to problems arising from the exponential growth of electronically stored information and the inconsistent sanctions regimes that developed through case law).

technologies assisted retrieval and analysis without displacing professional judgment or authorship.[9] As doctrine lagged behind technological capability, courts and litigants were required to improvise within existing rules and to develop sufficient technological literacy to retain, consult, and critically evaluate experts. Access to expertise, tools, and training was never evenly distributed, yet the system functioned. That balance would later prove fragile. Technology-Assisted Review arguably marked the first moment when technology altered not only efficiency, but the sequence in which information became visible, foreshadowing deeper epistemic divides.[10]

Courts were required to determine what constituted discoverable and admissible digital material long before the term "electronically stored information" entered the Federal Rules. Although the Federal Rules of Civil Procedure were not drafted with digital evidence in mind, the 1970 amendment to Rule 34 clarified that "data compilations" were discoverable, laying early groundwork for later doctrine.[11] The term "e-discovery" emerged only decades later, as digital data became the primary and then overwhelming source of relevant evidence in litigation. The lesson of this period is clear: doctrine absorbs technological change, but only after prolonged friction, uneven learning, and institutional lag.

---

9. Tolerance in this context did not connote uniform competence or effective mastery. During the emergence of electronic discovery, lawyers and judges experienced significant asymmetry in technological understanding and skill. *See* The Sedona Conference, *The Sedona Principles, Third Edition: Best Practices, Recommendations & Principles for Addressing Electronic Document Production*, 19 SEDONA CONF. J. 1, 179 (2018) (discussing the role of technological competence among counsel in electronic discovery). In the federal system, magistrate judges and a subset of district judges necessarily developed deeper expertise as they handled the bulk of discovery disputes, while many trial lawyers and Article III judges encountered eDiscovery issues less frequently and therefore learned more slowly. State courts faced an even steeper challenge, given broader dockets and fewer opportunities for sustained engagement with complex digital discovery. Although rule amendments and professional education eventually responded to the dominance of digital evidence, that process unfolded over many years. *See* FED. R. CIV. P. 34 advisory committee's note to 2006 amendment (addressing the dramatic growth in electronically stored information and in the variety of systems for creating and storing such information); FED. R. CIV. P. 37(e) advisory committee's note to 2015 amendment (addressing the serious problems resulting from the continued exponential growth in the volume of such information). The resulting competence divide narrowed unevenly but was never fully eliminated, a pattern that bears directly on the risks posed by generative AI.

10. *See* Da Silva Moore v. Publicis Groupe, 287 F.R.D. 182, 185, 193 (S.D.N.Y. 2012) (approving the use of a form of Technology-Assisted Review (TAR) in appropriate cases for the first time in a published federal decision, sometimes described in early case law and practitioner literature as "predictive coding"); Rio Tinto PLC v. Vale S.A., 306 F.R.D. 125, 127 (S.D.N.Y. 2015) (recognizing TAR as an accepted and often effective means of limiting discovery to relevant electronically stored information without undue burden). These decisions marked an inflection point in discovery practice by recognizing that AI-assisted technologies could be reliable, defensible, and appropriate tools for managing large-scale electronically stored information.

11. *See supra* note 8.

By the mid-1990s, email had become a primary means of communication for lawyers and judges, forcing courts to confront confidentiality, metadata, and preservation issues before procedural doctrine had fully adapted. The 2006 amendments to the Federal Rules of Civil Procedure formally recognized electronically stored information, and decisions such as *Zubulake v. UBS Warburg LLC*[12] clarified preservation duties, including duties of counsel. These developments illustrate the familiar pattern of technological adoption preceding doctrinal clarity. Yet the profession again adapted through litigation, expert testimony, and judicial education. What initially appeared disruptive became normalized because email and similar digital tools, while transformative, remained bounded. They altered how evidence was stored and reviewed without replacing professional judgment, factual synthesis, or legal authorship.

B. *Overcoming the Limitations of Paper-Based Practice with Technology*

The attacks of September 11, 2001, and subsequent natural disasters exposed the vulnerability of paper-based legal practice. Law offices lost files, pleadings, and archived evidence at a scale that exposed the absence of reliable disaster recovery systems. In the years that followed, cloud computing, electronic filing, and distributed digital storage expanded rapidly, transforming legal infrastructure. These technologies improved resilience rather than destabilized doctrine. They altered how legal information was stored and accessed without directly displacing professional judgment, evidentiary standards, or legal authorship. The profession's adaptation was uneven but largely constructive. This period reinforces a recurring lesson: doctrine does not anticipate technological change but absorbs it. But absorption requires time, education, and disciplined judicial engagement.

In this sense, the profession has previously responded to technological disruption reactively, often after crisis, and has successfully used technology to restore stability, but only because those technologies were infrastructural and bounded. The question now is whether that deliberate

12. *See generally* Zubulake v. UBS Warburg LLC, 216 F.R.D. 280 (S.D.N.Y. 2003) (*Zubulake III*); Zubulake v. UBS Warburg LLC, 220 F.R.D. 212 (S.D.N.Y. 2003) (*Zubulake IV*); Zubulake v. UBS Warburg LLC, 229 F.R.D. 422 (S.D.N.Y. 2004) (*Zubulake V*) (collectively articulating counsel's duty to identify, preserve, and supervise the handling of electronically stored information in light of the unique characteristics of digital data and information systems); *see also* The Sedona Principles, *supra* note 9; FED. R. CIV. P. 37(e) advisory committee's note to 2015 amendment (reflecting a shift from case-driven sanctions doctrine toward a rule-based framework). *Cf.* Barbera v. Grailed, Inc., No. 24-cv-3535 (LJL), 2025 WL 2098635, at *30 (S.D.N.Y. July 25, 2025) (applying post-2015 sanctions analysis and departing from earlier case-law regimes).

process remains viable as technological acceleration compresses the window for adaptation in the age of generative AI.

II. ALGORITHMIC DEPLOYMENT AND THE LIMITS OF INCREMENTAL ADAPTATION

The institutionalization of electronic discovery obligations accelerated in the early 2000s through a combination of sanctions decisions, judicial guidance, and rule amendments. Decisions such as *Zubulake v. UBS Warburg LLC*[13] clarified preservation duties for electronically stored information and emphasized the need for proactive identification and litigation holds. The Sedona Principles[14] further articulated best practices for defensible discovery conduct, and the 2006 amendments to the Federal Rules of Civil Procedure[15] formally incorporated "electronically stored information" into discovery doctrine. Subsequent amendments, including the 2015 revision to Rule 37(e), recalibrated sanctions standards to reflect evolving technological realities.[16]

These developments illustrate a recurring pattern in the profession's response to technological change. Competence was built incrementally, often in reaction to costly litigation failures rather than through anticipatory reform. Doctrine followed practice. Judges, lawyers, and institutions learned through experience, guided by reasonableness standards rather than technical mastery. That model of adaptation proved workable when technology assisted human judgment without altering authorship or analytical control.

As the volume of electronically stored information expanded exponentially, traditional linear review methods became economically unsustainable in large-scale litigation. Technology-Assisted Review (TAR) emerged as a defensible alternative. Courts approved its use within the proportionality framework of Rule 26, emphasizing transparency, sampling, validation, and cooperation. Judicial acceptance of TAR did not reflect uncritical confidence in algorithms. Rather, it reflected recognition that statistically validated methods could satisfy discovery obligations more reliably than exhaustive manual review.

Importantly, TAR remained bounded. It classified and ranked documents to assist human reviewers. It did not generate legal argument,

---

13. *See generally Zubulake*, 216 F.R.D. 280 (S.D.N.Y. 2003) (*Zubulake III*); *Zubulake*, 220 F.R.D. 212, 216 (S.D.N.Y. 2003) (*Zubulake IV*); *Zubulake*, 229 F.R.D. 422 (S.D.N.Y. 2004) (*Zubulake V*).

14. *See generally* The Sedona Conference, *supra* note 9.

15. Amendments to the Federal Rules of Civil Procedure, 2006 U.S.C.C.A.N. 699.

16. *See* Amendments to the Federal Rules of Civil Procedure, 2015 U.S.C.C.A.N. 131 (adopting Rule 37(e) to address preservation failures and to replace divergent sanctions standards that had developed through case law).

synthesize evidence, or produce authority. Human judgment remained central, and methodological transparency allowed courts to evaluate reasonableness without requiring technological parity between parties. Oversight shifted from completeness to defensibility. That shift proved manageable because TAR's failure modes were observable, testable, and correctable.

TAR nonetheless marked a subtle inflection point. It introduced the first epistemic asymmetry in discovery practice. Keyword searching, even when imperfect, was universally accessible. Ranking systems were not. Validated models allowed parties to see the most relevant information earlier, refine strategy more quickly, and reduce costs dramatically. Control over informational sequence affected litigation posture long before full review was complete. The advantage was not merely economic; it shaped how knowledge entered the case.

Courts absorbed that asymmetry by emphasizing transparency and methodological rigor rather than equal access to technology. The system adapted, but the precedent was established: technological architecture could shape litigation capacity in ways doctrine would tolerate, so long as human judgment remained primary and validation remained visible.

Generative AI extends this asymmetry into a fundamentally different domain. The disparity is no longer confined to document prioritization. It reaches authorship, factual synthesis, expert framing, and legal reasoning itself. What was once a sequence advantage has become a cognitive one. Unlike TAR, which is limited to classification or ranking in aid of human review, generative AI is not so constrained. It produces persuasive outputs that can obscure errors, compress verification time, and blur the distinction between assistance and authorship.

Unlike earlier discovery technologies, AI operates with reduced transparency. Its training data, internal weighting, and prompt-dependent behavior are opaque to end users, and its outputs lack standardized, court-recognized validation metrics. Verification therefore becomes more dependent on judgment and more time sensitive. Compounding this risk, AI tools are widely available and deceptively easy to use, inviting reliance by practitioners who may lack the expertise necessary to detect distortion rather than outright error.

The lesson of the TAR era is therefore instructive—but incomplete. Incremental adaptation succeeded when technology supported human judgment within visible constraints. AI removes many of those constraints. As tools evolve from assisting review to shaping narrative and analysis, the margin for uneven competence narrows and the consequences of error escalate. The profession can no longer rely on reactive absorption alone.

A critical but often underexamined reason for this breakdown lies in the profession's limited engagement with systems thinking. Generative

AI operates within probabilistic, multi-stage architectures that require attention to training inputs, validation, feedback loops, failure modes, and lifecycle governance. These concepts are familiar in engineering and other technical disciplines but have historically played a limited role in legal education and professional formation. As a result, many lawyers and judges approach AI outputs as discrete tools rather than as components of dynamic systems whose reliability depends on design, supervision, and context. The resulting mismatch between technological architecture and professional intuition magnifies error and obscures accountability in ways prior discovery technologies did not.

These limitations mark the point at which historical patterns of technological adaptation cease to be reliable guides. The implications of this shift extend well beyond discovery practice. They bear directly on how legal knowledge is formed, contested, and evaluated in an AI-mediated system. Part III examines how the resulting competence gap manifests in practice, not as isolated error, but as a systemic asymmetry shaping how legal knowledge is generated, tested, and presented.

III. EARLY SIGNALS OF A WIDENING DIVIDE FOLLOWING PUBLIC AI DEPLOYMENT

The pace and manner in which artificial intelligence entered legal practice underscores why the competence divide now presents heightened and immediate risk. Although research in artificial intelligence spans decades, the public release of large-scale generative AI tools in late 2022 marked a qualitative shift in accessibility. With the public availability of OpenAI's ChatGPT in November 2022, sophisticated generative capabilities became instantly usable by lawyers, judges, students, and litigants without technical training, institutional gatekeeping, or formal supervision. What previously required specialized infrastructure or vendor mediation could now be deployed by individual practitioners in ordinary legal workflows. Within months, distinct and diverging patterns of use emerged that presaged a widening divide across the legal profession.

On one side of that divide, lawyers, firms, legal technologists, and court systems that invested time and effort in understanding the tools began developing increasingly sophisticated and responsible applications. Early adopters identified productive uses across legal workflows, including drafting and revising correspondence and pleadings, summarizing records and deposition transcripts, organizing large document sets, assisting with discovery search and management, supporting legal research and issue framing, managing expert witness materials, and improving internal knowledge management. As competing models and platforms emerged from major technology providers and generative capabilities were integrated into legal research and practice

tools, those with the capacity to evaluate, test, and supervise these systems began leveraging them to improve efficiency, analytical rigor, and institutional function. For these users, AI operated as a force multiplier for professional judgment rather than a substitute for it.

At the same time, a contrasting pattern emerged. Lawyers unfamiliar with the limitations of generative systems began using AI for legal research and drafting without adequate verification or supervision, leading to filings containing fabricated citations, distorted summaries, and unsupported analysis.[17] Highly publicized incidents revealed not intentional misconduct, but misplaced confidence in tools whose probabilistic nature and failure modes were poorly understood. Institutional responses followed quickly but unevenly.[18] Courts, ethics bodies, and law firms sought to manage perceived risk through sanctions,[19] disclosure mandates, restrictions, and, in some instances, categorical prohibitions on AI use.[20] These reactions were understandable given the novelty of the technology and the seriousness of the errors, but they did not eliminate improper use or ensure responsible adoption.[21]

These early patterns reveal more than isolated misuse or transitional error. They expose a structural divergence in how the profession is absorbing AI—between actors with the resources, permission, and institutional support to develop competence, and those constrained by prohibition, fear, or lack of training. That divergence emerged rapidly, not because AI is uniquely dangerous, but because it entered practice without the buffers that previously allowed uneven competence to remain tolerable. Part IV examines why this divergence is qualitatively different

---

17. *See* Mata v. Avianca, 678 F. Supp. 3d 443, 448–49 (S.D.N.Y. 2023) (sanctioning attorneys for submitting hallucinated cases from ChatGPT); Johnson v. Dunn, 792 F. Supp. 3d 1241, 1246 (N.D. Ala. July 23, 2025) (sanctioning two attorneys who failed to supervise AI-assisted filings and who had no idea AI was used in drafting the submissions). *See generally* Ralph Artigliere, *When AI Policies Fail: The AI Sanctions in* Johnson v. Dunn *and What They Mean for the Profession*, JD SUPRA (Aug. 1, 2025), https://www.jdsupra.com/legalnews/when-ai-policies-fail-the-ai-sanctions-9043268/ [https://perma.cc/96LY-GWVR].

18. Artigliere, *supra* note 17.

19. *Id.*; *see generally* Ralph Artigliere & William F. Hamilton, *Reasonable or Overreach? Rethinking Sanctions for AI Hallucinations in Legal Filings*, EDRM (Aug. 18, 2025), https://edrm.net/2025/08/reasonable-or-overreach-rethinking-sanctions-for-ai-hallucinations-in-legal-filings/ [https://perma.cc/FLY8-2YMG] (arguing that sanctions for hallucinated citations must be principled, proportionate, and consistent with existing jurisprudence).

20. For a discussion of the range of court orders from prohibition to disclosure of AI use, *see* Gary E. Marchant, *AI in Robes: Courts, Judges, and Artificial Intelligence*, 50 OHIO N.U. L. REV. 473 (2024).

21. An increasing number of hallucinated citations and fabricated authority have now been cataloged across jurisdictions, reinforcing that such failures are neither isolated nor hypothetical. *See, e.g.*, Damien Charlotin, *AI Hallucination Cases Database*, https://www.damiencharlotin.com/hallucinations/ [https://perma.cc/A87R-VHET] (last visited Feb. 21, 2026) (collecting judicial decisions and sanctions involving AI-generated errors).

from earlier technological transitions: how generative AI reshapes analysis, authorship, and evidentiary judgment, and how those differences expose the institutional responses that either narrow or harden the competence divide.

IV. GENERATIVE AI AND THE EXPANSION OF EPISTEMIC ASYMMETRY

AI represents a qualitative break from prior legal technologies. Tools that once assisted retrieval, organization, or prioritization now directly participate in synthesis, framing, and analysis. As a result, generative systems operate not merely as efficiency aids, but as force multipliers that reshape how legal judgment is formed and exercised.

Lawyers and judges who develop the literacy necessary to understand and supervise these systems gain new capacity to engage directly with data, evidence, and work product that were previously delegated to specialists or external actors.[22] This shift has implications across litigation practice, from discovery strategy and expert preparation to evidentiary presentation and trial advocacy. The significance of these changes lies not in any single use case, but in their cumulative effect on how knowledge enters and moves through the legal system.

Because generative AI alters cognition rather than sequence, the profession's historical reliance on incremental adaptation, effective when technology supported human judgment within visible and testable constraints, is unlikely to prevent the existing competence divide from widening.

Initial judicial and professional attention has understandably focused on the most visible failures with AI, such as hallucinated citations or fabricated authority. Those failures are serious, but they do not define the full scope of risk. AI also introduces more subtle and therefore more consequential dangers. These include inaccurate summaries of evidence, distorted factual narratives, incomplete or misleading synthesis of the record, and AI-influenced expert analyses that appear authoritative but rest on flawed premises. In these contexts, the risk is not fabrication, but distortion: errors that may not be recognized as errors, yet can materially shape submissions, rulings, and outcomes.[23]

---

22. *See generally* Xavier Rodriguez, *Judging AI: How U.S. Judges Can Harness Generative AI Without Compromising Justice*, 109 JUDICATURE No. 2 at 17 (2025) (discussing judicial responsibility to supervise generative AI use and to guard against subtle distortions in analysis and decision-making); John Tredennick & William Webber, *Generative AI for Smart Discovery Professionals* (Merlin Search Techs., 4th ed. 2025), https://21577496.hs-sites.com/hubfs/4th%20Ed.%20Generative%20AI%20For%20Smart%20Discovery%20Professionals.pdf [https://perma.cc/5BVS-JZUM] (describing advantages as well as risks associated with AI-generated summaries, synthesis, and analytical outputs in legal workflows).

23. *See generally* The Sedona Conference, *The Sedona Canada Primer on Artificial Intelligence and the Practice of Law*, 26 SEDONA CONF. J. 99 (2025) (emphasizing that AI-related risks in legal practice often arise from subtle distortion, overreliance, and inadequate supervision

Whether AI enhances or undermines legal judgment and performance depends not on its availability, but on the user's understanding of its limits and the discipline with which its outputs are verified. Where that understanding is uneven, the result is not merely individual error, but structural asymmetry in how facts are analyzed and arguments are formed. To see why these risks arise, it is necessary to examine how AI differs in kind, not merely degree, from the discovery technologies courts have previously accommodated.

A. *From Ranking to Synthesis*

Technology-Assisted Review introduced ranking into discovery by reordering document visibility. Courts accommodated that shift within proportionality doctrine by emphasizing transparency, sampling, validation, and defensibility. Although ranking systems altered efficiency and review sequencing, they remained bounded. They classified documents to assist human reviewers. They did not generate legal argument, synthesize evidence, or replace professional judgment.

AI extends this architecture in a fundamentally different direction. These systems do not merely prioritize information—they synthesize, summarize, cluster, and explain patterns across large datasets. As a practical consequence, generative tools allow trial lawyers to reengage directly with discovery strategy and case theory formation—functions that, during the era of large-scale electronic discovery, were often delegated to specialists or external service providers.[24] This shift does not eliminate the need for expertise, but it alters who can meaningfully interrogate the record at early stages and how quickly strategic judgments can be formed.

Generative tools can generate chronologies, draft issue outlines, identify narrative themes, and simulate analytical reasoning. More importantly, they permit iterative engagement: lawyers can refine prompts, assess alternative hypotheses, request counterfactual summaries, and probe anomalies in real time. What began as statistical prioritization becomes dynamic interaction with evidence.

---

rather than obvious fabrication, and that such risks implicate competence, verification, and professional judgment); Varun Magesh et al., *Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools*, 22 J. EMPIRICAL LEGAL STUD. 216 (2025) (documenting that AI legal research tools produce hallucinations encompassing substantive distortion and false attribution, not only citation error, suggesting that visible fabrication understates the broader reliability risks of AI-assisted legal research).

24. *See generally* Tredennick & Webber, *supra* note 22 (describing the use of generative AI tools to synthesize, summarize, and iteratively analyze large legal datasets and contrasting those functions with earlier ranking-based discovery technologies). *Cf.* Surden, *supra* note 5, at 1312–14 (distinguishing between systems that assist information retrieval and those that perform higher-order analytical or synthetic functions beyond simple retrieval or ranking).

The distinction is structural. Ranking alters sequence; generative systems alter cognition. Proportionality is no longer confined to managing cost and volume. It becomes an analytical multiplier. Lawyers fluent in generative tools can compress investigative cycles, accelerate strategic assessment, and shape litigation narratives earlier than ever before. The informational landscape of litigation will change accordingly.

B. *Acceleration, Leverage, and Structural Divide*

During the TAR era, asymmetry reflected disparities in cost and efficiency. Parties with validated ranking models could identify key documents earlier and reduce review expense, but the underlying evidence remained human-authored and independently verifiable. Courts could evaluate disputes using familiar reasonableness standards.

AI extends asymmetry beyond efficiency into analytical leverage. Lawyers who deploy these tools effectively gain accelerated access not only to documents, but to framing. They can simulate opposing arguments, stress-test theories, identify inconsistencies in testimony, and surface latent connections across extensive evidentiary records. Control over informational sequence evolves into control over analytical trajectory.

This analytical acceleration extends beyond document review. Lawyers fluent in generative tools can more rapidly organize expert materials, test competing explanations, synthesize technical records, and identify weaknesses in opposing expert theories. The advantage lies not in automation, but in compression of the analytical cycle, allowing informed strategic decisions to be made earlier and with greater confidence.

This shift does not guarantee unfair outcomes. It does, however, recalibrate capability within adversarial proceedings. Where one party possesses generative fluency and the other does not, the disparity affects strategic foresight, negotiation posture, and litigation pacing. The risk is not that artificial intelligence replaces human judgment, but that uneven mastery amplifies institutional inequality in ways existing procedural doctrine was not designed to absorb.

C. *Integrity Failures as Signals, Not the Core Problem*

Public attention has focused on instances in which generative tools produce fabricated citations or inaccurate summaries. Courts have responded with sanctions and standing orders requiring human verification of AI-assisted filings. These measures do not prohibit innovation—they reaffirm that professional responsibility and verification remain nondelegable.

Such episodes are significant, but they are not the central concern. Hallucinated authority and careless reliance on automated outputs represent failures of supervision and competence, not evidence of inherent instability in the technology itself. The more consequential issue is uneven integration. Generative systems can enhance analysis when used by practitioners who understand their limits and validate outputs. The integrity risk arises when that competence remains unevenly distributed. Focusing reform efforts exclusively on visible integrity failures risks obscuring the more durable asymmetries in analytical capacity and institutional leverage that generative AI introduces.

D. *Evidence, Authentication, and Trial-Level Asymmetry*

AI introduces a related challenge at the evidentiary stage. Rules governing authentication and expert testimony require courts to assess authorship, provenance, and reliability. As synthetic text, images, audio, and video become more prevalent, those assessments grow more technically demanding.

Unlike earlier forms of electronic evidence, generative systems can produce persuasive but fabricated artifacts that are difficult to detect without specialized expertise. The burden of identifying and challenging such materials is unevenly distributed. Well-resourced parties may retain forensic experts to probe authenticity and interrogate model outputs. Smaller firms and pro se litigants may lack comparable capacity.

At trial, these disparities become more pronounced. Generative tools can assist in organizing themes, developing demonstratives, modeling juror responses, and synthesizing testimony in real time. When one party possesses the competence to deploy such tools responsibly and the other does not, the imbalance affects not only efficiency but narrative formation, placing increased pressure on courts to ensure reliability and fairness. Competence in generative systems therefore affects not only a party's ability to present evidence but also its capacity to challenge the authenticity and reliability of AI-modified or AI-generated materials offered by an opposing party.

This asymmetry thus extends beyond discovery. It reaches admissibility and narrative formation at trial. Judges must evaluate increasingly technical disputes while maintaining procedural fairness.

E. *Ethical Doctrine as Exposure, Not Cure*

Ethical doctrine thus functions less as a solution to AI-driven asymmetry than as an early warning system that exposes where institutional competence has failed. AI does not create new ethical duties. It exposes the depth of existing ones. Rules governing competence and confidentiality were drafted in technologically neutral language precisely

to endure change. In the generative era, those rules assume heightened structural significance.[25]

The ethical duty of competence in the context of artificial intelligence operates at multiple levels. At a minimum, lawyers must possess sufficient conceptual understanding to recognize that generative AI systems do not retrieve verified information, but instead generate probabilistic outputs based on patterns learned from training data. Without that baseline understanding, lawyers risk over-attributing reliability to fluent and persuasive text.

Competence, however, does not end with conceptual awareness. It also requires operational judgment. Lawyers who deploy AI tools must implement verification practices appropriate to the task at hand, must understand the limitations and failure modes of the systems they use, and must exercise independent professional judgment rather than deferring to machine-generated output. Existing rules governing supervision, candor, and diligence already impose these obligations; AI exposes whether those duties are being satisfied.[26]

The implications of AI competence extend beyond accuracy and verification to the preservation of attorney-client privilege. Generative systems differ widely in data retention, training use, and third-party access. Lawyers and their clients[27] who input information into AI tools without understanding those mechanics risk inadvertent disclosure or waiver under long-settled privilege doctrine. The law of privilege has not changed; what has changed is the ease with which privileged material can be transmitted beyond the lawyer's control through uninformed use of opaque systems. Competent lawyers can manage this risk through informed tool selection, contractual safeguards, and disciplined workflows. Incompetent use, by contrast, may compromise privilege invisibly and irreversibly. Once again, AI does not create a new ethical

---

25. *See* MODEL RULES OF PRO. CONDUCT r. 1.1 cmt. 8 (A.B.A. 2024); R. Regulating Fla. Bar 4-1.1 cmt. (2024); Mark D. Killian, *Court Approves CLE Tech Component*, FLA. BAR NEWS (Oct. 15, 2016), https://www.floridabar.org/the-florida-bar-news/court-approves-cle-tech-component/ [https://perma.cc/Q3BX-JZA6] (reporting that Florida was the first jurisdiction to impose a mandatory technology-specific CLE requirement, requiring three hours of technology CLE every three-year reporting cycle).

26. Safe use entails issues of confidentiality; human supervision, verification and responsibility oversight of output; and managing limitations inherent in AI models. *See* Fla. Ethics Op. 24-1.

27. Lawyers have a responsibility to advise clients about potential legal missteps for privilege. Now lawyers need to understand the implication of using AI tools to do so. *See* Michael D. Berman, *A.I. Documents Deemed Not Privileged*, EDRM (Feb. 12, 2026), https://edrm.net/2026/02/a-i-documents-deemed-not-privileged/ [https://perma.cc/3C4X-XUGV] (discussing a federal court decision holding that documents prepared using a commercial AI service and transmitted to counsel were not protected by attorney-client privilege).

problem—it exposes whether existing duties of competence and supervision are being met.

Finally, ethical competence in this setting includes knowing when reliance on specialized expertise is required.[28] Just as lawyers do not attempt to substitute themselves for forensic accountants or medical experts, responsible AI use may require consultation with technologists or information-governance professionals.[29] The duty of competence is therefore not a demand for technical mastery, but for informed oversight. Where that oversight is absent, the risk is not merely inefficiency, but distortion of the factual and analytical foundations on which legal judgments rest.

Competence now includes disciplined oversight of probabilistic systems that shape legal analysis. Verification is nondelegable. Confidentiality obligations similarly require informed evaluation of system architecture, data retention, and third-party access. Ethical doctrine functions as a stabilizing mechanism by translating technological missteps into enforceable consequences.

But ethical enforcement operates reactively. Sanctions and disciplinary proceedings correct individual failures after distortion has occurred, but they do not close structural divides. Ethical rules expose the competence divide; they do not resolve it. Addressing the structural effects of AI therefore requires coordinated institutional adaptation in legal education, firm practice, judicial engagement, and procedural governance.

F. *Institutional Responses That Widen Rather Than Narrow the Competence Divide*

Institutional responses to generative AI provide a clear illustration of how well-intentioned risk management can, if poorly calibrated, deepen rather than mitigate structural asymmetry. Early reactions by courts, firms, and regulators reveal a recurring pattern: visible failures prompted restrictive measures that were ethically unsound, operationally ineffective, and ultimately counterproductive to the system's integrity.

The public availability of large-scale generative AI systems in late 2022 marked an inflection point for the legal profession. Within months of ChatGPT's release, early adopters across practice settings and courts began integrating AI tools into legal workflows, including research,

---

28. *See* MODEL RULES OF PRO. CONDUCT r. 1.1 cmt. 2 (A.B.A. 2024); R. Regulating Fla. Bar 4-1.1 cmt. (2024) (A lawyer may achieve competent representation "through the association of a lawyer of established competence in the field in question.").

29. *See* A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512, at 3 (2024) ("[L]awyers should either acquire a reasonable understanding of the benefits and risks of the GAI tools that they employ in their practices or draw on the expertise of others who can provide guidance about the relevant GAI tool's capabilities and limitations.").

drafting, factual synthesis, document organization, expert witness management, and case analysis.[30] Competing models from major technology companies and the rapid integration of generative AI capabilities into legal research platforms accelerated experimentation and institutional adoption. These developments demonstrated that AI could materially enhance efficiency and analytical capacity across the profession.

At the same time, a parallel and contrasting trend emerged. Lawyers unfamiliar with the limitations of generative systems began submitting filings containing fabricated citations, distorted summaries, and unverified analysis. Courts responded with sanctions, disqualification, and referrals to disciplinary authorities.[31] In reaction, some firms and courts adopted restrictive or prohibitory measures, including bans on AI use or disclosure regimes that functioned as deterrents rather than accountability mechanisms. These responses were understandable. The technology was new, the risks were visible, and the consequences were real. Those reactions were driven not only by technological uncertainty, but by legitimate concerns about reputational harm, client trust, and judicial legitimacy in a system that depends on public confidence. But their institutional effects were counterproductive.

From an ethical and supervisory perspective, blanket prohibitions fail as a matter of professional responsibility. The duty of competence requires lawyers to understand the technologies they use—or reasonably should use—in modern practice.[32] The duty of supervision requires organizations to train, guide, and monitor professional work, not to disclaim responsibility by forbidding tools that are already shaping legal workflows.[33] Banning AI use does not satisfy these duties; it evades them. It treats ignorance as risk mitigation and substitutes fear for governance. Just as uncritical reliance on AI abdicates professional judgment, categorical bans abdicate the obligation to develop it.[34]

---

30. *See* Rodriguez, *supra* note 22.

31. *See, e.g.*, Mata v. Avianca, 678 F. Supp. 3d 443, 448–49 (S.D.N.Y. 2023); Johnson v. Dunn, 792 F. Supp. 3d 1241, 1246 (N.D. Ala. July 23, 2025).

32. MODEL RULES OF PRO. CONDUCT r. 1.1 cmt. 8 (A.B.A. 2024); R. Regulating Fla. Bar 4-1.1 cmt. (2024); A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512 (2024).

33. *See* MODEL RULES OF PRO. CONDUCT r. 5.3 (A.B.A. 2024); R. Regulating Fla. Bar 4-5.3 (2024); *see also* ABA Comm. on Ethics & Pro. Responsibility, Formal Op. 512 (2024); Fla. Bar Pro. Ethics Comm., Op. 24-1 (2024) (explaining that lawyers must supervise AI they use the same way they supervise nonlawyer assistants under Rule 5.3 and Florida Rule 4-5.3).

34. *Cf.* J.D. Morris, *Word Can't Even Spell-Check After 40 Years: You Trust AI with Your Law License?*, THE TECH. BLINDSPOT (Feb. 12, 2026), https://www.linkedin.com/pulse/word-cant-even-spell-check-after-40-years-you-trust-ai-jd-morris-0pk7c [https://perma.cc/FV5H-GRC4] ("Banning AI makes the same fundamental error as using it uncritically. Both approaches dodge the actual obligation: understanding the technology well enough to use it responsibly.").

These ethical shortcomings translate directly into structural consequences. Sanctions and prohibitions did not eliminate improper use.[35] Nor did bans prevent continued, unregulated reliance on AI by lawyers operating without guidance or supervision.[36] Instead, these measures displaced use into informal and invisible channels, outside institutional safeguards. Lawyers and institutions with the resources, technical support, and organizational commitment to develop AI literacy continued to advance rapidly. Others were discouraged from engagement altogether or driven underground. The result was not risk reduction, but asymmetrical risk concentration.

This dynamic widens the competence divide along multiple dimensions: capability, supervision, transparency, and accountability. Those who use AI competently and openly refine workflows, improve accuracy, and develop institutional knowledge. Those excluded from disciplined engagement either avoid AI entirely—ceding efficiency and analytical leverage—or use it irresponsibly in the shadows. The disparity is not merely economic. It is epistemic. It affects who can synthesize records, test theories, identify weaknesses, and shape litigation narratives early and effectively.

The lesson is not that institutions should refrain from managing risk, but that suppression is a poor substitute for competence formation.[37] Institutional responses that emphasize prohibition rather than education, supervision, and verification harden the divide between those able to use AI responsibly and those excluded from meaningful engagement.[38] As with earlier technological transitions, the profession faces a choice. It can respond by narrowing access and elevating fear, or it can respond by building shared standards of competence that preserve fairness, accountability, and legitimacy in an AI-mediated legal system.

---

35. Prohibition is ethically unsound because ethical requirements and guidelines around the country, starting with Comment 8 to Model Rule 1.1 and state equivalents, require that lawyers maintain competence by keeping "abreast of changes in the law and its practice, including the benefits and risks of relevant technology." MODEL RULES OF PRO. CONDUCT r. 1.1 cmt. 8 (A.B.A. 2024). *See* R. Regulating Fla. Bar 4-1.1 cmt. (2026); A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512 (2024); Fla. Bar Pro. Ethics Comm., Op. 24-1 (2024).

36. Morris, *supra* note 34. The expectation that overworked lawyers will not avail themselves of powerful, time-saving tools until an organization sorts out the details through months or years of study appears to be unrealistic.

37. *See* Jonathan Lent & Kyu Young Paek, *Common Issues That Arise in AI Sanction Jurisprudence and How the Federal Judiciary Has Responded to Prevent Them*, A.B.A. BUS. L. TODAY (Sep. 17, 2024), https://www.americanbar.org/groups/business_law/resources/business-law-today/2024-september/common-issues-arise-ai-sanction-jurisprudence/ [https://perma.cc/T3D4-CXCN].

38. *See* Rodriguez, *supra* note 22 (encouraging judges to master AI and to not yield to panic and excessive reaction over hallucinated citations in cases); Lent & Paek, *supra* note 37 (arguing that a local rule or court order that reminds parties that they are subject to sanctions if they misuse AI is a sufficient and more adaptable way to address concerns surrounding AI).

V. EMPIRICAL EVIDENCE THAT THE DIVIDE CAN NARROW

Recent empirical work suggests that the competence divide created by generative AI is neither fixed nor inevitable.[39] In a randomized controlled study of law students performing realistic legal tasks, researchers found that access to AI tools produced large and consistent reductions in time across all participants,[40] while quality improvements were uneven and concentrated among lower-performing users.[41] Higher-performing participants saw little or no quality improvement, but also saw no systematic degradation.[42]

In this study, law students were randomly assigned to complete four common legal tasks—drafting a complaint, preparing a contract, composing a section of an employee handbook, and advising a client through a short memorandum—either with or without access to GPT-4, following targeted training in its use. Outputs were blind graded for quality, and time to completion was measured. The study therefore tested not abstract exam performance, but the kind of drafting, synthesis, and advisory work that junior lawyers routinely perform in practice.

The study's most important finding for purposes of this Article is not that AI improved performance unevenly, but how it did so. Where AI assistance improved quality at all, the largest gains accrued to participants who began with lower baseline performance, while higher-performing participants saw little or no quality improvement. At the same time, time savings were distributed relatively evenly across skill levels. This pattern undermines the assumption that AI inevitably entrenches existing hierarchies. Properly trained and supervised, generative AI can function as a competence-leveling tool, reducing disparities in output quality without eliminating the need for professional judgment. The results suggest that competence with AI is not fixed or innate, but responsive to education, workflow design, and institutional expectations.

These findings reinforce that AI can function as a partial equalizer only when accompanied by discipline, training, and verification, not when adopted reflexively or without understanding. Importantly, participants who blindly relied on AI output performed worse than those who actively interrogated and edited it.[43] The study therefore supports a critical distinction that this Article emphasizes: AI does not eliminate epistemic asymmetry through access alone. It narrows the divide only when accompanied by competence, judgment, and structured adaptation.

---

39. Jonathan H. Choi, Amy B. Monahan & Daniel Schwarcz, *Lawyering in the Age of Artificial Intelligence*, 109 MINN. L. REV. 147, 206 (2024).

40. *Id.* at 182.

41. *Id.* at 181.

42. *Id.*

43. *Id.* at 166.

The study suggests that AI reduces disparities in time efficiency while simultaneously increasing disparities in judgment unless users possess sufficient professional competence. The findings reflect, in concentrated form, dynamics already playing out across the profession. Where AI competence is uneven, the divide widens; where it is treated as a professional priority, AI becomes a tool for narrowing that divide. Ensuring that outcome is an institutional responsibility shared by law schools, courts, firms, and the profession as a whole, one that cannot be deferred without consequence.

If generative AI can narrow the competence divide when integrated with discipline and training, but widen it when adopted haphazardly or suppressed, then the profession's task is no longer diagnosis but design. That task requires building institutional structures that allow AI to advance, rather than undermine, the just, speedy, and inexpensive determination of actions.

VI. INSTITUTIONAL REFORM: A CALIBRATED ROADMAP FOR CLOSING THE COMPETENCE DIVIDE

The legal profession has navigated technological disruption before, but its adaptations have rarely been anticipatory or broadly successful. Historically, doctrine has absorbed technology only after friction, failure, and uneven learning. AI compresses that cycle dramatically. It accelerates the pace at which errors propagate and magnifies the consequences of uneven technological competence. Addressing the resulting competence divide therefore requires a coordinated institutional response grounded in the core purposes of the procedural system.

Federal Rule of Civil Procedure 1 supplies the governing principle. The Rules are to be construed and administered to secure the just, speedy, and inexpensive determination of every action. AI has the potential to advance all three objectives—but only if its use is disciplined, transparent, and subject to meaningful oversight. If unevenly understood or poorly governed, it risks undermining fairness, distorting fact development, and eroding confidence in adjudication.

The competence divide that emerged during earlier waves of legal technology was not the product of indifference, negligence, or resistance to change. It was an understandable institutional response to mounting complexity under real-world constraints. Trial lawyers and judges were already operating within demanding, deadline-driven environments that required constant attention to substantive law, procedural developments, and case-specific demands. Mastery of electronic discovery and information technology was not a discrete task, but a rapidly evolving domain requiring sustained engagement at a time when professional bandwidth was already fully consumed.

Educational infrastructure compounded the problem. Comprehensive and practical eDiscovery education was not offered at scale when discovery became predominantly electronic. In many jurisdictions, including large state court systems, robust eDiscovery and technology CLE and continuing judicial education programs reached only a small fraction of the profession.[44] Law schools were similarly late to integrate eDiscovery as a core component of civil procedure and trial practice, even as discovery disputes increasingly turned on electronically stored information.

The result was not a failure of will, but a mismatch between the pace of technological change and the profession's capacity to deliver systematic formation and training. Because discovery is expensive and shapes settlement in the overwhelming majority of cases, diffusion of control at that stage produced outsized consequences for outcomes, leverage, and client expectations.

These educational gaps are compounded by institutional design choices that have lagged behind technological reality. Legal education, licensure, and professional incentives have historically emphasized doctrinal mastery and adversarial skill while largely treating technological competence as peripheral or optional. Bar examinations rarely test technological literacy, law school curricula address it unevenly, and continuing education requirements tend to respond to failures after they occur. At the same time, billing structures and professional norms often penalize efficiency and experimentation, discouraging the disciplined adoption of tools that could otherwise enhance accuracy and access. The resulting competence divide is therefore not simply a matter of individual aptitude or caution, but a predictable consequence of institutional frameworks that have not yet recalibrated for AI-mediated practice.

In that environment, specialization and delegation emerged as rational adaptations. Dedicated eDiscovery lawyers, consultants, vendors, and service providers filled an urgent need and brought technical expertise into cases. But over time, those adaptations carried unintended consequences. Trial lawyers and judges increasingly ceded direct

44. This assessment reflects the authors' experience teaching and observing electronic discovery and litigation technology over more than two decades across judicial education programs, law school courses, and continuing legal education for practicing lawyers. During that period, comprehensive and practice-oriented eDiscovery education was offered sporadically and reached only a limited portion of the profession at any given time, even as discovery disputes increasingly turned on electronically stored information. As a result, many trial lawyers and judges reasonably relied on specialists, consultants, or vendors to manage discovery workflows rather than developing direct operational competence themselves. That adaptive response was rational under conditions of time pressure and limited educational capacity, but it had the unintended effect of diffusing strategic control over the discovery phase of litigation, with significant implications for leverage and outcomes.

engagement with discovery—the phase of litigation that shapes outcomes in the vast majority of cases—to specialists and external actors. Strategic control over the development of the evidentiary record became diffused, even as ethical responsibility remained with counsel of record and the court. That pattern of delegation worked, imperfectly, because earlier technologies assisted retrieval and organization without participating directly in synthesis, reasoning, or judgment. Generative AI cannot be absorbed in the same way. If AI literacy is treated as a niche specialty rather than a baseline professional competence, the divide will not merely persist; it will harden. The profession cannot afford to repeat this adaptive pattern. AI competence must be treated as foundational, and that responsibility must be shared among law schools, CLE and CJE providers, courts, firms, and individual lawyers and judges alike.

This roadmap identifies four interlocking institutional levers—formation, oversight, incentives, and coordination—through which the profession can respond deliberately while preserving flexibility and encouraging responsible innovation. The reforms proposed below are not intended to be aspirational. They respond directly to the failure modes identified in the preceding sections and are designed to guide deliberate institutional adaptation at a moment when technological acceleration leaves little margin for delay.

A. *Formation: Legal Education and Licensure as Baseline Infrastructure*

Technological competence can no longer be treated as optional enrichment in legal education and licensure. If AI-assisted tools now shape how facts are synthesized, arguments developed, and evidence presented, then baseline professional formation must account for those realities.

Law schools should integrate AI-related competence[45] into core doctrinal courses rather than confining it to electives. Civil procedure courses should address AI-assisted discovery workflows and proportionality analysis under Rule 26 in technology-intensive cases. Evidence courses should engage authentication, provenance, and expert reliability challenges raised by AI-generated or AI-assisted materials. Professional responsibility courses should confront verification duties, confidentiality risks, and supervision obligations in AI-integrated practice.[46]

---

45. *See, e.g.*, MODEL RULES OF PRO. CONDUCT r. 1.1 cmt. 8 (A.B.A. 2024); R. Regulating Fla. Bar 4-1.1 cmt. (2024); Fla. Ethics Op. 24-1.

46. For example, at the University of Florida Levin College of Law, Master Legal Skills Professor William Hamilton integrates practical technology instruction into his eDiscovery courses through industry-standard tools and applied exercises, illustrating the feasibility of embedding technology competence within doctrinal instruction. *See* Univ. of Fla. Levin Coll. of

This integration need not train students to become technologists. Its purpose is conceptual and operational literacy: understanding what these systems do, where they fail, and how professional obligations attach to their use. Clinics and simulation courses are particularly well suited to reinforce disciplined workflows. For example, practice courses can require students to document verification steps for AI-assisted research memoranda and justify reliance on AI-generated summaries of records.

Licensure standards should reinforce this foundation. Bar examinations need not test facility with particular tools, but they should assess understanding of recurring issues: eDiscovery obligations, Rule 26 proportionality in technology-driven cases, authentication of digital evidence, AI-assisted drafting risks, and confidentiality exposure. Licensing has inevitably shaped curricula. If technological competence is examinable, it becomes foundational rather than discretionary, just as civil procedure and evidence once did.

B. *Oversight: Judicial Engagement as System Stabilization*

Judges occupy a central position in managing the effects of technological change on litigation practice. Discovery rulings, evidentiary gatekeeping, and case management orders shape how AI tools are deployed and constrained in practice.

Judicial education should therefore address, as part of core case-management instruction, AI-assisted discovery, authentication disputes involving synthetic or AI-modified evidence, proportionality arguments where AI materially alters cost and scope, and disclosure expectations regarding AI use.[47] Rule 26 conferences increasingly turn on

---

Law, *Electronic Discovery, Digital Investigations, and Evidence* (William Hamilton, Master Legal Skills Professor) (describing the course as providing practical instruction in eDiscovery issues and introducing students to industry-standard electronic discovery tools), https://law.ufl.edu/academics/courses-and-syllabi/courses/electronic-discovery/ [https://perma.cc/8WSJ-LT33] (last visited Feb. 21, 2026). Similarly, Professor Craig Ball, Adjunct Professor at University of Texas School of Law, uses hands-on, tool-based, practical eDiscovery instruction in his teaching and training programs. *See* Univ. of Tex. Sch. of Law, *Course Description: Electronic Discovery & Digital Evidence* (Craig Ball, Adjunct Professor) (noting that the course includes practical demonstrations and hands-on tool use) https://law.utexas.edu/courses/class-details/20262/29574/ [https://perma.cc/Z573-KKK9] (last visited Feb. 21, 2026).

47. State and federal judicial education programs have increasingly incorporated electronic discovery, litigation technology, and AI-related issues into core continuing judicial education curricula, often through scenario-based instruction, bench memoranda, and simulated hearings rather than standalone technology sessions. These programs suggest that judges can develop functional literacy in emerging technologies sufficient to manage cases, evaluate disputes, and supervise counsel without becoming technical specialists. Programs offered through state judicial colleges, the National Center for State Courts, the Federal Judicial Center, and national judicial conferences reflect a shift toward integrating technology competence into mainstream judicial formation rather than treating it as peripheral. The routine handling of complex electronic discovery and evidentiary disputes by federal magistrate judges further illustrates that sustained

technological preparedness. When one party arrives with a sophisticated understanding of AI-assisted review, expert analysis tools, or data architecture and the other does not, the imbalance affects more than efficiency. It shapes negotiation leverage, scope determinations, and litigation trajectory.

Evidentiary gatekeeping under Rules 901 and 702 likewise becomes more demanding. Judges must assess reliability and provenance where AI is used to summarize voluminous records, assist experts in forming opinions, or generate demonstratives and analyses. Judicial fluency allows courts to distinguish between legitimate assistance and impermissible substitution of judgment.

Importantly, calibrated oversight should not discourage responsible use. Courts that understand AI tools can regulate them without reflexive prohibition—permitting their use for tasks such as expert witness management, early case assessment, and document drafting improvement while insisting on transparency and verification where accuracy and reliability are outcome-determinative.[48]

C. *Incentives: Law Firm Practice, Client Expectations, and Access to Justice*

Law firm economics and client demands will significantly influence how AI is integrated into practice. Corporate clients increasingly expect efficiency, cost transparency, and technological sophistication. Firms that responsibly deploy AI to manage discovery, analyze records, assist in expert preparation, or improve drafting quality[49] can advance Rule 1 objectives while strengthening client trust and accountability.

Incentives, however, must reward accuracy rather than volume. Responsible AI use requires verification protocols, documentation practices, and supervision standards embedded into training and evaluation. Treated this way, AI competence becomes a professional skill rather than a productivity shortcut. Where incentives instead discourage engagement or reward unchecked output, the competence divide deepens.

---

exposure and education can produce judicial fluency in technologically complex matters without requiring technical specialization.

48. *See* Rodriguez, *supra* note 22, at 17; Lent & Paek, *supra* note 37.

49. The incentive effects discussed here are not limited to litigation practice. Office-based specialties such as mergers and acquisitions, real estate, trusts and estates, and transactional practice more generally face similar pressures regarding document quality, consistency, and efficiency. In those contexts, generative AI tools are increasingly used for document management, organization, draft comparison, issue spotting, and clarity of expression, raising parallel questions about competence, supervision, and incentive alignment. *See generally* Andrea Bucher, Comment, *Navigating the Power of Artificial Intelligence in the Legal Field*, 62 HOU. L. REV. 819, 827 (2025).

AI also holds substantial promise for improving access to justice.[50] Properly governed, it can assist with form drafting, issue identification, and preliminary legal analysis in settings long constrained by cost and resource limitations. Public interest lawyers, legal aid organizations, and state court systems assisting pro se litigants are already among early beneficiaries of responsible AI use. Empirical studies suggest that AI tools can increase productivity and reduce time burdens across a wide range of legal service providers, particularly when paired with training and oversight.[51] These advantages have direct implications for access to justice across the system. AI can assist individuals through self-help tools[52] and improved pathways to representation, enabling providers at all levels to deliver services more efficiently and economically.[53] It can also support courts and judges in modernizing processes that have historically impeded fair and timely outcomes,[54] but only where users possess the competence to deploy these tools responsibly.

The challenge is not whether AI should be used, but whether incentives encourage disciplined use rather than suppress engagement. Overly rigid restrictions risk concentrating effective AI use among already well-resourced actors while discouraging competence

50. *See generally* Drew Simshaw, *Access to A.I. Justice: Avoiding an Inequitable Two-Tiered System of Legal Services*, 24 YALE J.L. & TECH. 150 (2022).

51. Colleen V. Chien & Miriam Kim, *Generative AI and Legal Aid: Results from a Field Study and 100 Use Cases to Bridge the Access to Justice Gap*, 57 LOY. L.A. L. REV. 903, 909 (2025) (90% of legal aid professionals given free access to AI tools in the study indicated some level of productivity increase and 75% of participants indicated their intention to continue using tools).

52. Arming pro se litigants with self-help tools is not without potential complications. *See* Sean Steward, *Are A.I. Lawyers a Legal Product or Legal Service?: Why Current UPL Laws Are Not up to the Task of Regulating Autonomous A.I. Actors*, 53 HOFSTRA L. REV. 392 (2025) (stressing the importance of establishing a clear boundary between legal and non-legal work for autonomous AI systems); Maria E. Berkenkotter & Lino S. Lipinsky De Orlov, *Can Robot Lawyers Close the Access to Justice Gap?*, COLO. LAW. (Dec. 2024), https://cl.cobar.org/wp-content/uploads/2024/11/PROFESSIONAL-CONDUCT-AND-LEGAL-ETHICS-DEC2024.pdf [https://perma.cc/6XY9-SLB7] (arguing that AI could help litigants who are unable to afford legal services; but warning that autonomous AI systems risk providing inaccurate legal information or engaging in the unauthorized practice of law).

53. *See* Drew Simshaw, *Interoperable Legal AI for Access to Justice*, 134 YALE L.J. F. 795, 795–96 (2025) (arguing that it is time to confront the challenges to a gap in access to justice by tapping AI's potential to assist self-help litigants, lawyers who serve them, and the courts where they seek justice).

54. The National Center for State Courts (NCSC) has established the AI Policy Consortium for Law & Courts to provide resources and a collaborative forum for AI adoption in state courts. The Consortium convenes court leaders, technologists, and policy experts to develop governance frameworks and pathways for the use of AI in the judiciary. *See* Nat'l Ctr. for State Courts, *Artificial Intelligence (AI) Resources for Courts*, https://www.ncsc.org/resources-courts/artificial-intelligence-ai [https://perma.cc/5Y2L-Z83R] (last visited Feb. 18, 2026) (collecting materials and initiatives, including the AI Policy Consortium for Law & Courts, addressing policy development, governance, and implementation of AI technologies in state court systems).

development elsewhere.[55] If AI literacy is treated as a professional priority, the technology can narrow both efficiency and access gaps. If it is suppressed, delegated, or driven underground, the divide hardens along familiar lines of resources, training, and institutional support.

D. *Coordination: Standards, Guidance, and Institutional Alignment*

Fragmented responses invite inconsistency and strategic exploitation. Professional organizations, rulemakers, and courts should work toward shared baseline principles and guidance that align with existing doctrine rather than displace it.

Those principles are familiar: verification responsibility remains nondelegable; confidentiality obligations apply fully to AI-assisted workflows; AI-generated or AI-assisted materials are subject to discovery under established rules; and technological competence is integral to professional responsibility. Coordination does not require uniform rules or tool-specific mandates. It requires clarity about expectations and transparency about risk.

The profession has coordinated effectively in the past, most notably in its response to electronic discovery, through a combination of advisory-committee amendments,[56] judicial education, and the articulation of professional best practices. AI demands a similar collaborative approach, one that balances innovation with stability while preserving doctrinal continuity.

Taken together, these four institutional levers—formation, oversight, incentives, and coordination—define the conditions under which generative AI can be integrated without undermining fairness or legitimacy. None is sufficient alone. Education without oversight risks misplaced confidence; oversight without incentives breeds evasion; coordination without formation lacks traction. These interdependencies are illustrative rather than exhaustive: they identify the points at which institutional engagement can convert technological acceleration from a source of distortion into a source of stability. The Conclusion returns to why that conversion is no longer optional.

VII. CONCLUSION: COMPETENCE, LEGITIMACY, AND THE INTEGRITY OF THE SYSTEM

Generative artificial intelligence does not introduce an entirely new problem so much as it exposes a longstanding one. For decades, the legal profession has tolerated uneven technological competence, absorbing new tools incrementally through litigation, education, and doctrine. That model proved workable when technology assisted retrieval and

55. *See* Rodriguez, *supra* note 22, at 17*;* Lent & Paek, *supra* note 37.

56. *See supra* note 9.

organization without reshaping authorship, analysis, or judgment. AI alters that equilibrium. It operates directly on the core substrates of litigation: facts, narrative, analysis, and persuasion.

The central risk is not that AI will replace lawyers or judges. It is that uneven mastery of AI will recalibrate the adversarial system in ways doctrine alone cannot correct. When AI-competent participants can synthesize records, test theories, and frame narratives at machine speed, while others cannot reliably verify outputs or detect distortion, foundational assumptions begin to strain. Procedural fairness depends on informed participation. Adversarial testing presupposes a shared baseline of competence. Where that baseline fractures, legitimacy erodes.

Ethical doctrine exposes this competence divide, but it cannot close it. Sanctions and disciplinary proceedings operate after distortion has occurred. Standing orders and certification requirements reinforce existing duties of candor and diligence, but they do not build capacity. The structural risks posed by generative AI therefore require earlier and more deliberate institutional intervention—through education, judicial engagement, incentive alignment, and coordinated standards—rather than episodic enforcement alone.

Rule 1 provides the governing principle. Artificial intelligence can advance the just, speedy, and inexpensive determination of actions by improving accuracy, efficiency, and access when it is used with understanding, verification, and oversight. Rule 26 proportionality, evidentiary gatekeeping, and existing professional responsibility doctrines already supply the doctrinal framework. What is required is the competence to apply those doctrines meaningfully in an AI-mediated environment. That competence cannot be left to individual initiative or ad hoc adaptation alone. It depends on coordinated engagement by law schools, courts, firms, and the profession's licensing and training institutions.

AI competence also carries direct implications for access to justice. When responsibly integrated, generative tools can reduce cost, compress time, and expand the capacity of lawyers and courts to serve clients who have long been priced out of meaningful representation. That potential will not be realized if institutional responses emphasize restriction over disciplined use. Guardrails that discourage engagement or treat AI as presumptively suspect risk hardening the very competence divide they seek to prevent, concentrating effective use among well-resourced actors while leaving others behind.

The legal profession has navigated comparable technological transitions before. Electronic discovery, TAR, and digital evidence each required recalibration. The difference now is velocity. Adaptation must occur before distortion becomes structural. The choice is no longer whether AI will shape litigation; it already does. The question is whether

courts, educators, firms, and lawyers will engage AI deliberately, with shared responsibility for competence and verification, or allow uneven adaptation to redefine fairness by default.

AI can strengthen the legal system's core functions if integrated with discipline and care. Without that collective effort, the competence divide will widen along familiar lines of resources and access, and procedural legitimacy will erode predictably. The profession's task is therefore not to resist AI, nor to romanticize it, but to govern it—by ensuring that the tools now shaping legal judgment are matched by the competence required to use them responsibly.

APPENDIX

Because courts are already confronting these issues in real cases, we conclude with a practical framework for judicial case management.

APPENDIX A: PRACTITIONER AND JUDICIAL CASE MANAGEMENT GUIDANCE

(Non-Doctrinal Supplement)

This Appendix is intended as a practical supplement to the analysis in Parts IV through VI of this Article. It translates the doctrinal, institutional, and ethical considerations discussed in the main text into concrete case-management guidance for judges and litigants confronting the use of artificial intelligence in planning, discovery, evidentiary proceedings, and motion practice. Because its purpose is operational rather than analytical, the Appendix adopts a checklist and bullet-point format designed for practical reference and use. Nothing in this Appendix is intended to propose new procedural rules, alter existing doctrine, or impose tool-specific mandates; instead, it illustrates how established principles of proportionality, evidentiary gatekeeping, and professional responsibility may be applied in an AI-mediated litigation environment. This Appendix is illustrative, not exhaustive, and is intended to assist—not constrain—judicial discretion.

1. Framing the Issue Early: Case Management and Rule 26(f)

Generative AI should be addressed early, not reactively.

At initial case management conferences, Rule 26(f) conferences, and follow-up proceedings, courts may appropriately inquire—without mandating disclosure of work product—whether and how AI tools are being used in discovery workflows, document review, expert preparation, or factual synthesis. The goal is not surveillance, but informed proportionality.

Key considerations include:

- Whether AI materially affects the scope, cost, or sequencing of discovery.
- Whether AI-assisted review or summarization changes preservation or production strategies.
- Whether one party's technological preparedness significantly exceeds the other's in ways that may affect fairness.

Courts retain broad discretion to tailor discovery plans. AI competence disparities are relevant to proportionality analysis where they affect access to information or litigation leverage.

2. Discovery Oversight: Proportionality, Transparency, and Verification

Generative AI may legitimately be used to:

- assist with document review and prioritization,
- summarize voluminous records,
- identify themes or timelines,
- support expert witness preparation.

However, courts should remain attentive to **verification protocols**. Where AI-generated summaries or analyses are relied upon, parties should be prepared to explain:

- how outputs were verified,
- whether underlying source materials were preserved and reviewed,
- how human judgment was exercised.

The inquiry is familiar. It mirrors long-standing scrutiny of expert methodologies and TAR workflows. The focus remains on **reasonableness and defensibility**, not perfection.

3. Motion Practice and Submissions: Accuracy Over Volume

Courts may reasonably expect that filings—whether AI-assisted or not—reflect independent verification.

Standing orders or local practices that require certification of accuracy reinforce existing duties of candor and competence. They should be framed to **encourage responsible use**, not deter innovation. The emphasis should be on results and process, not tool selection.

Judges should remain alert to:

- unusually polished but shallow analysis,
- factual summaries untethered from record citations, and
- arguments that appear authoritative but lack doctrinal grounding.

These are not new concerns. AI increases their frequency and plausibility.

4. Evidence and Trial: Authentication and Expert Reliability

Generative AI intensifies traditional evidentiary questions.

When AI-generated or AI-assisted materials are offered at trial—whether summaries, demonstratives, expert analyses, or media—courts should apply familiar standards:

- Rule 901 for authentication and provenance,
- Rule 702 for expert reliability and methodology.

The presence of AI does not lower the threshold. It raises the need for clarity. Judges may require parties to explain:

- authorship and generation processes,
- source data integrity,
- steps taken to detect manipulation or error.

The burden of challenging AI-influenced evidence may be uneven. Courts should be attentive to that asymmetry in managing expert disputes and evidentiary hearings.

Addressing evidentiary issues early permits sufficient opportunity for both sides to address complex authentication and admissibility issues with expert assistance.

5. Judicial Posture: Regulation Without Overreaction

Effective judicial management of AI requires neither technical mastery nor categorical bans. It requires **conceptual fluency**.

Courts that understand what generative AI does—and does not do—are better positioned to:

- permit beneficial uses (efficiency, access, clarity),
- require safeguards where accuracy and reliability are critical,
- avoid chilling legitimate innovation.

The objective is stability, not stasis.

6. The Throughline: Competence as System Integrity

Generative AI is best understood as a **stress test** for the litigation system. It magnifies both skill and error. Judicial case management remains the primary stabilizing force.

By integrating AI awareness into local rules, case management, discovery oversight, evidentiary rulings, and ethical enforcement, courts can ensure that technological acceleration serves—rather than

undermines—the core commitments of fairness, accuracy, and public confidence.

THE RECOVERY OF MEANING IN A DIGITAL AGE: E-DISCOVERY AND THE RENEWAL OF JUDGMENT

William F. Hamilton[*]

Abstract

Modern litigation unfolds within a dense and expanding digital environment. Nearly every human activity now leaves behind emails, documents, metadata, collaborative communications, and machine-generated records that together form the factual foundation of contemporary litigation. This Article argues that electronic discovery is not merely a technical, managerial, or compliance process, but a core litigation practice concerned with the recovery, reconstruction, and rehabilitation of the factual world on which judgment depends.

Tracing the evolution of discovery from its twentieth-century roots to the present era of pervasive digital evidence, this Article shows how electronically stored information (ESI) continues to alter the dynamics of litigation itself. The volume, dispersion, and heterogeneity of digital artifacts means that no party can fully understand its own case without engaging the broader, shared digital record assembled through discovery and disciplined by judicial oversight. Preservation requirements, spoliation sanctions, authentication prescriptions, and proportionality principles operate not merely to regulate adversarial tactics, but to protect the recovery of lived experience and the court's capacity to ground its judgments in reality.

This Article then examines the role of generative artificial intelligence within this landscape. Unlike the earlier discovery tools of keyword searching and traditional technology-assisted review (TAR), generative AI can illuminate relationships among documents, events, and actors, situating individual artifacts within broader temporal and organizational contexts. Properly used, generative AI can assist lawyers in recovering structure, continuity, and perspective from vast corpora of digital material. But illumination is not judgment. Generative AI can organize information and reveal patterns, yet judgment requires human engagement, interpretation, and responsibility for meaning.

Drawing on discovery rules, principles, and practices, as well as Hannah Arendt's account of judgment and worldliness, this Article contends that the central challenge of the digital age is the recovery of context sufficient to support responsible judgment, perspective, and persuasion. As data and advocacy increasingly converge, the legal profession must resist the temptation to treat technological outputs as substitutes for judgment. Instead, e-discovery augmented by generative AI should be understood as a primary means by which courts and

* Master Legal Skills Professor, University of Florida Levin College of Law.

advocates restore factual depth, the plurality of perspectives, and temporal continuity to the processes of adjudication.

I. INTRODUCTION

A lawsuit does not begin only with the filing of a complaint. More fundamentally, it begins with a simple but profound premise: the world has left its mark. Nearly every human interaction—professional, commercial, or personal—now produces a residual digital shadow. But unlike the blurred silhouettes on the wall of Plato's cave, these shadows are not pale imitations. They are durable traces of lived experience. Emails accumulate quietly; text documents record their own evolution; metadata inscribes time and authorship; collaborative applications structure conversations; and automated systems preserve actions no one expected to be remembered.[1] These artifacts form the factual ground on which litigation proceeds. They do not declare their relevance; they simply wait to be understood.

1. *See* Rowe Ent., Inc. v. William Morris Agency, Inc., 205 F.R.D. 421, 428–33 (S.D.N.Y. 2002) (early case recognizing that electronically stored information presents discovery challenges including volume, accessibility, and cost); Zubulake v. UBS Warburg LLC, 217 F.R.D. 309, 313–18 (S.D.N.Y. 2003) (noting that digital information is created and stored across multiple systems and that traditional discovery assumptions do not readily translate to electronic data); George L. Paul & Jason R. Baron, *Information Inflation: Can the Legal System Adapt?*, 13 RICH. J.L. & TECH. 10, ¶¶ 1–5 (2007).

Discovery has never been a serene exchange. It arises from an adversarial tradition shaped by strategy, guardedness, and the fear that candor may arm one's opponent. ESI has nonetheless altered this dynamic in subtle but decisive ways. The volume, dispersion, fragility, and heterogeneity of digital material in the modern ESI environment means that no lawyer can fully grasp a client's story by looking inward alone.[2] Understanding one's own case increasingly requires looking outward—into the data repositories of others, including the opposition and third-party witnesses, into systems beyond one's control, and into perspectives framed by different roles, incentives, and memories. The factual world of litigation is irreducibly shared.

At the same time, the solidity of the world itself is no longer immediately apparent. Lived experience increasingly unfolds as a series of fragmented moments, severed from traditions, shared practices, and the durable things that once structured the space of the human world. We find ourselves suspended between the past and the future, in a present that overwhelms memory and offers little horizon for what comes next. As continuity erodes, meaning becomes harder to sustain, and judgment falters. Litigation, through the disciplined practices of discovery, can partially counter this disorientation by reconstructing sequences of events, preserving digital content, and staging facts to persuade others that a particular account of the world is the more faithful one.[3]

Generative artificial intelligence enters at this inflection point. For two decades, the working tools of e-discovery—keyword searches, Boolean operators, and predictive coding—produced results largely from fragments. Keywords yielded hits without context, and ranking algorithms sorted documents without explanation. Lawyers were left to infer meaning from decontextualized search results, encouraging case theories to be formed in advance of, rather than derived from, the emerging evidentiary record. This abstraction was compounded by the demands of modern litigation practice, which rely on standardized routines and forms that introduce distance and often obscure the lived experience animating the events in dispute. The data remained mute.[4]

---

2. *See* Paul & Baron, *supra* note 1, at ¶¶ 3, 6–15, 38–45 (describing how volume, dispersion, and interrelatedness of ESI require contextual analysis across data sets); Jason R. Baron, *Law in the Age of Exabytes: Some Further Thoughts on 'Information Inflation' and Current Issues in E-Discovery Search*, 17 RICH. J.L. & TECH. 1, at ¶¶ 2–7 (2010); The Sedona Conference, *Commentary on Proportionality in Electronic Discovery*, 18 SEDONA CONF. J. 141, 147–53 (2017).

3. *See* HANNAH ARENDT, BETWEEN PAST AND FUTURE 219–22 (rev. ed. 1968) (describing the modern condition as a gap between past and future that destabilizes inherited standards of judgment); *see also* HANNAH ARENDT, THE HUMAN CONDITION 50–58 (2d ed. 1998) (discussing the "world" as the human artifice that relates and separates people).

4. *See Zubulake*, 217 F.R.D. at 309, 317–20 (discussing cost shifting data under FED. R. CIV. P. 26 and limits of traditional retrieval).

Generative AI shifts the terrain. It does not review documents in isolation but uncovers the relations among them. It illuminates structure, continuity, and contradiction. It can situate a remark within a chain of communications, a decision within a workflow, or a revision within a developing pattern. Used with care, it enables what Hannah Arendt called an enlarged mentality: the discipline of holding before oneself the multiple perspectives that constitute a shared world.[5] But illumination is not judgment. Generative AI can organize information and reveal patterns, yet judgment and revealing meaning entail human engagement, interpretation, and responsibility.

This Article proceeds from the conviction that e-discovery is not merely a technical enterprise. The vocabulary of modern discovery—data volumes, indexing protocols, TAR models, and metadata extraction—risks obscuring a simpler truth: e-discovery is how litigation recovers the world. It is a sustained encounter with the world that parties have made. It captures hurried decisions, tentative drafts, casual remarks, and ordinary communications that later acquire extraordinary significance. A generation ago, most discoverable documents were curated artifacts: polished memoranda, formal letters, and annual reports. Today's ESI is the opposite. It is spontaneous, unfiltered, and embedded in the rhythms of everyday life. It is not a perfect portal into intention, but it anchors litigation in more reliable terrain.[6]

The judiciary plays a critical role in protecting this factual world. Through sanctions for spoliation, authentication doctrines, and evidentiary safeguards, courts guard against distortion. When evidence is lost or manipulated, the injury is not merely a tactical disadvantage; it is a wound to the court's ability to ground its judgments. Spoliation is a harm to the public world of adjudication, not simply to a party.[7]

This Article traces how e-discovery has adapted to the explosion of digital artifacts, examines the pressures that generative AI introduces, and argues that judgment remains the central task of the legal profession even as new tools reshape the terrain. The longstanding divide between "e-discovery" and "litigation" is collapsing. The data produced in discovery now reverberates through case analysis, briefing, oral advocacy, settlement evaluation, and cross-examination. Data and advocacy have merged, and the profession must adapt to a world in which context and

5. *See* HANNAH ARENDT, LECTURES ON KANT'S POLITICAL PHILOSOPHY 42–44 (Ronald Beiner ed., 1982).

6. *See* FED. R. CIV. P. 26(b)(2) advisory committee's note to 2006 amendment.

7. *See generally* Victor Stanley, Inc. v. Creative Pipe, Inc., 269 F.R.D. 497, 525–29 (D. Md. 2010) (Grimm, M.J.) (analyzing the court's inherent power to impose sanctions to preserve the integrity of the judicial process); *cf.* FED. R. CIV. P. 37(e) advisory committee's note to 2015 amendment (standardizing the framework for spoliation sanctions to address inconsistent applications of inherent authority).

lived experience, which were long missing from electronic discovery and litigation, are now increasingly available.

II. THE TRANSFORMATION OF DISCOVERY

A. *The Modern Era: The Rise of Discovery and Its Conundrums*

Discovery was not always a central feature of American civil litigation. Prior to the adoption of the Federal Rules of Civil Procedure in 1938, parties seeking information from their adversaries relied on cumbersome, indirect mechanisms, most notably bills of particulars, which served as quasi-independent proceedings. These methods were formalistic and poorly suited to reconstructing the factual circumstances underlying a dispute.[8]

The 1938 Rules fundamentally altered this landscape. By embedding discovery into the ordinary course of litigation and permitting parties to engage in discovery without prior court approval, the Rules transformed civil litigation into a process centered on the systematic reconstruction of facts about the dispute.[9] Requests for production, interrogatories, inspections, and depositions became instruments through which litigants could reassemble the factual world.[10]

This procedural shift, whether intended or not, relocated the center of litigation from pleading and trial to discovery. Trials became less arenas of surprise, revealing cross-examinations, and rhetorical performances, and more culminations of an extended factual investigation conducted by the parties themselves. The accumulation, organization, and

---

8. *See* Stephen N. Subrin, *How Equity Conquered Common Law: The Federal Rules of Civil Procedure in Historical Perspective*, 135 U. PA. L. REV. 909, 943–75 (1987) (explaining that the 1938 Federal Rules effected a deliberate merger of law and equity procedure, replacing formalistic pleading regimes with a unified system oriented toward factual development and judicial management); *see generally* Charles E. Clark, *The Handmaid of Justice*, 23 WASH. U. L.Q. 297 (1938) (defending the new Rules as instruments for reaching the merits rather than adjudicating procedural form).

9. Although the Federal Rules abolished separate equity procedure, they did not abolish equitable judgment. To the contrary, equity's flexible, context-sensitive orientation toward fairness and fact remained embedded within unified civil procedure. *See* Subrin, *supra* note 8, at 922–26 (describing how equitable modes of reasoning persisted after the merger).

10. Discovery was understood from the outset as a principal mechanism through which the Federal Rules' equitable, fact-sensitive orientation would operate in practice. As Judge Charles E. Clark, the principal architect of the Rules, consistently emphasized, the merger of law and equity and the abandonment of rigid pleading formalism were intended to shift adjudication away from technical barriers and toward the development of a full factual record prior to decision. *See* Michael E. Smith, *Judge Charles E. Clark and the Federal Rules of Civil Procedure*, 85 YALE L.J. 914, 927–33 (1976) (explaining Clark's view that liberal pleading and expansive pretrial discovery were necessary to ensure that cases were decided on their merits rather than on procedural traps, and that discovery and pretrial procedure were central to that project). Accordingly, discovery was not framed as an ancillary procedural device, but as a core instrument for revealing the factual context necessary for sound judgment.

interpretation of information increasingly shaped and revealed pertinent claims and defenses rather than merely supporting a highly formatted claim made at the beginning of the case. The drama of trial was replaced with pervasive factual grounding, revealing innumerable factual relationships. Discovery became a canvas, providing the colors and shapes to paint the client's claims and defenses.[11]

Litigation quickly evolved into a contest over information. While private investigations played a role, the principal battleground was formal discovery practice. Two competing axes governed discovery. First, discovery was expansive. Under the original formulation of relevance, parties could obtain any information reasonably calculated to lead to the discovery of admissible evidence.[12] Second, litigation's adversarial nature was embedded in the discovery process. Parties resisted overbroad requests through motions for protective orders, while requesting parties pressed forward with motions to compel. Although courts intervened episodically to restrict overbroad requests, discovery's overall structure favored broad information disclosure and relied heavily on counsel to manage, interpret, and present the resulting information.[13]

Within this framework, depositions quickly emerged as the central mechanism for factual reconstruction. Interrogatories, though formally answered under oath, were drafted by counsel and rarely yielded unexpected insight. Paper document production, largely limited to memoranda, correspondence, and contracts, typically consisted of materials already known to one or both parties and were often created with an awareness of potential future scrutiny. Depositions, by contrast, were major litigation events. Depositions allowed counsel to explore sequence, motivation, and causation. They also required extensive preparation by both counsel and witnesses. Key depositions often

---

11. Modern civil litigation increasingly takes shape during discovery rather than at trial. *See* FED. R. CIV. P. 16 advisory committee's note to 1983 amendment (explaining that pretrial management and discovery were intended to narrow issues, shape claims, and promote informed resolution prior to trial); FED. R. CIV. P. 26 advisory committee's note to 1993 amendment (emphasizing the central role of discovery in defining the factual scope of disputes and facilitating early evaluation and resolution). Courts and commentators have accordingly described trial as the culmination of a process in which factual narratives are largely constructed, tested, and refined during discovery rather than at the evidentiary hearing itself. *See also* Judith Resnik, *Managerial Judges*, 96 HARV. L. REV. 374, 402–13 (1982) (describing trial as the culmination of a process largely shaped through pretrial management and discovery).

12. *See* FED. R. CIV. P. 26(b) advisory committee's note to 1946 amendment (expanding the scope of discovery beyond admissible evidence and reflecting an expectation that counsel would bear primary responsibility for managing the development of the factual record).

13. *See* Stephen N. Subrin & Thomas O. Main, *The Integration of Law and Fact in an Uncharted Parallel Procedural Universe*, 79 NOTRE DAME L. REV. 1981, 1985–2001 (2004) (describing discovery as a largely lawyer-driven process outside of sustained judicial supervision).

unfolded over multiple days and served as the backbone for summary judgment strategy, trial cross-examination, and settlement leverage.[14]

Yet the primacy of depositions also revealed the significant limits of pre-digital discovery. Depositions depend on human memory, emotion, and interest. Witnesses forget, reconstruct, resist, and rationalize. The retrieval of information stored in a witness's mind is mediated by the witness's conscious strategy and unconscious biases. Even under ideal conditions, depositions are highly imperfect vehicles for recovering the texture of lived experience. Seemingly simple events such as a single telephone call required painstaking reconstruction: who was on the line, who was listening, who entered or left the conversation, and when. A substantial portion of deposition practice was devoted to establishing chronology, solidifying narratives, and narrowing factual ambiguity rather than uncovering new information.[15]

During the late twentieth century, discovery grew more contentious not merely because of expanding volumes of information driven by larger organizations and the proliferation of copy machines and printers, but because it lacked the consistent, moderating presence of the judge. Unlike trial, discovery unfolded largely outside the courtroom, supervised only intermittently through motion practice. By the time disputes reached the court, they were often hardened conflicts rather than manageable disagreements. Conduct that would never be tolerated in open court flourished in discovery's shadows, particularly at depositions. For many attorneys, discovery became less a means of uncovering foundational facts and more an effort to reveal nothing while burdening the opposition with escalating costs—a pathology the procedural system was ill-equipped to correct. Depositions were plagued by speaking objections that coached witnesses or encouraged selective memory. Meanwhile, interrogatories and requests for production were mired in boilerplate objections, such as that the requested discovery was overbroad, vague, ambiguous, or immaterial. The judiciary, wary of tipping the scales of

14. *See generally* John S. Applegate, *Witness Preparation*, 68 TEX. L. REV. 277, 298–323 (1989) (describing the role of depositions and other witness-preparation techniques in developing factual narratives and testing witness credibility in modern civil litigation); *see also* FED. R. CIV. P. 30(d)(1) advisory committee's note to 2000 amendment (recognizing depositions may extend beyond one day when necessary).

15. *See generally* John H. Langbein, *The Disappearance of Civil Trial in the United States*, 122 YALE L.J. 522, 549–52 (2012) (noting that modern civil litigation relied heavily on pretrial discovery and depositions to reconstruct facts, despite their cost and inefficiency); *see also* Steven S. Gensler, *Judicial Case Management: Caught in the Crossfire*, 60 DUKE L.J. 669, 671–77 (2010) at 671–77 (observing that case management reforms have shifted significant decisional authority into the pretrial phase).

justice, increasingly came to view discovery disputes as burdensome, cantankerous, perpetual battles often over marginal stakes.[16]

This was the pre-digital crisis of discovery: a system designed to recover the factual world of disputes, but strained by adversarial incentives, limited by human memory, and structurally detached from judicial presence. It is against this backdrop, not as a sudden rupture, but as a response to accumulated strain, that the digital transformation of discovery must be understood.[17]

B. *The Digital Turn*

The dramatically increased prevalence of ESI in litigation and the diminution of paper documents, which began at the beginning of this century, has fundamentally altered discovery. Suddenly, clients and counsel confronted a new electronic "information container." The "*Zubalake* Shock" was that a new sheriff had arrived in town.[18] This new sheriff was the new digital information format, previously barely recognized by attorneys, with its own demands, challenges, and risks. The litigation impact of the digital revolution became clear when the *Zubalake* cases clarified that "document" discovery meant the location and exchange of "invisible" digital files stored on innumerable devices. Suddenly, litigators were tasked with locating and exchanging digital files generated, transmitted, and stored on a plethora of devices that could be easily altered or obliterated, and that were coded and read by a variety of applications. Most importantly, this digital information was

16. *See* Steven S. Gensler, *A Bull's-Eye View of Cooperation in Discovery*, 10 SEDONA CONF. J. 363, 368–75 (2009) (describing cooperation methods to mitigate lawyer-driven discovery practices characterized by obstruction, boilerplate objections, and abusive deposition conduct); *see also* Lee H. Rosenthal, *From Rules of Procedure to How Lawyers Litigate: 'Twixt the Cup and the Lip*, 87 DENV. U. L. REV. 227, 233–40 (2010) (providing a history of discovery rules and emphasizing procedures under the rules to avoid hardened disputes and strategic excess before issues reached the court).

17. As synthetic content, strategic misinformation, and generative "AI slop" increasingly saturate public and commercial discourse, the judiciary's evidentiary gatekeeping function assumes renewed significance as a necessary defense of a shared factual world in which judgment remains possible. *See, e.g.*, Paul W. Grimm, Maura R. Grossman, & Gordon V. Cormack, *Artificial Intelligence as Evidence*, 19 NW. J. TECH. & INTELL. PROP. 9, 12–25, 88–89 (2021) (proposing a burden-shifting framework for the admissibility of AI-generated evidence grounded in reliability, transparency, and human accountability); *see also* Advisory Committee on Evidence Rules, Proposed FED. R. EVID. 707 (2023) (addressing the admissibility of evidence derived from machine learning systems). The proposal generated substantial public comment and institutional concern and has since been paused, reflecting judicial unease about both over-exclusion and under-regulation of synthetic evidence.

18. *See generally* Zubulake v. UBS Warburg LLC, 217 F.R.D. 309 (S.D.N.Y. 2003) (cost-shifting for electronic discovery); 220 F.R.D. 212 (S.D.N.Y. 2003) (duty to preserve electronically stored information); 229 F.R.D. 422 (S.D.N.Y. 2004) (sanctions, adverse inference, and counsel's preservation obligations); 231 F.R.D. 159 (S.D.N.Y. 2005) (attorneys' fees and cost calculation).

voluminous, and much of it was created "on-the-fly" without any attention to whether it could ever be recovered and used for litigation.

Humans have generated information with a wide variety of technologies. For example, the typewriter revolutionized writing. No longer must one tediously scratch letters in longhand, struggling to maintain legibility. All one needs is a typewriter (and the requisite skills and dexterity) to create perfectly readable text. Computers added various capabilities such as storage, but for the most part, computers were initially seen as production devices, typewriters on steroids. With the rise of the internet, computers ceased to function primarily as isolated tools for document creation and instead became nodes in a shared communications network. No longer does one print documents on paper and send them by ground to a recipient. Rather, electronic documents are sent to the recipient instantly, virtually anywhere in the world, via the internet. Today, apart from bills and formal notices, paper letters have nearly gone the way of the once ubiquitous slide rule. When the digital era fully emerged in litigation in the early years of the twenty-first century, the digital environment confronting litigators was unlike anything they had previously experienced. Whereas the modern discovery era was defined by the relative scarcity of paper documents, the digital age produced an explosion of evidence. But that explosion came at a cost: information was now stored in electronic files that required new methods of access and retrieval.[19]

This new format for discoverable information crashed against the training and habits of most litigators. Prior to the digital era, cases typically started with a collection of the relatively few relevant documents and interviews of few principal witnesses. The document collection was easily handled by a paralegal or a small paralegal team. Relevant documents were stored in cabinets; the cabinets had labeled drawers; the drawers had labeled folders; and the folders contained the paper documents. Locating the relevant case documents was relatively straightforward.

Unless preliminary relief was sought, the full collection of relevant documents was neither urgent nor central and was generally postponed until it became practically necessary. The complaint needed to be answered in twenty days, but enlargements were routinely requested and granted. Attacks on the complaint usually delayed discovery until the case was at issue. Finally, documents could also be requested. All this easily took months. Over the months typically consumed by pleadings and motions directed at the pleadings, paper documents posed little risk

19. *See* THE SEDONA CONFERENCE, THE SEDONA PRINCIPLES: BEST PRACTICES, RECOMMENDATIONS & PRINCIPLES FOR ADDRESSING ELECTRONIC DOCUMENT PRODUCTION 3–16 (1st ed. 2003) (recognizing that electronically stored information differs fundamentally from paper records in volume, persistence, mutability, and methods of access).

of inadvertent loss or degradation. Paper documents could be reviewed at the producing counsel's office, sent to a copying office, and reproduced. Each party controlled its own documents, usually at a limited number of locations. The limited number of documents was literally fenced and locked within the party's physical offices.

The digital era disrupted this comfortable dynamic. Electronic files increased exponentially and were no longer confined to handy file cabinets and drawers, but were dispersed across multiple "drives," systems, and platforms. The challenge of locating potentially relevant information was compounded by the fragility of digital files, which are routinely altered, overwritten, or deleted through ordinary use and automated retention practices. Delay became potentially dangerous and spoliation emerged as a major concern.[20]

At the same time, the expansion of third-party and cloud-based storage fundamentally altered the relationship between parties and their information. Documents were no longer housed solely within a party's physical offices or under its exclusive control, but distributed across email servers, cloud platforms, vendor systems, and globally networked databases. Information was no longer proprietary in the sense of being locked away in private cabinets or computers. Moreover, retrieving and decoding these materials was no longer a simple act. Each file was encoded in a particular format, accessible only through specific applications, and often embedded within poorly organized digital environments. For example, emails accumulated chronologically rather than by subject, and documents were deposited into sprawling folders or drives without meaningful structure. Locating key materials within these systems often felt like searching for the proverbial needle in a haystack, prompting demand for information management skills.

The dispersion and heterogeneity of ESI fundamentally altered the conditions under which factual reconstruction occurs in civil litigation. As business and personal activity migrated from discrete, locally controlled systems to distributed digital environments of shared servers, outsourced platforms, mobile devices, and cloud-based services, relevant evidence ceased to reside in any single, coherent repository. Instead, the factual record became fragmented across systems subject to different custodians, technical architectures, retention policies, and access permissions. No individual litigant, however diligent, could unilaterally identify, collect, and contextualize all potentially relevant data without cooperation from adversaries and non-party service providers. Discovery

20. *See, e.g.*, *Zubulake*, 220 F.R.D. at 216–18 (emphasizing that electronically stored information is uniquely vulnerable to alteration or destruction through routine use and automated systems); *see also* FED. R. CIV. P. 37(e) advisory committee's note to 2015 amendment (recognizing the distinctive preservation risks associated with ESI).

thus shifted from a unilateral act of production to a coordinated process of mutual orientation toward a dispersed informational field.

In this environment, cooperation is not merely a matter of civility or efficiency; it is a functional prerequisite for meaningful discovery. Identifying the universe of relevant information now requires shared disclosures about system design, data flows, and preservation constraints that are invisible from outside a party's digital infrastructure. Preservation decisions made in isolation risk either spoliation or disproportionate over-collection, and uncoordinated collection efforts can produce technically incompatible or context-stripped data sets that undermine reliability and usability. Digital discovery demands collaboration at the front end—through dialogue about sources, formats, scope, and sequencing—so that the parties can jointly reconstruct a factual record that no longer exists as a stable, pre-litigation whole. Cooperation emerges not as a concession to opposing counsel, but as the only viable response to the distributed nature of modern information.[21]

Unlike paper-era document review, where the relatively few case documents could be read by even the most senior partner, ESI requires search technologies to identify likely relevant documents located in large repositories of mostly noise. But searching was soon found not to be analogous to the near certainty gleaned from a professional reading a document; early search technologies merely predicted that a document might be relevant based on certain words, patterns, or other indicia, without reconstructing its context or meaning. These predictions were notoriously poor. Worse, attorneys were notoriously poor at using search techniques, which had never been taught at law school. Moreover, search compounded the attorney's distance from the information in the files because it required using search software. In short, the relatively simple process of document production in the paper era, fraught with its own dysfunctional dynamics, was replaced by a new world completely alien to most litigators' skills, talents, habits, and prejudices. At great cost, litigators soon began realizing that the cute games, tricks, and sharp practices of paper litigation, when applied to ESI, produced horrific results. Sanction cases mounted. New models were waiting to emerge.[22]

---

21. *See* The Sedona Conference, *Commentary on Legal Holds, Second Edition*: *The Trigger & The Process*, 20 SEDONA CONF. J. 341, 350–58, 386–87 (2019) (noting that modern information systems routinely involve third-party, cloud-based, and distributed storage, complicating traditional notions of possession, custody, and control); *see also* FED. R. CIV. P. 34 advisory committee's note to 2006 amendment (recognizing that electronically stored information may reside in multiple locations and systems beyond a party's immediate physical control).

22. *See* The Sedona Conference, *The Sedona Conference Cooperation Proclamation*, 10 SEDONA CONF. J. 331 (2009) (calling for cooperative discovery practices in response to the complexity and distributed nature of electronically stored information).

C. *The Rise of Necessary Collaboration*

The contrast between digital-era discovery and twentieth-century paper discovery is stark, and the resulting structural changes in litigation practice became equally apparent. In the digital environment, voluminous information must be culled and filtered to identify potentially relevant material. That process necessarily isolates digital artifacts from the context and connectivity that give them meaning. Relevant information is fundamentally relational.[23]

A text message, for example, is not a timeless object whose relevance can be predicted solely by the presence of particular words or patterns of data in the message. It acquires meaning only when its context is understood: why it was sent, when it was sent, and how it was received. What prompted it and what did it cause? Context is revealed not merely by the message itself, but in part by what the recipient did next. Context is established by the information possessed by both sender and recipient. In short, a full understanding of an email event depends on reconstructing its shared informational background.

Multiply this single interaction across thousands of different kinds of messages—emails, text messages, instant messages, and communications embedded in collaborative platforms—distributed across multiple systems, involving multiple recipients, producing cascading responses and subsequent actions, and it becomes clear that a client's story can be coherently constructed only when the relevant information is shared. This "collaborative" feature of modern discovery remains deeply counterintuitive to many "old school" practitioners: effective advocacy in digital litigation increasingly depends not only on access to one's own information, but also on the other party's documents and communications, as well as third-party data. This is not a claim about cooperation as an ethical ideal, but about the conditions necessary for persuasive advocacy in our digital informational environment.

The integrated nature of ESI follows directly from its character as a product of the computer as a communication device. Information does not exist in isolation; it is generated, transmitted, received, and acted upon for purposes embedded in human activity. Pre-digital litigators implicitly recognized this integration when they placed depositions at the center of discovery, seeking to reconstruct events through human custodians whose memories were partial, biased, and contingent. What the digital discovery era reveals is not that depositions have become obsolete, but

23. *See, e.g.*, The Sedona Conference, *Commentary on ESI Evidence & Admissibility, Second Edition*, 21 SEDONA CONF. J. 83, 93–104 (2020) (emphasizing that the meaning and evidentiary value of electronically stored information depend on context, relationships, and surrounding circumstances rather than isolated content); *see also* Grimm, Grossman, & Cormack , *supra* note 17, at 84–88 (2021) (explaining that artificially generated evidence requires knowledge of how it was produced to assess reliability).

that human testimony is no longer the sole—or even primary—means of reconstructing the connective tissue of events. When properly assembled, digital artifacts allow parties and courts to recover sequences, relationships, and actions that human recollection alone could never reliably reconstruct.

III. THE WORLDLINESS OF DIGITAL EVIDENCE

The human world is a shared domain. We are beings in the world in the sense that we share customs and traditions, inhabit common spaces, and express our cares and concerns with others. Lived human experience is not composed of self-contained objects moving independently in a magnetic field. Such a premise may be scientifically and analytically useful, but it obscures how meaning arises in human affairs.

Of course, we can adopt an objectifying attitude and treat people and things as having fixed locations at given moments in a multidimensional space, defined by primary and secondary qualities. I can objectify and locate my dog next to me according to Cartesian coordinates and belonging to the species *Canis lupus familiaris*, which is golden-colored and responds to specific commands. But any dog owner knows that the lived relationship with their dog—formed through shared routines, expectations, affection, mutual understanding, and history—is something quite different. To treat the digital record as a mere collection of "simple locations" is to commit what Alfred North Whitehead termed the "fallacy of misplaced concreteness," mistaking the abstract for the reality of lived experience.[24]

I share a world with my neighbors. We look out onto the same street. The children playing there belong to neighborhood families. We walk the same sidewalks, work together clearing debris after storms, and talk about the neighborhood and other matters. Yet we experience this same cul-de-sac differently. We recognize one another, but we bring different perspectives, memories, and expectations to the same shared place. The older neighborhood couple watches generations of new families move in and remembers their children playing basketball in the center of the circle. The new family welcomes new friends, new children for their children to play with, new school communities, and new opportunities. The neighbors celebrate in the circle with dramatic events of fireworks, song, and drink. The neighborhood is a human world full of meaning, goals, and relationships—a space where people appear, speak, and act in one another's presence. It is their neighborhood full of meaning. A common expression is "Remember when…." The neighborhood has

24. ALFRED NORTH WHITEHEAD, SCIENCE AND THE MODERN WORLD 53–55 (1925) (critiquing the Cartesian and Lockean doctrines of "simple location" and "secondary qualities" as the "fallacy of misplaced concreteness," which occurs when one mistakes an abstract, mathematical representation of reality for the concrete reality of experience itself).

stories of success on the high school basketball team, of job transfers, of neighbors lost, of marriages and divorces, of hospitalizations, and of enduring illnesses. The neighborhood exemplifies decency, friendship, and collective commitment that allow human judgment about what is important and what is wrong. These stories generate memory, continuity, and a shared sense of reality that is a past, present, and future. To withdraw from this common life, whether by choice or circumstance, is to lose access to the world as a lived experience. To step away from participation is not merely to change location, but to lose one's place in the world as it is lived and understood. This was the insight of Socrates, who chose hemlock rather than accept banishment from the Athenian *polis*. A life extracted from the shared stories and public presence of one's fellow citizens is a life severed from the very possibility of meaning and judgment.[25]

Lived experience is temporal, a "no longer" and a "not yet," notwithstanding the seemingly endlessly present demands of our current cultural environment and its endless distractions. The contemporary world has taken on a frantic horizontal pace. Events seem to unfold all at once across an ever-expanding present, rather than organically within a past, present, and future. Life can begin to appear not as an unfolding set of sustained actions and activities, but as a cacophony of discrete, jarring, transitory events. As Fredric Jameson observes, this shift toward a "series of pure and unrelated presents" weakens our access to the historical depth required for narrative. Meaning recedes. Memory weakens. Events that should stand out and endure fade quickly in the crush of immediacy.[26]

The work of discovery is to recover the loss of the everyday world of lived experience that has become occluded and, in a significant sense, vacated by the loss of meaning. The actual world appears only as a shrouded, unconscious background to the immediacy of our lives. But what might appear in the present as meaningless digital debris becomes, in litigation and electronic discovery, the material through which our actual world is recovered. For example, an employee does not receive a promotion. The employee feels aggrieved. The supervisor feels justified.

---

25. *See* HANNAH ARENDT, THE LIFE OF THE MIND: THINKING 166–93 (1978) (discussing Socrates as the paradigmatic example of the "two-in-one" of thought and the citizen who refuses to be alienated from the shared world of the *polis*); *see also* HANNAH ARENDT, *Philosophy And Politics*, *in* THE PROMISE OF POLITICS 5–39 (Jerome Kohn ed., 2005) (explaining Socrates' commitment to the *doxa*—the perspectives—of his fellow citizens as the foundation of political life); HANNAH ARENDT, *Personal Responsibility Under Dictatorship*, *in* RESPONSIBILITY & JUDGMENT 17–48 (Jerome Kohn ed., 2003) (arguing that judgment is the faculty that resists "automatic" processes and that the individual must maintain their "place in the world" to remain a responsible moral agent).

26. *See* FREDRIC JAMESON, POSTMODERNISM, OR, THE CULTURAL LOGIC OF LATE CAPITALISM 25–31 (1991) (describing postmodern culture as present-oriented and fragmented, with diminished connection to shared historical narratives).

Each has a perspective, and the event is composed of facts. What was the employee's performance? What training was provided? What did the human resources department say and do? What did the employee post on social media? What is the trajectory of the company and the career arc of the employee? How did the employee and their colleagues interact? What explanation has the supervisor offered? Electronic discovery is uniquely positioned to lay the factual foundation that sustains these competing perspectives. The actions and reactions of participants are reflected in the manifold of digital artifacts surrounding them. Information drawn from multiple sources, dense with detail and meaning, makes possible the reconstruction of lived experience. Electronic discovery places events on their own stage, allowing judgments to be formed about the conduct of the actors. This common foundation exists because we live in one world with clashing perspectives that persuasion seeks to order. Litigation tests differing judgments about the same shared reality.

Electronic discovery is well-suited to this reconstructive effort given the volume and richness of modern data. Data collected from multiple sources supplies the context necessary for reconstruction. Electronic discovery reveals not only the human-generated content of digital files, but also the metadata that situates those artifacts in time, place, and relationship. Metadata captures dates, authors, file paths, and other contextual markers. Rather than relying on decontextualized documents and often unreliable testimony to establish time, place, and meaning, the content and metadata of digital artifacts anchor and guide this interpretive work. For this reason, preserving and collecting key metadata, as well as document content, is critically important.[27] Modern e-discovery platforms do not merely index the text content of documents, but also rely on metadata in the construction timelines, relationship maps, and other contextual frameworks that assist interpretation.

While electronic discovery allows us to come closer to reconstructing the factual world, we must recognize the limits of this recovery. Electronic discovery can bring us closer to the factual record, but it remains a reconstruction. The point is not that there is a perfect reconstruction, as if we could get there if we worked hard enough. In the human world, there are only persuasive perspectives of the meaning of things. We cannot import into the human world of meaning an "objective" standard more appropriate to scientific investigation. A purely scientific reconstruction treats the digital record as a collection of fixed points in space and time, but it cannot breathe life back into the

27. *See* The Sedona Conference, *The Sedona Principles, Third Edition: Best Practices, Recommendations & Principles for Addressing Electronic Document Production*, 19 SEDONA CONF. J. 1, 172–77 (2018) (Principle 12 explains that metadata provides essential contextual information for understanding electronically stored information.).

relationships and intentions that gave those points meaning then and for us now.

Facts require interpretation, or more precisely, they call for judgments about meaning, responsibility, and right and wrong. Judgment, in this sense, concerns the evaluation of acts, the assessment of responsibility, and the determination of what ought or ought not to be done. Attorneys, courts, and juries all engage in such judgments, though at different moments and with different authority. Of course, in addition to the kind of thinking associated with judgment and meaning, litigators also engage constantly in instrumental thinking —what means are available to get to a particular end. They decide how to sequence depositions, when to file motions, how to staff a review team, and how to manage time and cost. These are questions of means, not meaning. Instrumental reasoning is essential to litigation, but it does not determine what a set of facts signifies, whether conduct is wrongful, or how responsibility should be assigned. Judgment is therefore not a matter of efficiency or strategy. It is neither purely logical nor deductive, nor is it mathematical cognition. Cognition tells me that if two angles of a triangle add up to one hundred twenty degrees, the remaining angle must be sixty degrees. Judgment, by contrast, asks what this act means in this situation, and whether it can be justified to others. It requires a temporal perspective drawing on memory, experience, and examples, as well as imagining the perspectives of others. Judgment must withstand the test of publicity, that is, whether others might reasonably agree.[28] Judgment depends on a shared factual world, because without common facts, there is nothing about which judgment can meaningfully occur.[29]

The common law is guided by exemplars, by reasoned decisions that explain how the court understands the facts before it. Those judgments inform future judging, which must nonetheless think anew about different facts, in different contexts, and at different times. Examples provide guidance, but that guidance is never dictatorial or compulsory.[30] No case

---

28. Hannah Arendt distinguishes judgment from detached observation by emphasizing the responsibility of the spectator. Judgment does not arise from an Archimedean standpoint that purports to see everything at once, but from the more demanding task of seeing a shared world from multiple perspectives without guarantees. *See* ARENDT, LECTURES ON KANT'S POLITICAL PHILOSOPHY, *supra* note 5, at 40–46, 67–72, 119–22 (describing judgment as the activity of a spectator who must "think from the standpoint of everyone else" without recourse to rules or certainties). Arendt insists that judgment remains grounded in shared appearance rather than Cartesian modes of abstractions or claims to universal perspective. *See id.*

29. *See* LINDA M. G. ZERILLI, A DEMOCRATIC THEORY OF JUDGMENT 38–43, 117–42 (2016) (interpreting Arendtian judgment as grounded in shared appearance and plurality rather than subjective preference or rule application).

30. *See* EDWARD H. LEVI, AN INTRODUCTION TO LEGAL REASONING 1–8 (1949) (describing common law reasoning as the use of examples that guide without compelling future decisions and emphasizing the responsibility of judges to address novelty in each case).

is decided simply because it is "on all fours" with another. Every judgment carries responsibility and must address novelty.

The special mission of electronic discovery is to envelop litigation in the relevant factual world, thereby laying the predicate for judgment. Litigation is not a process in which one side presents "its facts" and the other offers "alternative facts." It is an adversarial process in which both sides operate from a shared factual foundation, offering competing perspectives, persuading, and seeking an authoritative public judgment. Judging is a public act of disclosure, allowing what the facts reveal to appear.[31] Litigation has long functioned not only as a formal adjudicatory process, but as one of the few remaining structured spaces in which contested facts must appear, be tested, and be publicly accounted for, bounded by evidentiary rules that distinguish persuasion from assertion and truth-seeking from mere narrative, rhetorical force.

IV. FROM VOLUME TO RECONSTRUCTION

The practices that make judgment possible are fragile. They can be preserved, but they can also be displaced—often in the name of efficiency, objectivity, or progress.[32] From the outset of the e-discovery era, the sheer volume of ESI began to threaten the conditions under which judgment operates. Manual human review of vast data sets was no longer feasible. The data deluge had arrived, and with it a race toward increasingly efficient methods of culling and review. Modern litigation unfolds across sprawling networks of email, messaging platforms, shared drives, cloud services, collaboration tools, and structured databases, each with its own metadata and technical architecture. Virtually every dispute now exists across multiple, interlocking data systems, fragmenting the factual record even as it expands it.

Even in the early years of this century, collecting three to four gigabytes of data per custodian was common.[33] Printing a single gigabyte of text data could yield twenty to thirty Bankers® boxes of paper. At the same time, relevant information was typically buried in noise. Ten percent of an employee's email may bear on any particular dispute and identifying that ten percent is extraordinarily difficult. Litigation's short history of electronically searching gigabytes, and even terabytes, of data for relevant documents, often described as information recall, is not a happy story. Early electronic discovery methods relied heavily on search

31. *See* ARENDT, BETWEEN PAST AND FUTURE, *supra* note 3, 50–58 (1958) (describing the public realm as a space of appearance in which actions and facts are disclosed and judged).

32. *See generally* ARENDT, BETWEEN PAST AND FUTURE, *supra* note 3, at 197–226 (arguing that mass society transforms cultural objects into consumable entertainment, thereby eroding the conditions for disinterested judgment).

33. *See, e.g.*, Paul & Baron, *supra* note 1, ¶¶ 8–12 (2007) (discussing the rapid growth of electronically stored information and the infeasibility of traditional review methods).

terms. Search terms suffer from an inherent contradiction: the broader the term, the greater the likelihood of capturing relevant documents, but also the greater the volume of false positives. In short, increased recall comes at the expense of precision and an increased volume of subsequent human eyes-on review. Increasingly elaborate refinements sought to solve this paradox. E-discovery professionals deployed Boolean operators, fuzzy logic, wildcards, stemming and lemmatization, proximity searches, and de-duplication. Email threads were preserved; parent-child relationships were maintained. Nonetheless, teams of attorneys were required to review the documents on their screens, deciding one by one whether each document was relevant. E-discovery reviewers typically focused on a single binary question: responsive or not. Context, when considered at all, was used primarily for efficiency, for example, assigning reviewers by custodian or document type, but not for reconstruction. The outcome of this process was not a narrative or a story. It did not aim to recover lived experience. It produced classifications, not meaning.

After the first decade of electronic discovery, technology-assisted review (TAR) famously emerged, heralded by the *Da Silva Moore* cases. Increased computing power enabled software to be trained to predict relevance by comparing documents in the collection to sets of previously identified responsive documents. Thus emerged the familiar mantra of "find more like this." Early TAR systems relied on expert-generated document "seed sets" to train the software model. Once training was deemed sufficient through testing against control sets, the entire document population was scored for likelihood of relevance.[34]

Several years later, TAR 2.0 introduced a modified training method.[35] Rather than relying on a fixed expert-created seed set, TAR 2.0 systems learned continuously from reviewer decisions and dynamically re-ranked documents from more to less likely to be relevant as the review progressed. Despite this technical advance from keyword searching, TAR, whether in its 1.0 or 2.0 form, remained focused on the same binary endpoint: responsive or not. TAR assigns a numerical likelihood of responsiveness but review still culminates in a yes-or-no responsiveness determination, largely divorced from the document's broader context, origin, and meaning. Large teams of reviewers moving through hundreds

34. *See* Maura R. Grossman & Gordon V. Cormack, *Technology-Assisted Review in E-Discovery Can Be More Effective and More Efficient Than Exhaustive Manual Review*, 17 RICH. J.L. & TECH. 11, ¶¶ 2–4 (2011) (describing a TAR process where human review of documents is used to train a system that then ranks the remaining document population according to recall and precision metrics).

35. *See* The Sedona Conference, *The Sedona Conference TAR Case Law Primer, Second Edition*, 24 SEDONA CONF. J. 1, 16–17 (2017). TAR was a significant advancement in search methodology and made proportional review practicable. In favorable cases, reviewing a minority of documents was sufficient to identify a substantial share of relevant and important material. Review could then be curtailed once marginal returns no longer justified additional cost.

of gigabytes of data, marking documents responsive or not, do not reconstruct lived experience.[36] Generative artificial intelligence altered this dynamic.

A. *AI Review and the Waning of Human Document Review*

For most of the electronic discovery era, document review functioned as the principal drag on litigation. It was labor-intensive, expensive, and slow, consuming disproportionate resources. Regardless of how documents were collected, searched, or ranked, they ultimately passed through human hands. That condition is rapidly disappearing. AI can now perform the core tasks that once defined document review. It reads text and applies defined legal criteria consistently across very large collections. It does not tire, lose focus, or vary in judgment from one document to the next. Once a review protocol is defined, the system applies it identically. As a result, large-scale human review is no longer economically or operationally defensible in most matters. Industry data already reflects this shift, with traditional contract review volumes declining sharply as AI systems complete reviews faster, at lower cost, and with greater consistency than human teams.[37]

This development does not eliminate document review. It changes how review works. Review will no longer consist of large teams of attorneys clicking through documents making binary coding decisions. Instead, in AI review, attorneys design AI prompts, refine instructions, test outputs, and validate results. In many respects, this structure mirrors traditional review practice, where case litigators developed protocols, trained reviewers, and supervised their work. The difference is that the reviewers are no longer human.

Once AI review is complete, the real engagement with the digital artifacts begins. The task is no longer filtering noise. It is reconstructing events, sequences, and relationships and assessing what they reveal. Attorneys interrogate the data, test competing accounts, and assemble a persuasive account of the shared factual world. In Arendt's terms, this is the point at which judgment requires an enlarged mentality, the disciplined effort to consider events from multiple perspectives grounded in a common reality. AI can summarize key documents, generate reports, and answer substantive questions. It allows attorneys to engage in

---

36. *See* Xavier Rodriguez, *Artificial Intelligence (AI) and the Practice of Law*, 24 SEDONA CONF. J. 783, 795–806 (2023) (discussing the ethical and practical implications of AI in legal workflows).

37. John Tredennick, *Five Reasons Why Your Best Document Reviewer Isn't Human Anymore*, EDRM (Jan. 26, 2026), https://edrm.net/2026/01/five-reasons-why-your-best-document-reviewer-isnt-human-anymore/ [https://perma.cc/6LFJ-85G4] (reporting industry adoption of AI-based review systems and a corresponding decline in reliance on large-scale human document review teams).

continuing "conversations" with the digital artifacts. AI allows a "deposition" of the data. Through iterative and probing inquiry, attorneys can explore sequences of events, identify key actors, and test emerging accounts of what occurred. What is learned through this process does not substitute for the lawyer's engagement with the case. Rather, it assists the lawyer in situating the digital materials within the broader factual story of the dispute, including testimony, context, and judgment that extend beyond the documents themselves.[38]

For example, AI can identify the most important individuals involved in a dispute, when their involvement began, and what actions they took.[39] It allows iterative prompts (follow-up questions and requests) that continuously illuminate the world disclosed by the digital artifacts. In this way, AI begins to reveal the space in which people appeared, acted, and interacted, including when and how those interactions unfolded. AI helps attorneys find their voice in telling a persuasive client story, test hypotheses, and explore lines of inquiry. AI can "stress test" what an attorney might have early, blithely, and inaccurately conclude was a compelling story. The attorney remains in control, but the data is now deployed to reconstruct and challenge perspectives of the factual record. AI can illuminate the evolution of a dispute, the roles of its participants, and the sequence of events, and it can surface patterns and perspectives that might otherwise go unnoticed.

Document review is not the site of judgment. Review classifies; advocacy interprets. AI-assisted review establishes the shared evidentiary world by organizing disclosed materials according to human-defined criteria. At a subsequent stage, AI systems can facilitate human judgment, for example, by exposing multiple perspectives within the record, enabling the exercise of what Arendt described as enlarged mentality. Once discovery is understood as the progressive disclosure of a single evidentiary world, the distinction between "my documents" and "your documents" loses analytic force. Party documents, witness

38. *See* Craig Ball, *The Leery Lawyer's Guide to AI and LLMs in Trial Practice*, BALL IN YOUR COURT (Jan. 14, 2026), http://www.craigball.com/Leery_Lawyers_Guide_to_ AI_2026.pdf [https://perma.cc/3M2G-7BZF] (explaining how attorneys can use LLMs to query, summarize, and "chat" with large document sets to identify key facts and contradictions); *see also* Justin Smith, *Five Ways Generative AI Is Reinventing Modern Litigation Workflows*, EVERLAW (Apr. 14, 2026), https://www.everlaw.com/blog/ai-and-law/five-ways-generative-ai-is-reinventing-modern-litigation-workflows/ [https://perma.cc/K798-7TK8] (explaining that AI enables attorneys to synthesize large datasets into a coherent narrative of events).

39. *See* Alvarez & Marsal, *AI in eDiscovery*, https://www.alvarezandmarsal.com/expertise/ai-powered-ediscovery [https://perma.cc/9675-9PTN] (explaining that AI tools can identify key individuals and generate insights into the "who-what-when" facts of a case); *see also* Maura R. Grossman et al., *The GPTJudge: Justice in a Generative AI World*, 23 DUKE L. & TECH. REV. 1, 27 (2023) (noting that GenAI can help pro se litigants identify claims and draft pleadings tailored to their circumstances).

communications, and system-generated records are not separate categories but elements of a unified corpus reflecting the same underlying events. From an advocacy perspective, ownership is beside the point. At this production layer, AI-assisted review consolidates the shared world of disclosed information, while judgment remains a downstream human responsibility.

AI does not replace the litigator. Its outputs are provisional, dependent on prompts, shaped by training data, and never themselves accountable. The attorney must understand the craft of using AI and always bears responsibility, without excuse, for judgment and advocacy.[40] But subject to these limits, generative AI represents a significant step toward recovering the depth that has been missing from modern discovery practice and fulfilling e-discovery's early promises and aspirations.[41]

V. TECHNOLOGY AND RESPONSIBILITY

New technologies often arrive accompanied by misunderstanding and misuse. Technology assisted review produced nearly a decade of disputes over such esoteric matters as stabilization and seed sets. The temptation to misuse technology is especially strong in an adversarial litigation environment already marked by suspicion and the pursuit of tactical advantages.[42]

The most publicized litigation issue associated with the emergence of generative AI has been the perceived prevalence of citation "hallucinations."[43] The paradox of citation hallucinations is that they are

---

40. *See* HANNAH ARENDT, THE LIFE OF THE MIND 53–65 (1978) (analyzing the distinction between thinking and knowing and the inability of logic to generate meaning); *see also* ARENDT, LECTURES ON KANT'S POLITICAL PHILOSOPHY, *supra* note 5, at 42–44 (linking judgment to publicity and the standpoint of others). A complementary warning appears in Günther Anders's account of "Promethean Shame." *See* Günther Anders, *On Promethean Shame*, *in* CHRISTOPHER JOHN MÜLLER, PROMETHEANISM: TECHNOLOGY, DIGITAL CULTURE AND HUMAN OBSOLESCENCE, 29–95 (Christopher John Müller trans., 2016) (providing the seminal translation of Anders's analysis of technological self-subordination and the desire to become as "perfect" as one's machines).

41. *See* Maura R. Grossman & Gordon V. Cormack, *Comments on "The Implications of Rule 26(g) on the Use of Technology-Assisted Review*," 7 FED. CTS. L. REV. 286, 300–03 (2014) (arguing that attorneys satisfy their ethical duties not by reviewing every document, but by verifying the reasonableness of the automated process through statistical validation).

42. The procedural aftermath of *Da Silva Moore v. Publicis Groupe*, 287 F.R.D. 182 (S.D.N.Y. 2012), illustrates the profound misunderstanding and disputes that can occur when litigators attempt to deploy unfamiliar technologies in this adversarial environment. Rather than exploiting the power of the technology, the *Da Silva Moore* litigation devolved extensive disputes over the seed set, the methodology of the software itself, and the fairness of the judge. *See, e.g.*, Moore v. Publicis Groupe, 868 F. Supp. 2d 137 (S.D.N.Y. 2012) (denying a motion to recuse the judge).

43. While graphic, the term "hallucination" is misleading. Generative AI does not "perceive" or "imagine"; it predicts. A "hallucination" is an off-the-mark probabilistic prediction

extraordinarily easy to detect, since one needs only check the cited authority. The early hallucination cases arose largely because attorneys unfamiliar with how AI systems operate could not accept that a computer could err so confidently on something so basic as a case citation. A legal citation offered by powerful, loquacious AI does not appear to an inexperienced attorney as a constructed prediction that may or may not be accurate. More recent hallucination cases often involved breakdowns in human workflows. For example, an attorney provides a research prompt to an AI and forwards the draft output, assuming the receiving attorney will verify the AI's citations. The receiving attorney assumes the sending attorney would never forward unchecked citations. Erroneous citations then appear in a filed memorandum, and sanctions typically follow.[44] These errors are readily addressed through internal workflow controls that require citation verification in all filings, regardless of deadline pressure or litigation exigency.

We are accustomed to conceiving of computers as sophisticated calculators, a term derived from the human "calculators" of the nineteenth and early twentieth centuries. For that purpose, the computer is infallible. No one worries that a computer will produce the wrong answer to the question 2+2. The answer is always 4. This inheritance misleads us when we move to generative AI. AI does calculate, but its calculations result in predictions. Because these outputs arise from statistical weights and probabilistic distributions, they can be "off the mark." While AI can accurately summarize documents, its fluent and confident ability to synthesize data creates the danger that its predictive outputs will be taken uncritically.

A deeper and more enduring concern exists that remains even when the AI is technically accurate. This danger is the conflating of prediction and judgment. Even when used skillfully to test ideas or simulate perspectives, AI does not present a thinking judgment. AI does not think because it has no perspective of its own. Within the machine, there is nothing analogous to the human sense of "no longer" and "not yet." AI operates entirely in an ever-present now, applying statistical descriptors,

---

arising from the statistical weights of the model's training data. *See* Adam Tauman Kalai et al., *Why Language Models Hallucinate* (Sep. 4, 2025), https://cdn.openai.com/pdf/d04913be-3f6f-4d2b-b283-ff432ef4aaa5/why-language-models-hallucinate.pdf [https://perma.cc/RFB6-X73W] ("Hallucinations need not be mysterious—they originate simply as errors in binary classification. If incorrect statements cannot be distinguished from facts, then hallucinations in pretrained language models will arise through natural statistical pressures.").

44. *See* William Hamilton & Ralph Artigliere, *Reasonable or Overreach? Rethinking Sanctions for AI Hallucinations in Legal Filings*, EDRM (Aug. 18, 2025), https://edrm.net/2025/08/reasonable-or-overreach-rethinking-sanctions-for-ai-hallucinations-in-legal-filings/ [https://perma.cc/FLY8-2YMG] (arguing that most hallucination incidents reflect failures of human review and workflow design rather than inherent defects in generative AI systems).

vectors, and embeddings to generate striking outputs. But judgment requires more than the recombination of circumstances or the echoing of perspectives. It requires what Hannah Arendt identified as the "two-in-one" of internal dialogue, the capacity to be in a relationship with oneself and to answer for one's actions.[45] Because the AI cannot "appear" to itself or others as a person, its outputs can never be more than sophisticated data processing. Judgment is the exercise of a human faculty that no amount of predictive power can replace.

Generative AI functions as a machine for reproducing habits. It ingests the statistical "mores" of the legal past to predict text in the legal present. In Arendtian terms, it operates entirely within the realm of repetitive behavior, reproducing responses dictated by its training data. But the essence of judgment is the confrontation with the new, with what has not yet been decided. Because AI is tethered to the statistical mean of what *has happened*, it cannot recognize the singular event that breaks the pattern. It cannot be surprised. It lacks the "two-in-one," the silent partner of conscience, the ability to step back, ponder, and think of the meaning of things.[46] AI rushes forward, driven by the rolling momentum of probability. It does not pause. It requires a human attorney to provide the necessary caution, to interrupt the habit, to recognize the novelty of the specific case, and say, "Not here. Not this." The attorney provides the pause that allows judgment to enter the room.

Judgment is a human labor of discipline and conscience that no amount of predictive power can replicate. We must be wary of the illusion that the hard, albeit rewarding, work of judgment can be skipped or that AI can produce anything instantly to meet immediate demand. This illusion gives rise to what is now called "AI slop," content that is flat, deadening, and devoid of meaning.[47] The irony is that, when used

---

45. *See* HANNAH ARENDT, *Some Questions of Moral Philosophy*, *in* RESPONSIBILITY & JUDGMENT 71, 90–91 (Jerome Kohn ed., 2003) (explaining that the "two-in-one" is the condition of being a person who can be held responsible for their actions).

46. *See generally* ARENDT, THE LIFE OF THE MIND: THINKING, *supra* note 25. Arendt distinguishes judgment from rule-following cognition and habit, describing judgment as a reflective activity oriented toward particular situations rather than generalizable rules. She emphasizes the human capacity to respond to what is new, and the "two-in-one" of thinking as an internal dialogue that grounds conscience and responsibility. *See* ARENDT, THE LIFE OF THE MIND: THINKING, *supra* note 25, at 179–93 (thinking and the "two-in-one"); *id.* at 213–16 (judgment as reflective and worldly). Arendt contrasts judgment with behavior and habit, warning that systems which operate solely through rule application or statistical regularity can reproduce outward conformity without responsibility. *See also* HANNAH ARENDT, EICHMANN IN JERUSALEM 135–37 (rev. ed. 1965) (distinguishing judgment from unthinking compliance).

47. *See* Ella Feldman, *Merriam-Webster's Word of the Year for 2025 Is 'Slop,' the A.I.-Generated Junk That Fills Our Social Media Feeds*, SMITHSONIAN MAG. (Dec. 17, 2025), https://www.smithsonianmag.com/smart-news/merriam-websters-word-of-the-year-for-2025-is-slop-the-ai-generated-junk-that-fills-our-social-media-feeds-180987887/ [https://perma.cc/DHL6-XLDX] (defining "AI slop" as digital content of low quality produced in quantity by artificial

properly, AI can help attorneys to reconstruct stories, to grasp temporality, to understand the perspectives of others, and to prepare for judgment—that is, in thinking with what Hannah Arendt described as an "enlarged mentality."[48] When misused, it becomes yet another contributor to distraction, clutter, and the erosion of meaning.

VI. THE RECONSTRUCTION OF ADVOCACY

The electronic discovery era introduced a structural division within litigation practice. E-discovery teams responded to requests for production, ingested productions from opposing parties and third parties, and managed data repositories. Litigation and trial attorneys, by contrast, strategized cases, drafted and argued motions, and occasionally tried cases.[49] The split arose because trial and litigation attorneys were steps removed from the data. E-discovery teams supplied data when requested, for example, by producing emails or documents for a particular deponent, but the trial team's engagement with the data was largely derivative. Cases were built and tried without sustained engagement with the rich digital artifacts capable of reconstituting lived experience.

The typical document review process reinforced this separation. Reviews were conducted in large warehouse-like spaces, with each reviewer stationed at their own computer. Tens or dozens of reviewers were trained to identify relevant documents. Reviewers were organized into small teams supervised by a senior reviewer, who answered questions and relayed information deemed important by review managers. Documents were assigned to reviewers in batches and "checked out" from the collection of documents being reviewed. Reviewers clicked through documents, coding each as relevant or not, occasionally flagging privilege or tagging broad issues such as contract formation or interpretation. Review speed was closely monitored. Depending on the kind of data (emails, text messages, long documents, spreadsheets), reviewers could be expected to code a document per minute. After completing a batch, the documents were checked back in and replaced with another batch delivered to the reviewer. Team leaders and managing attorneys tracked metrics throughout the day and night, running quality-control checks. Efficiency and speed were constant

---

intelligence).

48. ARENDT, LECTURES ON KANT'S POLITICAL PHILOSOPHY, *supra* note 5, at 42–43.

49. *See generally* MICHAEL I. QUARTARARO, PROJECT MANAGEMENT IN ELECTRONIC DISCOVERY: AN INTRODUCTION TO CORE PRINCIPLES OF LEGAL PROJECT MANAGEMENT AND LEADERSHIP IN EDISCOVERY (2d ed. 2021) (providing a practical account of e-discovery workflows and team structures). While such accounts emphasize efficiency, coordination, and risk management, they do not address the deeper epistemic consequences of separating litigation strategy from sustained engagement with the underlying digital record.

imperatives.[50] Doubling review time doubled review costs. The objective was a defensible production of properly coded documents. While this process was optimized for defensibility and speed, it necessarily avoided narrative, causation, and meaning. The goal was not the meaningful reconstruction of human experience. Defensibility, not narrative coherence, became the governing standard.[51] For the most part, e-discovery review was a grueling, largely outsourced process, with junior members of the case litigator team often relegating the review oversight.

Although litigators had access to documents produced by and to opposing parties, those materials were not organized by the opposition in the production to create an integrated story. More importantly, it was nearly impossible for litigation teams to apprehend the depth embedded in tens of thousands of documents and their associated metadata within a review designed for classification rather than comprehension. The result was that trial teams tended to focus on key documents identified by clients and leverage these documents through deposition, cross-examination, and trial. Immersion in the full data environment was rare.

Advocacy, however, now requires more. The volume and dispersion of modern data mean that lived experience can be reconstructed only when information is aggregated and understood as an integrated web. Emails emerge from issues discussed in Teams messages, developed in PowerPoint presentations, and reflected in cloud-hosted databases accessed by multiple stakeholders at different times. For e-discovery, there is no meaningful distinction between "my" data and "your" data. There are only shared digital artifacts that together form the foundation for persuasive judgment. These artifacts can confirm or undo once-

---

50. *See generally* HERBERT MARCUSE, ONE-DIMENSIONAL MAN (1964); FREDRIC JAMESON, POSTMODERNISM, OR, THE CULTURAL LOGIC OF LATE CAPITALISM (1991); *see also* MARK FISHER, CAPITALIST REALISM: IS THERE NO ALTERNATIVE? 39–53 (2009) (analyzing the tendency of bureaucratic metrics and audit culture to displace substantive professional judgment). These works describe the substitution of administered routines, metrics, and behavioral compliance for substantive judgment, a dynamic reflected in large-scale document review practices organized around speed, defensibility, and process validation rather than reason, meaning, or interpretive understanding.

51. *See generally* The Sedona Conference, *Commentary on Achieving Quality in the E-Discovery Process*, 15 SEDONA CONF. J. 265, 271–74 (2014) (observing that large-scale document review has historically been oriented toward defensible production rather than substantive understanding of the data corpus). Modern document review developed primarily in response to adversarial requests for production, incentivizing reviewers to classify individual documents accurately rather than to understand relationships within the producing party's own data. In this posture, meaningful engagement with the corpus was often viewed as inefficient and outside the immediate objectives of production review, given cost pressures and workflow constraints. The paradox was that producing parties frequently gained little substantive understanding of their own documents during review, while trial teams later relied on downstream tools such as timelines, clustering, or contextual visualization to reconstruct meaning from both sides' productions. In practice, these tools were often underutilized due to time constraints and the press of litigation.

plausible narratives. We inhabit a shared world of actions, practices, customs, and mores. Digital artifacts are the durable traces of our participation in that world. They are not merely data points. To understand ESI is to recognize it as the evidence of a lived reality that is constantly being woven and re-woven through human action.

E-discovery compels adversaries to reconstruct a common factual world within which effective advocacy becomes possible. With AI, the litigator can now engage directly with the data, questioning it, probing it, and assessing responses in an exploratory dialogue. Litigation strategy and judgment emerge through this engagement, which has become a craft the modern litigator must understand and master. Litigation judgments require dwelling with the data long enough for its patterns, tensions, and silences to disclose meaning as a clash of human narratives and perspectives. AI allows the litigator to work with a rich canvas, recovering and restoring depth, perspective, and meaning.

VII. WHAT E-DISCOVERY REVEALS AND WHAT IT DEMANDS

The story today of electronic discovery is the story of a promise come home. From the outset, the emergence of digital information promised a recovery of depth that had long been frustrated by traditional paper document and deposition discovery. The digital archive appeared to offer access to the lived texture of events: how decisions were made, how understandings evolved, and how actions unfolded over time. However, that promise proved difficult to realize.

Each phase of the electronic discovery process presented its own challenges. Attorneys resisted, adapted, and at times gamed the e-discovery processes articulated by EDRM and The Sedona Conference. Keyword searching revealed an inherent tension between precision and recall. Expansive collections of ESI strained resources and fostered a kind of institutional schizophrenia, separating document review teams from litigation and trial attorneys. The promise of depth gave way to practices focused on cost management, defensibility, and speed.

The Rules Committee and the judiciary responded as best they could. Courts sought to curtail abusive tactics and to discipline discovery practice through doctrine. The federal 2006 and 2016 rule amendments addressed the explosion in volume and the unworkable scope of requests seeking "any and all documents about, concerning or related" to case topics and emphasized proportionality as a governing principle. These interventions mattered. They imposed limits and restored some measure of control. But they did not, standing alone, recover lived experience.

Technologists also did their part. Keyword search became increasingly sophisticated. TAR introduced a new dynamic, replacing blunt keyword hits with probabilistic predictions making proportionality operational. TAR allowed parties to manage volume and cost more

effectively. Yet even as efficiency improved, something essential remained elusive. The focus remained on classification rather than reconstruction, on ranking rather than meaning.

Lived experience continued to evade the industry until the emergence of generative AI. We now stand at a threshold. AI extends the keywords to TAR trajectory. AI can not only predict and rank likelihood of responsiveness, but also articulate reasons, trace connections, organize narratives, and respond to inquiries. These capabilities, working with the rich and expansive fabric of digital artifacts, open the world of pertinent lived experience. The digital record becomes something that can be explored, questioned, and understood as a fabric of human activity rather than a warehouse of isolated artifacts. This is not a utopian moment. Challenges of cost, time, and discipline remain. But AI makes it possible to engage the factual record persuasively.

What e-discovery now reveals is that facts are not inert objects. Digital artifacts reflect human action and the human world. Emails are written under pressure, drafts revised in response to uncertainty, messages shaped by audience, timing, and expectation. These artifacts acquire meaning only when reconstructed within a shared factual world. Treated as isolated documents, they flatten experience. Reconstructed in relation to one another, they recover depth, temporality, and plurality, and allow meaning to emerge.

In the end, electronic discovery is both the reconstruction of the factual world and the recovery of its meaning. In an era increasingly dominated by synthetic information, strategic misinformation, and unaccountable persuasion, civil litigation remains one of the few institutional settings in which contested claims must be grounded in a common evidentiary record and subjected to public standards of proof, with courts serving as stewards of the procedural conditions that make such judgment possible. In litigation, contested factual claims must appear, be tested, and be subjected to publicly accountable standards of relevance, reliability, and admissibility. Unlike political discourse, advertising, or many forms of administrative decision-making, litigation is anchored in a shared evidentiary record and constrained by rules that distinguish persuasion from assertion and proof from narrative force. Electronic discovery plays a critical role in sustaining this space by providing a common factual ground against which competing interpretations can be evaluated. Properly practiced, e-discovery *preserves the shared factual ground on which judgment depends and flourishes.*

ARTIFICIAL INTELLIGENCE IN DOCUMENT REVIEW: STATISTICALLY DEFENDING THE PROCESS AND RESULTS OF E-DISCOVERY WORKFLOWS

Aron Ahmadia,[*] *Nathan Reff*,[**] *Matthew Reiber*,[***] *Chad Roberts*[****] & *Cristin Traylor*[*****]

Abstract

The use of artificial intelligence (AI) for document review allows legal teams to process and produce electronic documents on a scale, speed, and cost that would have been unheard of only a few years ago.

But as legal teams become increasingly dependent on AI to locate and prioritize responsive documents, courts and opposing parties rightly demand heightened verification that the workflow is well designed and that the evidentiary output is sufficiently accurate to ensure a full and fair adjudication of the facts at trial. This Article explores the processes and statistical tools available to litigants to respond to related objections and concerns, focusing on the methods available to measure the completeness and accuracy of an e-discovery workflow. Drawing on principles of information retrieval and machine learning, this Article examines common metrics and explains their application within the framework of a defensible e-discovery workflow. This Article also addresses statistical sampling techniques, including random and stratified sampling, and discusses how confidence intervals and margins of error can be used to assess the reliability of the AI-driven output. It further examines the validation of technology-assisted review protocols, including continuous learning, simple learning, and other iterative workflows that adapt to human input and subsequent machine analysis. By examining both quantitative approaches and qualitative safeguards and by drawing on interdisciplinary literature from academia, industry, and government, this Article evaluates each technique's strengths, limitations, and situational suitability, thereby providing courts and practitioners with an actionable guide for selecting defensible validation strategies across diverse review contexts. This Article provides relevant stakeholders with a better understanding of how to evaluate and measure the performance of an AI-driven workflow and verify whether it meets the requirements of the Federal Rules of Civil Procedure. This Article is particularly helpful

* Dr. Ahmadia leads Applied Science at Relativity.

** Dr. Reff leads Applied Science Capabilities for Relativity in North America.

*** Matthew Reiber is an Associate Professor of Law at Jacksonville University College of Law.

**** Chad Roberts is the principal at eDiscovery CoCounsel, pllc, a law firm that serves as discovery counsel in complex civil litigation.

***** Cristin Traylor is the Senior Director of AI Transformation & Law Firm Strategy at Relativity and is a member of both the Virginia and DC Bars.

because it provides a roadmap for litigants and judges who must confront AI-related issues at critical junctures in the litigation stream. It is particularly timely because continued advances in AI, particularly generative AI, will bring this tool to a broader range of litigants who will use it and judges who must evaluate its merits.

INTRODUCTION

One of the more significant challenges confronting civil litigants is marshaling ever-increasing volumes[1] of electronic evidence[2] through the discovery process.[3] While the sources of this evidence are often readily discernable, reviewing it prior to production can be burdensome and

1. *See, e.g.*, *In re* Aspartame Antitrust Litig., 817 F. Supp. 2d 608, 614 (E.D. Pa. 2011) (defendants collectively gathered the equivalent of eighty million pages of potentially responsive documents); Tampa Bay Water v. HDR Eng'g, Inc., No. 08-CV-2446, 2012 WL 5387830, at *21 (M.D. Fla. Nov. 2, 2012) ("This was a lengthy, highly technical case which involved 17 million pages of documents."); *In re* Online DVD-Rental Antitrust Litig., 779 F.3d 914, 925 (9th Cir. 2015) ("Netflix ultimately produced almost 15 million pages . . . ."); *In re* Allergan Biocell Textured Breast Implant Prods. Liab. Litig., MDL No. 2921, 2022 WL 16630821, at *1 (D.N.J. Oct. 25, 2022) ("Defendants . . . collected and indexed 9.371 terabytes of data . . . ."); Fast Memory Erase, LLC v. Spansion, Inc., No. 10-cv-0481, 2010 WL 5093945, at *4 (N.D. Tex. Nov. 10, 2010) (Intel defendants collected and processed "more than 2,100 gigabytes of electronically stored information."); Brief of Defendants-Appellees Cisco Ironport Systems, LLC and Return Path, Inc., CBT Flint Partners, LLC v. Return Path, Inc., 737 F.3d 1320 (Fed. Cir. 2013) (No. 2013-1036), 2013 WL 1741921, at *12 (Cisco collected "1.2 terabytes of data – or the equivalent of 12% of the printed volume of the Library of Congress.").

2. The increasing volume is driven largely by the proliferation of unstructured data. Unstructured data is "free-form data that either does not have a data structure or has a data structure not easily readable by a computer without the use of a specific program designed to interpret the data; created without limitations on formatting or content by the program with which it is being created." The Sedona Conference, *The Sedona Conference Glossary: eDiscovery & Digital Information Management, Sixth Edition*, 27 SEDONA CONF. J. 1, 131–32 (2026) [hereinafter *The Sedona Conference Glossary*]. Common examples of unstructured data include email, text messages, social media posts, audio files, and videos. *See, e.g.*, Thought Leadership Team, *Unstructured Data: The Black Hole of Ediscovery*, KLDISCOVERY (Oct. 29, 2013), https://www.kldiscovery.com/blog/unstructured-data-the-black-hole-of-ediscovery [https://perma.cc/TQN9-8QVY].

3. The obligation to marshal data through the discovery process is most frequently encountered in the context of a response to a request for the production of documents and electronically stored information pursuant to Rule 34 of the Federal Rules of Civil Procedure. FED. R. CIV. P. 34(a)(1)(A). The term "electronically stored information" refers to "information that is stored electronically, regardless of the media or whether it is in the original format in which it was created . . . ." *The Sedona Conference Glossary*, *supra* note 2, at 45–46.

costly.[4] This is not a new problem of course,[5] and the Federal Rules of Civil Procedure have been amended several times since the early 1980s to make discovery more efficient and less expensive.[6] But even the best procedural system—with the most advanced rules—will be tested from time to time as individuals and entities create increasing volumes of electronic information, store it in wider array of locations, and render it discoverable in a broader range of cases.[7] Rule 26 laudably defines the *scope* of discovery in relation to the needs of the case (as well as to relevance and privilege) and thus allows the breadth of discovery to ebb and flow from case to case.[8] But Rule 26 does not address the *success* of

---

4. *See, e.g.*, Lawson v. Spirit AeroSystems, Inc., No. 18-1100, 2020 WL 6343292, at *1, 1 (D. Kan. Oct. 29, 2020) (collection of 322,524 electronic documents resulted in an end-to-end cost of $791,700.21).

5. The antitrust battles of the 1970s were the first to confront large-scale document discovery. *See, e.g.*, IBM Corp. v. United States, 471 F.2d 507, 523 (2d Cir. 1972) (Mulligan, J., dissenting) ("It is indeed mind-boggling to contemplate [the production of] 17 million document pages which in bulk weigh 87 tons and would stretch from coast to coast.").

6. Since the late 1970s, the scope of discovery has been a recurring issue for the Advisory Committee on Civil Rules, the subject matter committee of the Judicial Conference tasked with responsibility for studying and recommending changes to the Federal Rules of Civil Procedure. The Advisory Committee's work resulted in several amendments to Rule 26 and the scope of discovery. *See* FED. R. CIV. P. 26(b)(1)(iii) (1983) (authorizing court to limit the frequency and extent of discovery when it is "unduly burdensome or expensive, taking into account the needs of the case, the amount in controversy, limitations on the parties' resources, and the importance of the issues at stake in the litigation"); FED. R. CIV. P. 26(b)(1)-(2) (1993) (rearranging content of Rule 26(b) by moving specific limitations on discovery into subdivision (b)(2) and authorizing courts to limit discovery where "the burden or expense of the proposed discovery outweighs its likely benefit" and based on "the importance of the proposed discovery in resolving the issues"); FED. R. CIV. P. 26(b)(1) (replacing "relevant to the subject matter involved in the pending action" with "the claim or defense of any party," although authorizing district court to allow broader subject matter discovery upon a showing of good cause); FED. R. CIV. P. 26(b)(2)(B) (2006) (adding additional limitation on scope of discovery with respect to electronically stored information that is not reasonably accessible because of undue burden or cost absent a court order); FED. R. CIV. P. 26(b)(1)-(2) (2015) (rearranging the text regarding the scope of discovery, and the limitations on discovery, by moving certain text from Rule 26(b)(2) to Rule 26(b)(1) and providing that "the scope of discovery" is defined in relation to relevance, privilege, as well as proportionality to the needs of the case, considering "the importance of the issues at stake in the action, the amount in controversy, the parties' relative access to relevant information, the parties' resources, the importance of the discovery in resolving the issues, and whether the burden or expense of the proposed discovery outweighs its likely benefit.").

7. George L. Paul & Jason R. Baron, *Information Inflation: Can the Legal System Adapt?* 13 RICH. J. L. TECH. 10, 13 (2007) ("[T]he amount of information in business has increased by thousands, if not tens of thousands of times in the last few years. In a small business, whereas formerly there was usually one four-drawer file cabinet full of paper records, now there is the equivalent of two thousand four-drawer file cabinets full of such records, all contained in a cubic foot or so in the form of electronically stored information. This is a sea change.").

8. Rule 26(b)(1) provides that parties may obtain discovery "regarding any nonprivileged matter that is relevant to any party's claim or defense and proportional to the needs of the case . . . ." FED. R. CIV. P. 26(b)(1). The extent to which discovery is "proportional to the needs of the

discovery, specifically whether it unearthed an optimal corpus of evidence from which the parties can support their factual contentions at trial—optimal in the sense of completeness and accuracy after considering the cost and benefit of doing more or better.[9] To compound this situation, the federal courts have not addressed what constitutes successful e-discovery with sufficient frequency or at a meaningful level of generality[10] to provide a measuring stick against which litigants can evaluate their own efforts and courts can adjudicate claims of inadequacy or insufficiency. The scholarship regarding successful discovery is equally sparse.[11]

---

case" is measured against six specific factors: "the importance of the issues at stake in the action, the amount in controversy, the parties' relative access to relevant information, the parties' resources, the importance of the discovery in resolving the issues, and whether the burden or expense of the proposed discovery outweighs its likely benefit." *Id.* Several state procedural systems similarly define the scope of discovery with respect to both relevance and proportionality. *See, e.g.*, AZ. R. CIV. P. 26(b)(1) (scope of discovery is identical to that in Rule 26 of the Federal Rules of Civil Procedure); DEL. CH. CT. R. 26(b)(1) (scope of discovery is substantively the same as defined in Federal Rule 26).

9. Completeness or "recall" refers to the percentage of responsive items within a document population that are correctly identified as such and thereafter produced in discovery. Accuracy or "precision" in turn refers to the percentage of items retrieved by a particular search methodology that are in fact responsive. *See, e.g.*, Da Silva Moore v. Publicis Groupe, 287 F.R.D. 182, 189 (S.D.N.Y. 2012) ("Recall is the fraction of relevant documents identified during a review; precision is the fraction of identified documents that are relevant."); Alivecor, Inc. v. Apple, Inc., No. 21-cv-03958, 2023 WL 2224431, at *1 (N.D. Cal. Feb. 23, 2023) (same).

10. *See, e.g.*, *In re* Insulin Pricing Litig., No. 23-md-3080, 2025 WL 1112837, at *5–7 (D.N.J. Apr. 11, 2025) (describing validation protocol); *In re* Uber Techs., Inc., Passenger Sexual Assault Litig., MDL No. 3084, 2024 WL 3491760, at *5–6 (N.D. Cal. Mar. 15, 2024) (requiring disclosure of recall and prevalence but declining to specify a percentage for either); *In re* Soc. Media Adolescent Addiction/Pers. Inj. Prods. Liab. Litig., No. 22-md-03047, 2024 WL 1786293, at *7 (N.D. Cal. Feb. 20, 2024) (ordering that each party "shall take reasonable steps to validate its review process" but not specifying them); BCBSM, Inc. v. Walgreen Co., No. 20-cv-01853, 2023 WL 6852533, at *8–10 (N.D. Ill. July 21, 2023) (providing validation worksheet "to obtain an estimate of the quantity and quality (*i.e.*, materiality) of responsive documents identified and missed by a particular review effort, and to suggest the source of any inadequacy that may be indicated by such estimate"); Deal Genius, LLC v. O2COOL, 682 F. Supp. 3d 727, 734–35 (N.D. Ill. July 14, 2023) (describing use of elusion testing as a validation metric); *In re* Diisocyanates Antitrust Litig., Misc. No. 18-1001, 2021 WL 4295729 (W.D. Pa. Aug. 23, 2021) (evaluating parties' competing approaches to validation); City of Rockford v. Mallinckrodt ARD Inc., 326 F.R.D. 489, 492–95 (N.D. Ill. 2018) (describing utility of elusion test to validate adequacy of production); *In re* Broiler Chicken Antitrust Litig., No. 16-cv-08637, at *4–6 (N.D. Ill. Jan. 3, 2018) (identifying validation protocol "to ensure a reasonable production"); Rio Tinto PLC v. Vale, No. 14 Civ. 3042, 2015 WL 13956130, at *12 (S.D.N.Y. Sep. 8, 2015) (establishing validation protocol to determine whether the parties' respective productions were adequate).

11. *See, e.g.*, Karl Schieneman & Thomas C. Gricks III, *The Implications of Rule 26(g) on the Use of Technology-Assisted Review*, 7 FED. CTS. L. REV. 239, 269–73 (2013) (discussing the use of recall and validation statistics to satisfy the requirements of Federal Rule 26(g)); Maura R. Grossman & Gordon V. Cormack, *Comments on "The Implications of Rule 26(g) on the Use of*

An evolving collection of AI technologies, however, offers the possibility of achieving successful e-discovery by maximizing completeness and accuracy while minimizing burden and cost, even when applied to the largest data collections.[12] These technologies—which predict the responsiveness of individual documents—allow human reviewers to focus on the responsive documents in a collection and avoid the often-greater number of non-responsive ones. The resulting reduction in human involvement lowers the cost of discovery,[13] accelerates its completion, and shortens the time to trial, thus promoting the "just, speedy, and inexpensive determination of every action and proceeding."[14]

Although litigants have used AI in discovery for more than a decade,[15] disputes over its ability to reliably distinguish responsive from non-

Technology-Assisted Review," 7 FED. CTS. L. REV. 285, 301–13 (2014) (questioning the utility of relying on recall to evaluate the success of discovery).

12. *See* Maura R. Grossman, Gordon V. Cormack & Jason R. Baron, *Does the LLMperor Have New Clothes? Some Thoughts on the Use of LLMs in Ediscovery*, 109 THE ADVOC. (TEXAS) 22, 23 (2024) ("Over the past 15 years or so, the legal profession has become increasingly aware of the availability of various forms of AI used specifically to find responsive documents in complex litigation. The two most-established methods – commonly dubbed 'TAR 1.0' and 'TAR 2.0' – employ supervised machine learning to distinguish responsive documents from non-responsive documents."). *See also* Lea Malani Bays, *Harness the Power of GenAI in Document Review*, 61 TRIAL MAG. 34, 37 (May 2025) ("GenAI is the new hope for solving e-discovery challenges. . . . [I]nitial testing shows promising results, with GenAI identifying relevant documents with increased accuracy. Early studies suggest that GenAI can identify over 95% of relevant documents in a dataset and may even outperform both human reviewers and TAR.").

13. Publicly available information regarding the burden and expense of attorney review of electronic evidence prior to production is limited. *Lawson v. Spirit AeroSystems, Inc.*, No. 18-1100, 2020 WL 6343292 (D. Kan. Oct. 29, 2020) is one of the few cases to address the issue. In that case, the district court awarded hundreds of thousands of dollars in review costs for first and second level attorney review based on 3,620.9 hours. Publicly available information regarding the cost of storing or "hosting" electronic evidence during review is more readily available. *See, e.g.*, United States *ex rel.* King v. Solvay S.A., Civil Action H-06-2662, 2016 WL 3523873, at *15 (S.D. Tex. June 28, 2016) ($429,964.95 for data storage); Abbott Point of Care, Inc. v. Epocal, Inc., No. CV-08-S-543, 2012 WL 7810970, at *2 (N.D. Ala. Nov. 5, 2012) ($340,498 total for hosting); Murphy v. Precision Castparts Corp., No. 3:16-cv-00521, 2021 WL 4524153, at *6 (D. Or. Oct. 4, 2021) ($168,069 for data hosting).

14. FED. R. CIV. P. 1. The reduction of human involvement should also reduce the frequency of disputes regarding the proportionality of particular discovery by eliminating a common justification for reducing the breadth of discovery in particular cases. *See, e.g.*, Ruggles v. WellPoint, Inc., No. 08-CV-201, 2010 WL 11570681, at *5 (N.D.N.Y. Dec. 28, 2010) ("Ever since the advent of storing monumental quantities of information and data electronically and its dynamic impact upon discovery, by all measures, the cost of producing e-discovery, which is generally borne by the responding party, consumes an inordinate quantum of most discovery debates. Here, the cost of e-discovery dominates our discussion.").

15. *See, e.g.*, Roshanak Omrani et al., *Beyond the Bar: Generative AI as a Transformative Component in Legal Document Review*, RELATIVITY ODA LLC (2024) at 2 https://resources.relativity.com/rs/447-YBT-249/images/Beyond-the-Bar_preprint_2-9%20%

responsive documents persist,[16] regardless of the expertise of the lawyers in the case.[17] The requesting party contends that the process was not adequately executed, the resulting production is incomplete, and the

281%29.pdf?version=0 [https://perma.cc/Z5R2-BWSN] ("Supervised machine learning first came into use in eDiscovery around 2005, with courts in the United States, England, Ireland, Australia and other jurisdictions explicitly encouraging its use starting in 2012."); *see also In re* Valsartan, Losartan, & Irbesartan Prods. Liab. Litig., 337 F.R.D. 610, 616 (D.N.J. 2020) ("We are past the time when parties and courts view TAR as an outlier."). For early cases involving the use of technology in the discovery process, see Da Silva Moore v. Publicis Groupe, 287 F.R.D. 182, 183 (S.D.N.Y. 2012) ("This judicial opinion now recognizes that computer-assisted review is an acceptable way to search for relevant ESI in appropriate cases."); Glob. Aerospace Inc. v. Landow Aviation, L.P., No. CL 61040, 2012 Va. Cir. LEXIS 50, at *2 (Va. Cir. Ct. Apr. 23, 2012) ("[I]it is hereby ordered Defendants shall be allowed to proceed with the use of predictive coding for purposes of the processing and production of electronically stored information . . . ."); W. Penn Allegheny Health Sys., Inc. v. UPMC, No. 2:09-CV-0480-JFC, 2013 WL 12134102, at *1 (W.D. Pa. Feb. 28, 2013) (authorizing use of "a technology-assisted review process, including, for example, one that predictively codes documents as subject to or not subject to production . . . ."); W Holding Co. v. Chartis Ins. Co. of P.R., No. 11-2271, 2013 WL 1352562, at *5 (D.P.R. Apr. 3, 2013) ("The Parties shall meet and confer in good faith about any other technology or process that a producing party proposes to use to streamline the culling, review and production of ESI (e.g., email threading, near de-duplication, technology assisted review)."); Edwards v. Nat'l Milk Producers Fed'n, No. 11-CV-04766, at *1 (N.D. Cal. Apr. 17, 2013) (authorizing use of predictive coding "for a more cost efficient and higher quality review"); United States v. Educ. Mgmt. LLC, No. 07-cv-00461, 2013 WL 12140442, at *8 (W.D. Pa. Nov. 24, 2013) (approving use of computer assisted review); Rio Tinto PLC v. Vale S.A., 306 F.R.D. 125, 127 (S.D.N.Y. 2015) ("[I]t is now black letter law that where the producing party wants to utilize [technology assisted review] for document review, courts will permit it.").

16. *See, e.g.*, *In re* Insulin Pricing Litig., No. 23-md-3080, 2025 WL 1112837, at *5 (D.N.J. Apr. 11, 2025) ("The parties do not dispute that there must be a validation process for the TAR model. However, they dispute the scope of the data set that would comprise the validation process."); *In re* Uber Techs., Inc., Passenger Sexual Assault Litig., MDL No. 3084, 2024 WL 3491760, at *5 (N.D. Cal. Mar. 15, 2024) (resolving a dispute where parties agreed defendant would disclose recall and prevalence but disagreed about whether defendant would also have to disclose "the input quantities used to calculate" them); *In re* Soc. Media Adolescent Addiction/Pers. Inj. Prods. Liab. Litig., No. 22-md-03047, 2024 WL 1786293, at *7 (N.D. Cal. Feb. 20, 2024) ("In general, the Parties agree that some form of validation of the review procedures is appropriate. The disputes center on how much information to share about validation and the burden required to seek additional information from a Producing Party."); Deal Genius, LLC v. O2Cool, LLC, 682 F. Supp. 3d 727, 738–39 (N.D. Ill. 2023) (calling into question the reliability of elusion testing called into question based on the results from the application of an additional search term against the null set); *In re* Diisocyanates Antitrust Litig., Misc No. 18-1001, 2021 WL 4295729, at *2–5 (W.D. Pa. Aug. 23, 2021) (dispute regarding particular methodology for performing elusion test); City of Rockford v. Mallinckrodt ARD Inc., 326 F.R.D. 489, 491–92 (N.D. Ill. 2018) (resolving a dispute regarding whether defendant would be required to conduct elusion test at conclusion of its document review or whether plaintiff would be required to identify potential gaps and omissions in document production).

17. *See, e.g.*, Civil Docket, *In re* Insulin Pricing Litig., No. 23-md-3080 (D.N.J. 2023) (Davis Polk; Jones Day; Kirkland & Ellis); Civil Docket, *In re* Diisocyanates Antitrust Litig., Misc. No. 18-1001 (W.D. Pa. 2018) (Eckert Seamans, Mayer Brown; Morgan Lewis; Reed Smith); Civil Docket, *In re* Broiler Chicken Antitrust Litig., No. 16-cv-08637 (N.D. Ill. 2016) (Simpson Thacher; Skadden Arps; White & Case).

requirements of Rules 26(b)(1)[18] and 26(g)(1)[19] are therefore not satisfied. The producing party counters that the process was sound, the resulting production is proportionate to the needs of the case, and that, in any event, it is best situated to evaluate the "'procedures, methodologies, and technologies appropriate for preserving and producing'" its own information.[20] The court must then sift through arguments from lawyers and testimony from experts to craft an appropriate resolution.

This Article attempts to reduce the frequency of these disputes and to facilitate their resolution when they cannot be avoided by providing litigants and judicial officers with a thorough, readily accessible explanation of the use of AI in discovery and the statistical measures that can be used to validate its successful performance (or, alternatively, identify its shortcomings) regardless of the selected workflow. Part I describes the evolution of AI in e-discovery, highlighting how technology can be deployed to reduce the burden and expense of e-discovery while improving the quality of the output. Part II then provides a "user guide" of sorts regarding the validation of e-discovery workflows and a step-by-step process for selecting and implementing the more appropriate methodology for individual cases. While this part acknowledges perfection is not attainable,[21] it demonstrates that a well-designed workflow, coupled with the use of appropriate technology, increases the likelihood that e-discovery will be successful and the "truth will out."

---

18. *See supra* note 8.

19. Rule 26(g)(1) imposes a signature requirement on all discovery responses, which signature carries with it deemed certifications that the response is "consistent with these rules" and is not "unreasonable . . . considering the needs of the case, prior discovery in the case, the amount in controversy, and the importance of the issues at stake in the litigation." FED. R. CIV. P. 26(g)(1). These deemed certifications are best understood in relation to the evolving norms of the profession regarding discovery obligations. *See, e.g.*, Schieneman & Gricks III, *supra* note 11, at 246–47 ("[T]he producing party must conduct a reasonable, good faith, and diligent search for responsive ESI, recognizing that the objective is not perfection. That obligation, however, is tempered by the proportionality concept . . . and the ultimate goal of 'secur[ing] the just, speedy and inexpensive determination of every action and proceeding' . . . .") (quoting FED. R. CIV. P. 1).

20. *In re* Insulin Pricing Litig., No. 23-md-3080, 2025 WL 1112837, at *1 (D.N.J. Apr. 11, 2025) (quoting The Sedona Conference, *The Sedona Principles, Third Edition: Best Practices, Recommendations & Principles for Addressing Electronic Document Production*, 19 SEDONA CONF. J. 1, 52, 118 (2018)). *See also* Nichols v. Noom Inc., No. 20-CV-3677, 2021 WL 948646, at *2 (S.D.N.Y. Mar. 11, 2021) ("[A] producing party is best situated to determine its own search and collection methods so long as they are reasonable.").

21. *See, e.g.*, Fed. Hous. Fin. Agency v. HSBC N. Am. Holdings Inc., No. 11 Civ. 6189, 2014 WL 584300, at *2 (S.D.N.Y. Feb. 14, 2014) ("Parties in litigation are required to be diligent and to act in good faith in producing documents in discovery. The production of documents in litigation such as this is a herculean undertaking . . . . [N]o one could or should expect perfection from this process. All that can be legitimately expected is a good faith, diligent commitment to produce all responsive documents uncovered when following the protocols to which the parties have agreed, or which a court has ordered.").

I. THE EVOLUTION OF ARTIFICIAL INTELLIGENCE IN DOCUMENT REVIEW AND ITS VALIDATION FOR USE

Although the procedure for obtaining documents in civil litigation has changed dramatically over the years,[22] the process for reviewing and making them available to others remained constant until recently—manually collecting, manually reading, and manually identifying them as responsive or not in relation to a particular request. The process was rarely questioned because it was mutually known and understood by all involved. It was mutually known because it was the only process available. And it was mutually understood because it was based on the everyday experience of everyday lawyers. There was no "black box technology" or opaqueness to the review and classification process, and the parties rarely questioned its capacity to generate a complete and accurate evidentiary record.

By contrast, a producing party in litigation today has a wide selection and combination of technological tools that can automate large segments of a discovery production workflow, including the task of classifying documents for responsiveness. When these tools are used, their internal operations are largely unseen and unrelatable to everyday human experience. The verification of their outcomes is no longer a self-evident process.

While the notion of *validating* the outcome of a document review workflow predates—and is independent of—the employment of AI in a particular workflow, it has been the "black box" opaqueness of advanced AI technologies that has accelerated the need for validating the process as part of a defensible document production.

---

22. *See, e.g.*, 8B CHARLES ALAN WRIGHT, ARTHUR R. MILLER, FEDERAL PRACTICE AND PROCEDURE § 2201 (3d ed. 1998) ("As originally adopted in 1938, Rule 34 was patterned after the rules under the English Judicature Act, the former federal Equity Rule 58, and American state statutes authorizing the court to order parties in possession or control of documents to permit other parties to inspect and copy them before trial. To the extent that it was applicable, it superseded the former bill in equity for discovery in aid of an action at law. Under practices that prevailed before the rule was adopted, only a very limited form of production of documents was authorized and this was so hedged about with technicalities that actually little effective discovery was obtainable."); 7 James W. Moore et al., Moore's Federal Practice – Civil § 34App.100 (2025) ("The only effective method of obtaining a pretrial inspection of documents in an action at law was to bring a bill of discovery in equity in aid of the action at law . . . . Pretrial inspection of documents was permitted only if the bill of discovery showed that there was a real need for the inspection and that it could not be obtained in a practical, useful way in the action at law."); FED R. CIV. P. 34 advisory committee's note to 1970 amendment ("Rule 34 is revised to accomplish the following major changes in the existing rule: (1) to eliminate the requirement of good cause; [and] (2) to have the rule operate extrajudicially . . . .").

A. *The Evolution of AI Technologies in Document Review*

1. Eyes on Documents—Manual Review and the Emergence of Document Review Platforms

Following the widespread commercial adoption of electronic mail messaging, word processing, and document scanning into portable digital format, emerging discovery technologies were focused on managing document review workflows. Initial technologies simply attempted to accelerate traditional document review activities, namely looking at documents and making decisions about them.

A class of software applications—document review platforms—automated the rapid extraction of metadata and presentation of document images to human reviewers with the ability to add reviewer-created designations of the documents. The result of this type of review workflow was to increase the organization and structure of the document collection and, in turn, increase the ability to search, retrieve, and produce targeted content.

The review process that involves the manual and sequential classification of all documents by human review alone is referred to as *linear review*. Other than logistically accelerating the visual presentation of the documents and the application of user created designations, a linear review does not otherwise leverage any AI functionality on the document classification process. While other technology features of a document review platform, such as Boolean word searches or filtering by metadata, can result in the culling of a larger data set into a smaller one more likely to contain responsive documents, there is no functionality in linear review that duplicates or replaces human judgments about the document itself.

Because of the outsized impact that document review and classification tasks have on the overall costs of discovery,[23] the commercial development of technologies that address document review and classification has tended to justify the most sophisticated features. Commercial AI technologies in e-discovery first emerged in document review platforms to support the tasks of document review and classification for Rule 34 productions.

---

23. *See, e.g.*, Nicholas M. Pace & Laura Zakaras, *Where the Money Goes: Understanding Litigant Expenditures for Producing Electronic Discovery*, RAND INST. FOR CIVIL JUST. xiv (2012) (survey results revealed that, on average, the proportion of total costs associated with an e-discovery production cycle were approximately: 73% for document review and classification activities, 19% for processing (reformatting) the data for use by a review platform, and 8% for the collection of the data from the producing party's information environment).

2. Machine Learning and Predictive Coding: Replicating Human Judgments and Applying Them to Unseen Documents

In the 2010s, commercial applications of two related sub-sets of AI technology began to emerge in document review platforms: (1) machine learning algorithms used as document classifiers and (2) natural language processing (NLP) algorithms. NLP is a type of machine learning algorithm designed to identify complex patterns in the textual content of language. NLP functionality is a building block of more complex machine learning algorithms used in predictive coding.

Once an NLP algorithm is used to create a linguistic profile of all the documents in a data set (*indexing*), the stage is set for what has become one of the most impactful discovery technologies to date: predictive coding in document review. Predictive coding involves the computer developing an algorithmic model of the coding decisions made by human reviewers on a subset of documents (*training*) and then using the developed model to replicate and apply those decisions upon a much larger set of non-reviewed documents. It does this by identifying the common linguistic attributes of documents receiving a particular coding decision in a training set and then identifying the non-reviewed documents that most closely reflect the same linguistic attributes. When this identification and replication can be accomplished with reasonable accuracy—when testing shows the coding values predicted by the machine closely match the coding values of the human reviewer—it results in the defensible avoidance of the need to review large volumes of non-responsive documents. Document review using these machine learning AI tools came to be referred to as technology-assisted review, or simply "TAR."

As the earliest TAR technologies came to market, two methodologies emerged for the process by which the algorithmic coding decision model was developed, and both methodologies remain highly relevant to the understanding and emergent use of generative AI in document review: simple learning (TAR 1) and continuous learning (TAR 2).

3. "TAR 1" means "Simple Learning" means "Two Stage Protocol"

The first commercial TAR platforms used a two-step process to build the algorithmic model (a *classifier*) that predicted coding decisions. The first stage of what came to be referred to as a "TAR 1" protocol involves training and testing on a control set.

An initial *training set* of documents is first coded by human reviewers, and these coding decisions are delivered to the machine-learning software. The software then analyzes the training set and makes a first attempt to build a classification model capable of discerning those linguistic attributes that result in a document receiving a particular coding

decision. It then uses the classification model to make, for all the many other non-reviewed documents, a prediction as to whether a particular document is likely to be responsive or non-responsive.

A separate *control set* of documents is also coded by human reviewers, but these coding decisions are not used for training the predictive model. The control set is used to test the model's accuracy by comparing the human reviewer's coding decisions to the decisions predicted by the model. If testing shows the accuracy of the model's predictions to be within acceptable limits, the training process ends and the TAR 1 workflow advances to the second stage.

If testing shows the model's predictions are outside an acceptable degree of accuracy, the training cycle is repeated: additional documents are added to the training set,[24] reviewed, and the new training set coding decisions are again delivered to the machine-learning software. A larger and carefully curated training set tends to refine and improve the accuracy of the predictive classification model. The model's predictive accuracy is tested again. This training–testing–training–testing sequence is repeated iteratively until the model achieves an acceptable degree of accuracy and testing shows that additional documents and coding decisions will not appreciably improve the model's accuracy.

At this point, the predictive model is considered to have "matured," and the workflow advances to the second stage of the TAR 1 protocol. In the second stage, the model computes the predictive value of all documents in the data set. Documents with a sufficiently high predictive value for responsiveness are reviewed, whereas those with sufficiently low predictive value for responsiveness are not reviewed.

4. "TAR 2" means "Continuous Learning" means "Single Stage Protocol"

After the first predictive coding software came to the commercial market for document review, a second TAR protocol emerged. Instead of the two-step iterative training followed by static classification workflow of TAR 1, this protocol utilized a single continuous process of reviewer generated coding decisions. In the TAR 2 protocol, human reviewers begin what is essentially a linear review of documents in a large data set. In the background, their coding decisions are continuously delivered to the machine-learning software which, in turn, continuously builds and refines its predictive classification model. For each document in the data set, the classification model then ranks, typically by a scale of 1 to 100, the degree to which a particular document is likely to receive a particular designation.

---

24. A variety of workflows and methods exist that involve *how* the additional training documents are selected, and the criteria used for their selection.

The classification model frequently updates with new coding decisions from the reviewers, and as it does so, a choice can be made concerning what types of unreviewed documents are served to the human reviewers in their review queue. When the most highly predictive documents are placed in the queue first, the most important documents are likely to be identified sooner, but the overall time for the predictive model to reach maturity may be prolonged. Alternatively, the review process can be configured to select for review those documents which the classification model is unable to make a prediction with confidence. Coding those documents first may result in the model "maturing" more rapidly and may shorten the overall completion rates of document review. Each data set (and each case) has its own characteristics, which dictate optimal workflow.

5. Generative AI in Document Review: TAR With a New Classification Engine

One of the most promising applications of generative AI in discovery is its use in document review as the classification engine of a TAR workflow. Instead of a traditional machine learning algorithm being used as the document classifier to predictively rank documents, a large language model (LLM) is used to classify (or predict) documents for responsiveness, or some other trait. The first commercial applications of generative AI in document review utilize the TAR 1 (two stage protocol) workflow described above (where machine learning algorithms have previously been used as the classification engine instead of an LLM).[25]

In the training phase of the workflow, a user constructs an LLM prompt that describes the responsive content sought to be retrieved from the larger data set. The prompt is then used by the LLM to predict a responsiveness classification to a control set of documents in the data set. In the second testing phase of the workflow that follows, the control set is blind coded by human reviewers and compared to the LLM's predictions for the coding classifications of the same documents. The training and testing iterations continue with sequential refinements of the LLM prompt until there is an acceptable alignment between the human reviewer coding decisions and the coding predicted by the generative AI model.

25. Tara Emory, Jeremy Pickens & Wilzette Louis, *TAR 1 Reference Model: An Established Framework Unifying Traditional and GenAI Approaches to Technology-Assisted Review*, 25 SEDONA CONF. J. 109, 110–15 (2024) [hereinafter *TAR 1 Reference Model*].

B. *Quality Assurance, Quality Control, and Validation Terminology in e-Discovery*

Although most litigants are unlikely to characterize it as such, a successful response to a document request is a complex, process-driven workflow executed according to traditional project management principles, including *quality assurance*, *quality control*, and outcome *validation*.

"Quality Assurance" refers to workflow design activities calculated to prevent errors in the document review process. These are activities that primarily occur prior to the execution of the document review itself: workflow design, search methodology selection, the creation of coding decision protocols, reviewer training, etc.

"Quality Control" refers to activities calculated to identify and correct errors that occur during the execution of the document review process: erroneous document coding decisions, inconsistent document coding decisions, substandard reviewer performance, etc. The process of identifying and testing for errors may include statistical sampling and analysis.

"Validation" refers to outcome confirmation that the overall workflow achieved its intended results and that the accuracy of the process outcome was within acceptable error rates. It answers the question: did the overall production workflow perform as it was intended to? Validation occurs *after* the execution of the document review activities and includes statistical sampling and analysis.

This nomenclature is not uniform in either the literature or the case law, however. Many judicial decisions and much commentary (including the decisions and commentary quoted in this Article) tend to conflate the terminology and the steps taken to enhance the defensibility of the process.

C. *The Migration of Data Science Validation Principles Into Defensibility Standards for Legal Document Review*

Well before the widespread use of AI in document review workflows, significant disparities between the perceived quality and the actual quality of outcomes had already been documented. The landmark study published by Blair and Maron in 1985[26] demonstrated unexpected and shockingly low levels of recall in a full-text document retrieval experiment. Highly experienced legal practitioners with substantial knowledge of a real-world data set believed they were reliably identifying

26. *See generally* David C. Blair & M. E. Maron, *An Evaluation of Retrieval Effectiveness for a Full-Text Document-Retrieval System*, 28 COMPUTING PRACS. 289 (1985).

at least 75% of responsive documents using keyword search technology.[27] In reality, their rate was approximately 20%.[28]

Yet, more than twenty years later, in the introduction to a 2007 survey[29] of search and information retrieval methods in e-discovery, the notions of validation in the outcome of a document review process and quality control as part of review process itself were still being described as aspirational and "relatively unknown in the typical review processes in use today."[30] The gap between emerging technology and practical adoption is further illustrated as follows:

> For example, advanced forms of search techniques, including various forms of fuzzy logic, text mining and machine learning all automatically organize electronically stored information in new ways not achieved by past more familiar methods, including the simple use of "keywords" as the only automated aid to conducting manual searches. Although we are at the dawn of a new era, these new techniques hold the potential to increase both accuracy and efficiency. Through statistical sampling and validation techniques we can then confirm the accuracy of the results of either traditional or alternative forms of search, retrieval, and review.[31]
>
> . . . .
>
> The development of process control logs and second-level review techniques can also help the review team to ensure that the designed process is consistently applied to all of the information to be reviewed. Additionally, a second-level review process based on statistical sampling techniques can ensure the achievement of acceptable levels of quality. While these techniques are relatively unknown in the typical review processes in use today, their widespread adoption in businesses of all types should drive their implementation in large document review projects in the near future.[32]

Many data scientists and early adopters of AI technologies in document review workflows have consistently contemplated the notions of quality, validation, and defensibility. Historically, there has been recognition that the rigors of formal academic scrutiny would inform the

---

27. *Id.* at 293.

28. *Id.*

29. The Sedona Conference, *The Sedona Conference Best Practices Commentary on the Use of Search & Information Retrieval Methods in E-Discovery*, 15 SEDONA CONF. J. 217 (2014).

30. *Id.* at 215.

31. *Id.* at 192.

32. *Id.* at 215.

methodologies relied upon by legal practitioners to demonstrate defensibility.

An enormous academic literature already exists on the general subject of information retrieval (IR) and "searching." A number of computer-science oriented organizations and conferences are devoted in whole or in part to the subject, including with respect to comparing or evaluating different alternative search methodologies.[33]

However, so far as these authors are aware, there has been little reference in the academic literature to the special type of data mining problems faced by lawyers as they confront demands from opposing parties in civil discovery to find all responsive documents in large electronic databases under arbitrary, externally imposed deadlines. It must be conceded that "the legal system is remarkable in its reliance on both precise and imprecise concepts."[34] Nevertheless, the prime measures of accuracy analyzed in the information retrieval literature arguably map to the litigation support context.[35]

While notions of quality control and quality assurance in e-discovery document review workflows were clearly established in the legal literature by 2010, there remained comparatively little commentary on the validation of the workflow's final output. The emphasis tended to be

---

33. *See generally* CHRISTOPHER D. MANNING, PRABHAKAR RAGHAVAN & HINRICH SCHÜTZE, AN INTRODUCTION TO INFORMATION RETRIEVAL (Cambridge Univ. Press 2009) (surveying information retrieval and search); (McGraw-Hill 1983) (same); Special Interest Group on Information Retrieval (SIGIR), Ass'n for Computing Mach., https://sigir.org [https://perma.cc/JH4U-7Z9X] (last visited May 2, 2026) (describing a professional organization devoted to information retrieval research); Text Retrieval Conference (TREC), Nat'l Inst. of Standards & Tech., https://trec.nist.gov [https://perma.cc/ZNC3-6TYL] (last visited May 2, 2026) (describing a long-running conference designed to evaluate and compare information retrieval systems and methodologies).

34. Jason R. Baron, *Toward a Federal Benchmarking Standard for Evaluating Information Retrieval Products Used in E-Discovery*, 6 SEDONA CONF. J. 237, 241 (2005) (quoting Daniel E. Rose & Richard K. Belew, *Legal Information Retrieval: A Hybrid Approach*, *in Proceedings of the Second International Conference on AI and Law*, Vancouver, 138–46 (1989)).

35. *Id.*; *see also* Mia Mazza, Emmalena K. Quesada & Ashley L. Sternberg, *In Pursuit of FRCP 1: Creative Approaches to Cutting and Shifting the Costs of Discovery of Electronically Stored Information*, 13 RICH. J.L. & TECH 1, 34 (2007) ("Unlike search technologies, however, the goal of sampling is not solely to find potentially relevant data in a way that lives up to performance measurement standards, but rather to make a qualitative judgment as to whether it is worthwhile to conduct further searches of stored electronic data."); Douglas W. Oard et al., *Evaluation of Information Retrieval for E-Discovery*, 18 A.I. & L. 347, 381 (2010) (This is a comprehensive survey of the Text Retrieval Conference - Legal Tracks (2006–2009) research that foreshadowed more sophisticated Information Retrieval workflows to come: "One broad class of approaches that has gained currency in recent years for achieving that focus is broadly known as 'process quality.' Essentially, the idea is that the important thing is that we agree on how each performer of E-discovery services should design measures to gain insight into the quality of the results achieved by their particular process.").

on the adoption by e-discovery practitioners of project management and quality best practices prevalent in other industry sectors.[36]

D. *The Jurisprudence of Document Review Quality and Validation*

While there has been only modest judicial commentary on what documented success looks like in a discovery workflow that relies on AI, notions of quality and validation have been foundational to judicial acceptance and its continued use. In the presence of uncertainty and mistrust by a requesting party concerning a producing party's competent use of AI to identify responsive documents, it was the transparent integration of quality and validation protocols into the document production workflow that buttressed the original judicial approval of AI's use.

Da Silva Moore v. Publicis Groupe[37] was the first case to expressly consider and approve the use of machine learning AI as part of a document production workflow. When explaining the basis of its approval, the court reviewed and relied upon the process safeguards that would be used by the producing party to address quality assurance, quality control, and end-of-process validation:

> Several other lessons for the future can be derived from the Court's resolution of the ESI discovery disputes in this case. First, it is unlikely that courts will be able to determine or approve a party's proposal as to when review and production can stop until the computer-assisted review software has been trained and the results are *quality control verified.* Only at that point can the parties and the Court see where there is a clear drop off from highly relevant to marginally relevant to not likely to be relevant documents.[38]

The court continued:

> Counsel no longer have to worry about being the "first" or "guinea pig" for judicial acceptance of computer-assisted review. As with keywords or any other technological solution to ediscovery, counsel must *design an appropriate process*, including use of available technology, *with appropriate quality control testing*, to review and produce relevant ESI while adhering to Rule 1 and Rule 26(b)(2)(C) proportionality. Computer-assisted review now can be

36. *See, e.g.*, The Sedona Conference, *Commentary on Achieving Quality in the E-Discovery Process*, 10 SEDONA CONF. J. 299, 307 (2009).

37. *See generally* Da Silva Moore v. Publicis Groupe, 287 F.R.D. 182 (S.D.N.Y. 2012).

38. *Id.* at 192 (emphasis added).

> considered judicially-approved for use in appropriate cases.[39]

Subsequent decisions have reinforced the use of quality and validation metrics to make judicial determinations about the adequacy of a document production workflow.[40] While these cases are instructional and reflect growing depth on quality assurance, quality control, and end-of-process validation, they do not purport to be comprehensive practical frameworks that could be used by a practitioner to craft and execute a defensible document production workflow. The following section provides such a framework.

II. VALIDATING AI-ASSISTED REVIEW: A PRACTICAL FRAMEWORK

The following sections establish that AI offers significant potential for improving the efficiency and effectiveness of document review, but that disputes over its reliability persist and courts have provided limited guidance on how to measure success. Here we address this gap by providing a practical framework for validating AI-assisted review workflows.

The framework proceeds in logical sequence. Section A establishes what completion means in the context of document review, defining the target before discussing how to measure progress toward it. Section B introduces the core validation metrics: recall, precision, F1 score, richness, and elusion. Section C addresses statistical sampling, the mechanism by which these metrics are estimated. Section D provides workflow-specific guidance for the principal review methodologies. Section E addresses end-of-review validation as the critical defensibility gate. Section F discusses how to present validation results to courts and opposing counsel. Section G summarizes key recommendations.

Throughout the following sections of this Article, the focus is on practical guidance: what to measure, how to measure it, and how to present the results in a manner that satisfies legal standards and withstands challenge.

---

39. *Id.* at 193 (emphasis added).

40. *See supra* note 10. In some instances, those metrics disclosed significant deficiencies in a party's document production workflow. *See In re* Domestic Airline Travel Antitrust Litig., 2018 WL 4441507 (D.D.C. 2018). In more recent cases, courts have critiqued proposed quality and validation methodology with more granularity, addressing issues such as stratified sampling and the appropriate end points for overall workflow metrics. *See In re* Broiler Chicken Antitrust Litig., No. 1:16-CV-08637, 2018 WL 1146371, at *4–6 (N.D. Ill. Jan. 3, 2018) (identifying validation protocol "to ensure a reasonable production"); *In re* Diisocyanates Antitrust Litig., Misc. No. 18-1001, 2021 WL 4295729 (W.D. Pa. Aug. 23, 2021) (evaluating parties' competing approaches to validation).

A. *Defining Completion: What Does Success Look Like?*

Before measuring whether a review succeeded, legal teams must define what success means. This may seem obvious, but many discovery disputes arise because the parties never agreed on a definition of "done." The producing party believes it conducted a reasonable review. The requesting party believes the production is incomplete. Without a shared understanding of completion criteria, these disputes become intractable.

1. The Meaning of Completion

Completion is the point at which the producing party can credibly stop review. It is not perfection. No review process, whether conducted by humans or machines, will identify every responsive document. The Federal Rules of Civil Procedure do not require perfection; they require a reasonable search proportional to the needs of the case.[41]

Completion is best understood as an agreed stopping criterion informed by two types of evidence. First, quantitative metrics such as recall, precision, and elusion provide numerical estimates of review performance. These metrics, discussed in detail below in Section B,[42] allow parties to state with statistical confidence that the review achieved specified performance levels. Second, qualitative assessment through sampling and error analysis of the discard population (documents predicted by the AI tool to be non-responsive) provides a reality check on the quantitative metrics. If reviewers examine a sample of discarded documents and find patterns of missed responsive material, this signals a problem regardless of what the metrics suggest.

The combination of quantitative and qualitative evidence provides a defensible basis for stopping review. The producing party can explain not only what the numbers show, but also what human reviewers found when they examined the documents the system recommended discarding.

2. Completion is Not a Fixed Recall Target

A common misconception holds that review is "complete" when recall reaches some fixed threshold, often 75% or 80%. This view is mistaken for several reasons.

First, recall targets should be proportional to the needs of the case. A high-stakes patent litigation involving core technical documents may warrant different recall targets than a routine employment dispute. The proportionality analysis required by Federal Rule of Civil Procedure

---

41. FED. R. CIV. P. 26(b)(1)); *see* The Sedona Conference, *The Sedona Principles, Third Edition: Best Practices, Recommendations & Principles for Addressing Electronic Document Production*, 19 SEDONA CONF. J. 1, 65 (2018) (citing FED. R. CIV. P. 26(b)(1)).

42. *See* discussion *infra* Part II.B.

26(b)(1) applies to validation decisions just as it applies to scope decisions.[43]

Second, recall alone does not capture review quality. A review with 95% recall but 10% precision has identified nearly all responsive documents, but it has also included nine non-responsive documents for every responsive one. Depending on the case, this may represent wasted effort, unnecessary privilege review, or excessive production costs.

Third, recall is difficult to measure directly. Estimating recall requires knowing the total number of responsive documents in the population, which is precisely what the review is trying to determine. As discussed in Section B, recall is estimated in several ways, including indirectly through elusion testing, which introduces its own considerations.[44]

For these reasons, parties should agree at the outset on what metrics will be measured and, if appropriate, the stopping criteria based on the needs of the case. These determinations will be constrained by additional factors such as proportionality and time. A single recall threshold is rarely sufficient; understanding richness, recall, precision, and the associated confidence intervals provides a more complete picture of review performance.

3. Key Questions to Answer Before Review Begins

Legal teams should answer the following questions before beginning document review. These answers form the basis of a validation plan that guides the review and provides the framework for demonstrating completion.

What recall and precision targets are appropriate for this matter? The answer depends on the stakes of the case, the nature of the documents, and the consequences of missing responsive material or over-producing non-responsive material. In some matters, the priority is finding every potentially responsive document regardless of precision. In others, the priority is efficiency, and a somewhat lower recall may be acceptable if it significantly reduces review burden.

What confidence level and margin of error are acceptable? Statistical estimates come with uncertainty. A validation sample might estimate recall at 78%, but the true recall could be higher or lower. Confidence intervals quantify this uncertainty. A 95% confidence interval of plus or minus 5% means that if the sampling process were repeated many times, 95% of the resulting intervals would contain the true value. Legal teams

---

43. FED. R. CIV. P. 26(b)(1); Rio Tinto PLC v. Vale S.A., 306 F.R.D. 125, 127 (S.D.N.Y. 2015) (holding that parties must design discovery processes, including quality control measures, consistent with proportionality).

44. *See* discussion *infra* Part II.B.

should decide in advance what confidence level and margin of error they will accept.

What is the richness of the document population? Richness (also called prevalence) is the proportion of responsive documents in the population. Richness profoundly affects validation strategy and interpretation. In a low-richness population (for example, 2% responsive), even small elusion rates can still leave a large relative proportion of the total responsive documents in the discard set. Understanding richness is essential to interpreting validation results and setting realistic metric and confidence interval expectations. If richness is unknown, it should be estimated early in the review process by drawing a random sample from the full population and having qualified reviewers assess responsiveness. This initial richness estimate then informs sample size planning and provides a baseline for interpreting validation metrics.

How will the review be validated? Validation typically involves sampling to estimate performance metrics. Legal teams should decide in advance what sampling approach they will use (whether from the discard set, the responsive set, or across the full population), who will review the sample, and what criteria will trigger remediation if the sample reveals problems.

4. Documenting the Validation Plan

The answers to these questions should be documented in a validation plan before review begins. This plan serves multiple purposes: it aligns the legal team on completion criteria, provides a framework for making consistent decisions throughout the review, and creates a contemporaneous record demonstrating the reasonableness of the process if challenged.

The validation plan need not be elaborate. For straightforward matters, a one-page summary may suffice. For complex matters with multiple review streams and high stakes, a more detailed plan may be appropriate. What matters is that the plan exists, that it reflects thoughtful consideration of the relevant factors, and that the team follows it.

B. *Validation Metrics: What We Measure and Why*

This Section introduces the core metrics used to evaluate document review performance. These metrics originate in information retrieval research but have direct application to legal discovery. Understanding what each metric measures, and equally important, what it does not measure, is essential to designing effective validation protocols and avoiding common interpretation errors.

1. Core Metrics Defined

Five metrics form the foundation of document review validation. Each measures a different aspect of performance, and together they provide a comprehensive picture of review quality.

Recall measures the proportion of identified responsive documents in the population to the total number of responsive documents that exist in the population. In plain terms, recall answers the question: "How much did we find?" If a population contains 10,000 responsive documents and the review identifies 8,000 of them as responsive, recall is 80%. The remaining 2,000 responsive documents represent the "miss" rate: documents that should have been identified as responsive but were not. Recall is the metric most closely tied to the producing party's obligation to conduct a reasonable search. Low recall means responsive documents were missed. However, recall alone does not tell the complete story.

Precision measures the proportion of documents predicted by the AI tool to be responsive that are, in fact, responsive. In plain terms, precision answers the question: "How accurate were we in what we selected for review?" If the AI tool predicts 20,000 documents as being responsive but only 16,000 of them are actually responsive, precision is 80%. The remaining 4,000 non-responsive documents represent "false positives": documents included in the responsive set but should not have been. Precision affects efficiency and cost. Low precision means the legal team is reviewing, potentially privilege-checking, and producing many documents that are not actually responsive.

F1 Score is the harmonic mean of recall and precision. It provides a single number that balances both metrics, useful when comparing different review approaches or tracking performance over time. The formula is: F1 = 2 x (Precision x Recall) / (Precision + Recall). An F1 score of 1.0 represents perfect recall and perfect precision.

Richness (also called prevalence) measures the proportion of responsive documents in the population. In plain terms, richness answers the question: "How dense with responsive material is the population?" If a population contains 1,000,000 documents and 50,000 are responsive, richness is 5%. Richness is not a measure of review performance; it is a characteristic of the data. However, richness profoundly affects how other metrics are interpreted, as explained below.

Elusion measures the proportion of responsive documents in the discard set (documents classified as non-responsive). In plain terms, elusion answers the question: "What proportion of documents we discarded were actually responsive?" If the review classifies 800,000 documents as non-responsive and 8,000 of those are actually responsive, elusion is 1%. Elusion is one approach to validating review performance,

particularly useful for workflows where most documents are classified as non-responsive.[45]

2. How These Metrics Relate

The metrics described above are interrelated, and understanding these relationships is crucial to interpreting validation results.

High recall with low precision means the review found most of the responsive documents but also swept in many non-responsive documents. This pattern often occurs when classification thresholds are set too low. The producing party is unlikely to miss responsive documents, but review costs may be higher than necessary, and the production may contain many documents that should not have been produced.

High precision with low recall means the documents classified as responsive are mostly correct, but many responsive documents were missed. This pattern often occurs when classification thresholds are set too high. The production may be incomplete, although the documents that are produced are more likely to be responsive.

Recall requires knowing the total number of responsive documents in the population. Since exhaustively determining this number is impractical, recall is typically estimated through one of several approaches:

- *Control set validation*: In TAR 1.0 workflows, a randomly selected control set coded independently of training provides direct estimates of recall and precision.

- *Elusion-based estimation*: Sampling from the discard set (documents classified as non-responsive) and combining that elusion rate with a richness estimate allows recall to be calculated indirectly.

- *Direct sampling of the full population*: Random sampling across the entire document population, with qualified reviewers assessing each sample document, provides estimates of both richness and overall performance.

Each approach has tradeoffs in terms of cost, complexity, and precision. The choice depends on the workflow, the population characteristics, and the validation goals.

Note on elusion and recall: A common interpretive error is treating elusion as the inverse of recall. For example, assuming that a 3% elusion

45. *See* Relativity, *RelativityOne Review Center Guide*, 35–37 (2026), https://help.relativity.com/PDFDownloads/R1_PDF/RelativityOne%20-%20Review%20Center%20Guide.pdf [https://perma.cc/SE7J-VZZG] (providing detailed technical guidance on how validation statistics, including elusion, are calculated in practice).

rate implies 97% recall is incorrect because elusion and recall have different denominators. Elusion is measured against the discard set, while recall is measured against all responsive documents in the population. In low richness populations, even a small elusion rate can correspond to substantially lower recall once the size of the discard set and the estimated prevalence of responsive documents are taken into account.

3. Worked Example

To illustrate how these metrics work together, consider the following scenario:

Scenario: A population contains 1,000,000 documents. Richness is 10%, meaning 100,000 responsive documents exist. The review classifies 100,000 documents as responsive and 900,000 as non-responsive (the discard set). Validation sampling finds that 55% of the responsive set is actually responsive (precision = 55%), and 5% of the discard set is responsive (elusion = 5%).

Calculating recall from elusion: Responsive documents in the discard set equal 5% multiplied by 900,000, which is 45,000. Responsive documents found equal 55% multiplied by 100,000, which is 55,000. Recall equals 55,000 divided by 100,000, which is 55%. A second check here is that the total number of responsive documents found (55,000) and missed (45,000) sum to the estimated number of responsive documents in the population. In practice, these numbers will not sum perfectly, but they should be close.

This example illustrates that the same elusion rate can correspond to different recall levels depending on richness and population proportions. A 5% elusion rate might sound low, but in this scenario, it corresponds to a recall rate of 55%.

C. *Statistical Sampling: How We Measure*

The metrics described in Section B[46] cannot be calculated exactly for large document populations. No legal team can review every document to determine the "true" recall or precision of a review process. Instead, these metrics are estimated through statistical sampling.

This Section explains the fundamentals of statistical sampling as applied to document review validation. It covers random sampling, sample size determination, stratified sampling, and practical considerations for designing defensible validation samples.

46. *See* discussion *supra* Part II.B.

1. Random Sampling Basics

A random sample is a subset of documents selected from a population such that every document has an equal probability of being selected. Random sampling is the foundation of statistical inference because it allows conclusions drawn from the sample to be generalized to the population with quantifiable uncertainty.

Simple random sampling selects documents without regard to any document characteristics. Each document in the population has the same chance of being included in the sample. Simple random sampling is appropriate when the population is relatively homogeneous or when there is no reason to believe that different subgroups have meaningfully different characteristics.

Stratified random sampling divides the population into subgroups (strata) and draws separate random samples from each stratum. This approach is appropriate when the population contains identifiable subgroups that may have different responsiveness rates or error patterns. For example, a validation sample might be stratified by custodian, date range, document type, or review method.

The choice between simple and stratified sampling depends on the population and the validation goals. Stratified sampling can provide more precise estimates when strata differ meaningfully, but it adds complexity. For most end-of-review validation scenarios, simple random sampling across the entire reviewed population is sufficient and recommended.

2. Sample Size Determination

The precision of a sample estimate depends on the sample size. Larger samples yield more precise estimates. The question is: how large must the sample be to achieve acceptable precision?

Sample size determination involves three parameters. First, the confidence level is the probability that the confidence interval contains the true population value. A 95% confidence level is standard in most contexts, meaning that if the sampling process were repeated many times, 95% of the resulting confidence intervals would contain the true value. Second, the margin of error is the acceptable range around the sample estimate. A margin of error of plus or minus 5% means the true value is likely within five percentage points of the sample estimate (at the stated confidence level). Third, the expected proportion is the anticipated proportion of the characteristic being measured. When estimating richness or elusion, this is the expected percentage of responsive documents. When this proportion is unknown, using 50% produces the most conservative sample size.

For large populations, the sample size depends almost entirely on the desired precision, not the population size. However, a common

misconception in the discovery community requires clarification: the oft-cited figure of "approximately 385 documents" for a 95% confidence interval of *recall* with plus or minus 5% margin of error refers to 385 *responsive* documents in the sample. This means the total sample size for tightly estimated recall depends critically on richness. This is a primary reason why it is difficult to provide tight statistical estimates or recall when richness is very low.

3. Estimating Recall in Low-Richness Populations

Low-richness populations (those in which responsive documents constitute a small percentage of the total) present particular challenges for sampling. When few responsive documents exist in the population overall, random samples will contain few responsive documents, making it difficult to estimate all three key metrics with tight confidence intervals. In these scenarios, recall remains the key metric to pay attention to, and additional options should be considered.

The sample sizes discussed in Section C.2[47] assume that the proportion being estimated is around 50%. At lower proportions, the same sample size yields fewer responsive documents and therefore wider confidence intervals. Consider a population with 2% richness: a random sample of 1,000 documents would contain only about twenty responsive documents on average. This affects validation metrics differently. For precision estimation, if the responsive set has high precision, say 90%, a sample will contain mostly responsive documents and precision can be estimated reasonably well. For sampling from the non-responsive set, if the proportion of responsive documents there is low, say 2%, a sample of 500 documents would yield only about ten responsive documents, producing a wide confidence interval. For recall estimation, uncertainty in the underlying proportions propagates to the recall estimate.

Practical approaches for low-richness sampling:

Option 1: Larger samples. In low-richness populations, substantially larger samples are required to find enough responsive documents for tight estimates (i.e., small confidence intervals). When richness is below 10%, samples of 1,000 to 3,000 or more documents are mathematically *necessary* to keep confidence intervals smaller than the commonly requested plus/minus 5%. Larger samples will provide narrower confidence intervals, which may be appropriate depending on the needs of the case and whether the additional costs are justified by the need for tighter estimates of review performance.

47. *See* discussion *supra* Part II.C.2.

Option 2: Accept wider confidence intervals. In some cases, acknowledging uncertainty is appropriate and cost-effective. Consider an example where an initial random sample estimated richness at approximately 2%, meaning approximately 10,000 responsive documents exist in a population of 500,000. The review classifies 8,500 documents as responsive. To estimate recall, we sample 1,500 documents from the 491,500 documents classified as nonresponsive and find 10 responsive documents. The elusion point estimate is 0.67%, with a 95% confidence interval of approximately 0.35% to 1.2%. Extrapolating to the full non-responsive set yields an estimated 3,290 missed responsive documents (with a range of 1,720 to 5,900). Using the known richness estimate of 10,000 total responsive documents, the recall point estimate is 67%, but the 95% confidence interval spans approximately 41% to 83%—a range of 42 percentage points.

This wide confidence interval is not due to an insufficient sampling effort; sampling 1,500 documents represents substantial validation work. Rather, the wide interval is an unavoidable consequence of low richness: when the non-responsive set contains few responsive documents, any reasonable sample will capture only a small number of them, and small counts produce wide confidence intervals. As discussed in Section B, practitioners should also remember that even the point estimate for recall can be counterintuitively low relative to elusion, a 0.67% elusion rate does not imply 99.3% recall. A validation report should transparently communicate both the point estimate and the confidence interval, explaining that the uncertainty reflects population characteristics rather than methodological deficiency.

Option 3: The Rule of Three. When a randomly drawn elusion sample constructed from the documents predicted to be non-responsive actually contains zero responsive documents, the "Rule of Three" provides an upper bound on the true proportion of responsive documents among those predicted non-responsive. To be 95% confident that the true proportion is at most X%, the required sample size is approximately 3 divided by X. For example, to confirm that a proportion of predicted non-responsive documents is at most 2% responsive, review approximately 150 predicted non-responsive documents; if none are responsive, the proportion is likely below 2%. To confirm the proportion of responsive documents is at most 1%, review approximately 300 of the documents predicted non-responsive. Practically, this approach is useful for confirming that elusion is below some threshold with 95% confidence, which is often more useful in practice than a point estimate.

The appropriate choice depends on the stakes of the matter, the available budget, and the level of precision required for defensibility.

4. Stratified Sampling

Stratified sampling divides the population into subgroups and samples each subgroup separately. This approach may be useful in high-stakes litigation where different subpopulations may have meaningfully different characteristics or may have been reviewed through different processes.[48]

Stratification adds complexity to both the sampling design and the analysis. For most discovery scenarios where validation is needed, simple random sampling from the entire document population will suffice. Stratification may be reserved for cases where there are specific reasons to believe different subgroups behave differently and where answering subgroup-specific questions is important to defensibility.

5. Designing the Validation Sample

Beyond sample size and stratification, several practical considerations affect the defensibility of a validation sample.

Independence from training data. If the review used machine learning or generative AI, the validation sample should be independent of any data used to train or tune the model. Using training data to validate performance produces optimistically biased estimates.

Blinding and conflict review. Reviewers assessing the validation sample should not know how the documents were classified by the review process during their initial pass. If a reviewer knows that a document was classified as non-responsive, this knowledge may bias the responsiveness assessment. Blinded review produces more objective initial assessments.

However, we also recommend that after the initial blinded pass, validators take a second, unblinded pass through the documents to review any conflicts between their assessment and the original classification. This allows validators to determine their "self-overturn" rate and, if appropriate, change their designation after seeing the original classification. Importantly, a log should be maintained of all documents where designations changed during this unblinded phase. This two-pass approach balances the objectivity of blinded review with the opportunity to catch and reconsider genuine errors.

Qualification of reviewers. Validation samples should be reviewed by individuals qualified to make responsiveness determinations. For complex matters, this typically means attorneys with subject matter knowledge. For more routine matters, other legal professionals trained on the review parameters may be acceptable. The key is that the validation

48. *See generally In re* Broiler Chicken Antitrust Litig., No. 1:16-cv-08637, 2018 WL 1146371 (N.D. Ill. Jan. 3, 2018) (Order Regarding Search Methodology for Electronically Stored Information).

reviewer must be capable of making accurate responsiveness assessments.

Documentation. The validation process should be documented contemporaneously. Record the sample selection method, the sample size, the reviewers, the assessment criteria, the overturn numbers, and the results. This documentation supports defensibility if the production is challenged.

D. *Workflow-Specific Best Practices*

Sections A through C[49] provided the conceptual foundation: what to measure and how to measure it. This Section translates those concepts into practical guidance for specific review workflows. Each subsection addresses a common review approach, explains its characteristics, and provides recommended practices for quality control and validation.

The recommendations in this Section are best practices, not requirements. Different matters warrant different approaches, and legal teams should adapt these recommendations to their specific circumstances. The common thread is that every workflow should include mechanisms for measuring and demonstrating performance.

1. Linear Human Review

What it is: In linear human review, attorneys or paralegals review documents one by one, typically working through queues organized by saved searches or other criteria. Documents are coded for responsiveness and other issues based on reviewer judgment. This is the traditional approach to document review and remains common, particularly for populations requiring nuanced legal judgment.

Characteristics: Linear human review is thorough but slow and expensive. Quality depends heavily on reviewer training, consistency, and attention. Studies have shown that even qualified reviewers miss responsive documents, so human review should not be treated as a "gold standard" against which other methods are measured.[50]

Best practices for quality control: Reviewer training should occur before review begins, training reviewers on the responsiveness criteria, privilege issues, and coding instructions. Quality control sampling should occur early on and should involve a senior reviewer or subject matter expert sampling and re-reviewing a percentage of each reviewer's work. Common quality control rates range from 5% to 15% of documents, with

---

49. *See* discussion *supra* Part II.A–C.

50. Maura R. Grossman & Gordon V. Cormack, *Technology-Assisted Review in E-Discovery Can Be More Effective and More Efficient Than Exhaustive Manual Review*, 17 RICHMOND J. L. & TECH. 1, 37, 48 (2011) (study showing that "human assessors missed between 20% and 75% of all relevant documents").

higher rates for less experienced reviewers or more complex matters. Escalation protocols should establish clear criteria for escalating difficult documents to senior reviewers or subject matter experts. Reviewers should not guess on documents they do not understand. If there are substantial overturns during the quality control process, investigate the potential cause: unclear criteria, inadequate training, or inherently ambiguous documents, and then remediate accordingly.

Additionally, overturn rates of specific first-pass reviewers may be monitored to identify those who demonstrate deficiencies in review protocol training or understanding of criteria. If additional training does not address consistently higher than average overturn rates, it may be appropriate to cease the reviewer's continued participation in making review decisions.

Validation approach: For linear human review, although rarely done in practice, validation should assess precision, recall, and elusion. Sample from the responsive set to estimate precision (are the documents coded as responsive actually responsive?), and sample from the non-responsive set to estimate elusion (what proportion of discarded documents are actually responsive?). The assumption that human review of every document ensures low miss rates is not supported by empirical evidence. Research has shown that human reviewers can miss a substantial proportion of responsive documents; one study found that human reviewers missed approximately half of the responsive documents that were identified by a generative AI system.[51]

2. TAR 1.0 (Control Set/Simple Learning)

What it is: TAR 1.0, sometimes called "simple learning," is a workflow which uses a *seed set* of coded documents to *train* a machine learning model. After training occurs, a separate set of documents, called a *control set*, is used to *test* the model's predictive accuracy.

The initial seed set is typically created by drawing a statistically valid sample of documents from the larger collection. Known responsive documents identified by other means can also be added to the seed set for training.

A control set is created by selecting a separate random sample of documents that are coded by subject matter experts (SME's). These SME coding decisions, however, are specifically withheld from the training of the machine learning model. Only the coding decisions of documents in the seed set are delivered to the machine learning model for training.

After the documents in the seed set are coded and delivered to the machine learning model for training, the model computes the predicted

51. Omrani et al., *supra* note 15, at 6–7 (reporting that human reviewers had substantially lower recall—approximately 55%—compared to 96% for the generative AI system).

coding values of all other documents in the collection—including those of the control set. At this point, the machine's predicted coding values of the control set documents are compared to the SME's actual coding decisions of the control set. If the machine's predicted coding decisions do not reliably match the SME coding decisions, additional iterations of training are conducted.

Additional training is accomplished by selecting more documents to be added to the seed set and coded for responsiveness by the SMEs. A variety of methods can be used to select the additional documents for coding and inclusion in the seed set. Documents can be selected based on the predictive values calculated by the current state of training. These can be documents currently assigned with high confidence predictive values at the extremes of the responsive/non-responsive ranking, or documents with predictive rankings in between, indicating the machine's comparative inability to discriminate between responsive and non-responsive.

Additionally, documents can be selected for supplemental training based on judgmental selection—for example, choosing documents based on the subject matter of inconsistent coding observed in an earlier seed set comparison, or identified by the presence of specific terms or phrases likely to identify responsive content. The numerical size of coded documents in the seed set can also be increased by the inclusion and coding of an additional random sample of new documents.

In practice, combinations of these methods to select documents for additional training are typically used. After each iteration of additional training, the model's predicted coding values for the seed set are again compared to the SME's actual coding decisions of the seed set. With each round of iterative training, the model's predictive accuracy generally improves. When the model can consistently predict SME coding decisions with acceptable accuracy, training ends and the final predictive values of all documents in the collection become static. A linear review of all documents down to a cut-off range is conducted. Documents below the cut-off range—those reliably predicted to be non-responsive—can be defensibly omitted from linear review, creating substantial economic savings.

Characteristics: Because of the manual steps involved, a TAR 1.0 review workflow requires a review manager with substantial training and proficiency in the underlying document management platform. A TAR 1.0 workflow may be less efficient than a continuous learning workflow (described below) because a smaller set of documents informs the model. The primary advantage of a TAR 2.0 continuous learning workflow is that the model is continuously updated with the latest reviewer coding decisions, allowing it to learn from a larger and more diverse set of training examples over time.

Best practices for implementation: Seed set selection and size should include documents representative of the responsive population. Larger seed sets generally improve model performance, but improvements in accuracy can be observed to plateau and hit diminishing returns as the numerical size of the seed set is increased over the course of iterative training rounds. A common approach is to start with a random sample plus judgmentally selected "key" documents. Control set design requires the control set to be selected randomly and coded independently of the seed set. It should be large enough to provide statistically meaningful performance estimates. Rank cutoff determination involves reviewing documents in rank order and monitoring the rate at which responsive documents are found. When the rate drops below a threshold, such as when precision falls below some target, consider establishing a cutoff. Documents below the cutoff become the discard set.

Validation approach: TAR 1.0 offers multiple validation options. The control set, if sufficiently large and reserved exclusively for final validation, can provide direct recall and precision estimates. However, if the control set is examined during the seed set development process, for example, to decide whether to add more training documents, it may no longer serve as an unbiased measure of final performance. In such scenarios, it is common practice to draw and label a fresh simple random sample for post-review validation.

When end-of-review validation discloses some type of systemic quality failure, such as a subject matter coverage gap, one remediation technique is to re-open training, ingest all coded documents into the seed set, and then identify and code any additional documents ranked above the cut-off. Alternatively, a lower cut-off point threshold may need to be selected.

3. TAR 2.0 (Continuous Learning)

What it is: Continuous learning, sometimes called TAR 2.0, uses a rapid iterative process where the machine learning model continuously trains and learns from all reviewer coding decisions. Unlike TAR 1.0, where training ends at a discreet point in the workflow, a TAR 2.0 workflow continuously trains and re-computes the predictive ranking of all documents as reviewers make sequential coding decisions. With some exceptions,[52] a TAR 2.0 workflow creates a review queue comprised of the highest-ranked uncoded documents for review next. This "continuous learning" approach focuses reviewer effort on documents that the model

52. Well-conceived TAR 2.0 AI tools deliver a mix of randomized and targeted content to the review queue to prevent early review decisions from causing subject matter clusters to be overlooked.

predicts are most likely to be responsive or most likely to promote the accuracy of the algorithm.

Characteristics: Continuous learning is highly efficient for finding responsive documents quickly. By continuously presenting the most promising documents, it typically achieves high recall with fewer documents reviewed than linear review or TAR 1.0 workflows. Determining when to stop can be a moving target because the model continuously improves.

Best practices for implementation: Initial document selection should begin with a random sample or judgmental sample to provide initial training data. Richness estimates from an initial random sample help set expectations. Continue reviewing until the rate of finding responsive documents drops substantially. Because continuous learning tends to prioritize high-scoring documents for sequential review, it creates a risk of missing responsive documents that the model scores poorly throughout the process. Periodically sample from lower-ranked documents to check for systematic gaps. Stopping criteria are typically based on the declining rate of responsive documents found during review. When the yield becomes very low, the model has likely identified most of the responsive documents it can find. The stopping rationale should be documented.

Validation approach: For continuous learning, the discard set (low-scoring documents that were not reviewed) is the primary concern for validation because the model guides review toward high-probability documents. Sampling approaches include elusion testing (sampling the discard set to assess the proportion of responsive documents missed) or a coverage sample from across the predictive score distribution to identify whether the model has systematic blind spots. The validation approach should be tailored to the specific workflow and defensibility requirements.

4. Generative AI Document-by-Document Classification

What it is: Generative AI classification uses large language models to classify documents for responsiveness. Unlike traditional machine learning, which uses statistical patterns learned from labeled examples, generative AI can apply natural language instructions (prompts) that describe the responsiveness criteria. The model can also extract verifiable citations and generate rationales explaining why each document is or is not predicted as responsive.

Characteristics: Generative AI classification is a newer approach with promising early results. The ability to apply natural language prompts without extensive labeled training data is appealing, particularly when training data is scarce or expensive. The rationale-generation feature aids quality control because reviewers can assess not only the

classification but also the reasoning. However, generative AI models are not infallible, and their performance can vary depending on the prompt design and document characteristics. Current generative AI classification models exhibit minor non-deterministic variations between runs. In practice, generative AI classification systems using large language models configured for greedy decoding, where the model always selects its most likely response rather than introducing randomness, show greater consistency and self-agreement than human review teams.

Establishing defensibility through measurement and transparency: Defensibility is established by the measurement of the system's performance and the transparency in its use, not by the choice of which system to deploy. The choice of review methodology (whether human linear review, discriminative machine learning, or generative AI) is a decision based on cost, efficiency, and appropriateness for the task. The courtroom question is whether the results can be measured and demonstrated to be reasonable. The *TAR 1 Reference Model* article provides the framework: regardless of whether the underlying algorithm is discriminative machine learning or generative AI, the process of building a predictive model and demonstrating its effectiveness through sampling and statistics remains constant.[53]

The three-phase workflow: A key advantage of generative AI classification is the ability to validate performance before committing to full-scale review. Traditional methods typically require material up-front costs (training reviewers, building models through iterative labeling) before performance can be validated. Generative AI, by contrast, requires fewer up-front costs before an estimate of review performance can be calculated, enabling a more proactive approach.

Develop. Subject matter experts write and iterate on the prompt criteria that define responsiveness. They test the prompt on a small, diverse sample of documents, examine the model's predictions and rationales, and adjust the prompt until the model's classifications align with expected outcomes. As demonstrated in recent case studies, this iterative prompt development process involves tracking two key goals: increasing true positives (finding more responsive documents) and reducing false positives (decreasing incorrect predictions of responsiveness).[54]

Validate. Before running the model at scale, evaluate the prompt against a statistically rigorous random sample from the target document

---

53. *See TAR 1 Reference Model*, *supra* note 25, at 117.

54. *See* Tara S. Emory, *Case Study: Practical Testing for Effectively Organizing Documents with GenAI Review*, EDRM (June 24, 2025) (describing iterative prompt refinement process focused on "adjusting our prompt to find more of the relevant documents (i.e., increasing true positives)" and "decreasing aiR's incorrect predictions of responsiveness for documents that were not responsive (i.e., reducing false positives)").

population. Calculate recall, precision, and elusion with confidence intervals. If the metrics meet the validation plan criteria established in Section A,[55] proceed. If not, return to the development phase and refine the prompt.

Apply. Once validated, apply the prompt to the full document population with confidence that performance will be consistent with the validation results.

This proactive validation approach offers significant advantages. Problems are discovered before the cost of full-scale review is incurred. If validation does not achieve the desired results, the team can iterate on the prompt and revalidate without wasting effort on a full review that would need to be redone. From a defensibility standpoint, proactive validation allows the producing party to demonstrate expected performance before committing resources.

Best practices for implementation: Iterative prompt development requires clear, specific prompts modeled on the kind of detailed review protocol that would be provided to human reviewers. Test initial prompts on a small sample, examine disagreements between the model and subject matter experts, and refine the prompt to address gaps or ambiguities. If a control set is used, "the prompt writer should not be the same person who reviewed control documents and should not be exposed to information from the control set's contents, to avoid contaminating the objectivity and correctness of the control set."[56]

Proactive validation before scale should occur once the prompt is developed, validating on a random sample from the full target population before running at scale. A random sample of approximately 1,000 documents is typically sufficient to estimate recall and precision with reasonable confidence intervals, assuming adequate richness.

Threshold selection involves determining how to handle documents where the model is uncertain. Options include setting a threshold below which documents receive human review or treating "uncertain" as a separate category.

Human-in-the-loop for edge cases should establish criteria for routing documents to human reviewers, including documents where the model expresses low confidence, documents with unusual characteristics, or documents in categories known to be difficult.

Rationale review for quality control involves periodically reviewing the model's rationales, not just its classifications. Rationales that rely on incorrect facts or flawed reasoning signal problems even if the ultimate classification happens to be correct. LLM rationales and citations are useful quality control artifacts, but they are not guarantees of correctness;

---

55. *See* discussion *supra* Part II.A.

56. *See TAR 1 Reference Model*, *supra* note 25, at 124.

explanations may be unfaithful and may increase reviewer overreliance. Therefore, rationales should be treated as additional insights into the model's internal behavior and labeling process, and not as substitutes for rigorous sampling-based validation.

Revalidation triggers should be recognized because validation represents the accuracy of a specific model configuration on a defined document population. If the prompt criteria change, if the underlying model is updated, or if the document population changes substantially, revalidation may be necessary.

Hybrid approaches: Generative AI classification often works best when used in combination with other methods. As practical experience has demonstrated, different legal and factual issues have different characteristics that align better with different review methods. Issues requiring nuanced content analysis may perform well with generative AI, while issues with rules-based components (such as multiple criteria or date dependencies) may benefit from other review methods.[57] Strategic review design involves understanding each tool's capabilities and matching them to the characteristics of different review needs.

Validation approach—proactive plus end-of-review: The recommended approach for generative AI classification involves two validation stages: proactive validation and end-of-review validation. Proactive validation occurs before running the model at scale, as described above. This validates that the prompt and model are performing adequately on the target population and provides confidence before committing to full-scale review. End-of-review validation is still recommended in most scenarios: document populations often change during the course of the review process, and the understanding of the matter itself may evolve. Final validation is always recommended in multiphase and hybrid workflows. End-of-review validation is particularly important when the document population changed during review, when the generative AI results are combined with other review methods in a hybrid workflow, when the producing party wants to confirm that applied results match validation expectations, or when compliance or opposing party requirements call for final production validation. In many situations, after running the model at scale, counsel will perform quality control and may review certain documents, changing the designations from the original predictions. End-of-review validation makes sense when you want to validate the overall production as opposed to just the generative AI classification, which may have changed during human quality control.

---

57. *See* Emory, *supra* note 54 ("Our testing showed that GenAI review worked better for some issues than others. Issues requiring nuanced content analysis generally performed well. Issues with a rules-based component such as multiple criteria (e.g., mentioning a company plus discussing a topic), or dependent on date, had more limited performance.").

5. Hybrid and Multimethod Workflows

What it is: In many large matters, legal teams may use different review methods for different document populations. Hybrid workflows route different document subsets to different review methods based on practical considerations such as data source, custodian groupings, document type, timeline, budget, or available resources.

Characteristics: Hybrid workflows can provide flexibility in managing large-scale review projects. However, they introduce complexity in tracking, quality control, and validation. The primary challenge is ensuring that all portions of the population are validated and that the overall production is defensible.

Best practices for implementation: Document routing logic should define the criteria for routing documents to each review method before review begins. Routing criteria might include custodian, document type, date range, or data source. Document the routing logic so it can be explained. Consistent coding fields should be used across all review methods so that the final production can be unified. If different review streams use different field names or coding values, consolidation becomes error-prone. Quality control within each stream should ensure that each review method has its own quality control appropriate to that method. Human review should have quality control sampling; machine learning should have control sets or coverage checks; generative AI should have sample validation.

Validation approach: For hybrid workflows, the recommended approach is a single random sample across the entire document population at the end of review. This unified validation answers the question that matters most: across all methods combined, what are the recall and precision of the overall production?

This approach is simpler than validating each method independently and reflects how courts and opposing parties will evaluate the production: as a whole, not method by method. If the unified validation reveals problems, more detailed investigation can then determine which method or methods caused the issue.

In some cases, legal teams may want to validate individual methods in addition to the unified validation. This can be useful for internal process improvement or when different methods have very different expected performance levels. However, unified validation should be the default, with method-specific validation as an optional additional step.

E. *End-of-Review Validation: The Final Gate*

End-of-review validation is the moment when the legal team answers the question: did our review work? It is the critical defensibility gate

where the producing party generates the evidence needed to demonstrate that the production is complete and accurate.

This Section provides specific guidance on conducting end-of-review validation. It describes the recommended standard approach, addresses sample size considerations, explains how to respond if validation fails, and discusses documentation requirements.

1. The Standard Approach (Recommended)

For most matters, end-of-review validation should follow a straightforward process.

Step 1: Create a saved search or other mechanism that encompasses all reviewed documents. The validation must cover the entire population that will be the basis for production. If documents were reviewed through multiple methods, the saved search should include all of them. The search defines the universe from which the validation sample will be drawn.

Step 2: Pull a single random sample from the full document population. Draw a random sample from the entire reviewed population. The sample should include documents classified as responsive and documents classified as non-responsive in proportion to their presence in the population. Alternatively, if the goal is specifically to estimate precision and elusion separately with similar confidence, draw separate samples from the responsive set (to estimate precision) and the non-responsive set (to estimate elusion).

Step 3: Have qualified reviewers make independent responsiveness decisions. Reviewers assessing the validation sample should be qualified to make responsiveness determinations. They should apply the same responsiveness criteria used during the review. Importantly, validation reviewers should not know how the documents were originally classified during their initial assessment. After the initial blinded pass, reviewers should then conduct an unblinded review of any conflicts between their assessment and the original classification, with the opportunity to change their designations if warranted. A log of all designation changes during the unblinded phase should be maintained. This two-pass approach accounts for the adjudication phase in designation disputes between review systems while preserving auditability and transparency of the validation.

Step 4: Calculate recall, precision, and elusion. Using the validation sample results, calculate precision as responsive documents in responsive sample divided by total responsive sample size, elusion as responsive documents in non-responsive sample divided by total non-responsive sample size, and recall using the elusion rate, richness, and population

proportions as described in Section B.[58] Report these metrics with their confidence intervals.

Step 5: Document results and rationale. Record the validation methodology, sample sizes, reviewer qualifications, and results. This documentation supports defensibility. If the results meet the completion criteria established in the validation plan (see Section A),[59] the review is complete. If they do not, remediation is required.

2. Sample Size Guidance

As discussed in Section C,[60] sample size depends on three parameters: the desired confidence level, the acceptable margin of error, and the expected proportion being estimated. The appropriate sample size varies significantly depending on these factors, and practitioners should determine their sample size using these specific inputs rather than relying on general guidelines.

Sample size calculators are available from academic institutions and statistical organizations. To use these calculators effectively, practitioners should have an estimate of the expected proportion, for example, the richness estimate obtained early in review as discussed in Section A.[61] Enter the desired confidence level (typically 95%), the acceptable margin of error (e.g., plus or minus 5%), and the expected proportion into the calculator to determine the required sample size.

For low-richness populations, refer to the guidance in Section C.[62]

3. What if Validation Does Not Yield the Desired Results?

Validation may reveal that the review did not meet the completion criteria. This is disappointing but can provide valuable information. It is far better to discover problems through validation than to have opposing counsel discover them later.

Remediation options: Additional review may be needed if elusion is too high. Options include lowering the score cutoff, reviewing additional random samples from the non-responsive set, or targeted review of categories where errors are concentrated. Prompt or model adjustment may be needed for generative AI workflows where validation failures may indicate problems with the prompt or model behavior. Re-review by different reviewers may be warranted if precision is too low (many nonresponsive documents classified as responsive), as the issue may be with reviewer calibration.

---

58. *See* discussion *supra* Part II.B.
59. *See* discussion *supra* Part II.A.
60. *See* discussion *supra* Part II.C.
61. *See* discussion *supra* Part II.A.
62. *See* discussion *supra* Part II.C.

Revalidation after remediation: After remediation, repeat the validation process on the affected population. Do not assume that remediation solved the problem; verify it.

Documentation for defensibility: When the initial validation does not achieve the anticipated results, the metrics and remediation steps should be documented. A production that was remediated and subsequently met validation criteria may be more defensible than one that was never validated at all. The documentation shows that the legal team took quality seriously and corrected problems when they were discovered.

Renegotiation options: It may be prudent to avoid pre-review commitment to specific numerical validation results; if an aspirational commitment is made, a process for modifying the numerical value for good cause shown may be included. The characteristics of document sets vary widely and can be unpredictable. Parties may be well served by discussing acceptable completion criteria while leaving open the possibility of modifying those criteria due to issues that may arise during the review, or as more is learned about the underlying dataset. If the initial validation does not achieve the anticipated results, but the efforts to remediate may be too costly, burdensome, or not proportional, or may not yield much better results (i.e., "juice not worth the squeeze"), counsel may agree to accept the validation results, even though they were not as originally agreed upon as the level for "completion." This is when it becomes important for the parties to be reasonable, understanding the limitations of document review and statistics. Ongoing collaboration and measured transparency of process between requesting and producing parties tend to promote the successful resolution of unanticipated outcomes.

4. Documenting the Validation

Validation documentation serves two purposes: it supports defensibility if the production is challenged and it provides a record for future reference.

What to record: The validation plan, including completion criteria and target metrics; the definition of the validation population; the sample selection method (e.g., random, stratified, etc.) and sample size; the qualifications of validation reviewers; the responsiveness criteria applied; the raw results, including the number of responsive and nonresponsive documents in each sample; the calculated metrics with confidence intervals; any remediation performed and subsequent revalidation results; and the final conclusion indicating whether the review met completion criteria.

What to produce versus protect as work product: The distinction is between explaining what was done and how well it worked (which

supports defensibility) and revealing litigation strategy and internal assessments (which are protected). A validation summary that reports methodology and results without revealing privileged details is appropriate for disclosure.

The validation methodology and results are typically producible if challenged. They demonstrate the reasonableness of the review process. However, some details may be protected, including training decisions, the rationales of detailed error analysis, and internal quality discussions, which may be considered protected work product. When a requesting party challenges a review and production workflow with evidence strongly suggesting quality shortcomings, a court may require the producing party either to produce process documentation or to conduct additional remedial measures. The option of coming forward with process documentation showing good faith efforts to achieve quality will serve a producing party well under such circumstances.

F. *Presenting Results to the Courts and Opposing Counsel*

Validation is only useful if it can be communicated effectively. Legal teams must be prepared to explain their review process and its results to courts, opposing counsel, and clients. This Section provides guidance on what to disclose, what to protect, and how to respond to challenges.

The continued adoption of emerging AI technologies, and the beneficial efficiencies they bring to the administration of justice, are likely to be inhibited if potential users are overly concerned that their interactions with discovery technologies may be subject to compulsory disclosure to adverse parties. For example, crafting a textual prompt created for use by a generative AI review tool may expressly disclose the mental impressions and legal analysis performed by the producing party's counsel.

At the same time, producing parties using established AI workflows routinely make *voluntary* but measured decisions to disclose certain workflow and validation information, even when that information could be subject to a meritorious assertion of work product protection. This is typically done because disclosure and collaboration tend to strengthen a producing party's defensibility posture. When both the requesting party and producing party essentially "risk share" with agreements concerning workflow criteria and validation protocols, it may decrease the likelihood that a producing party will be compelled to undertake an expensive review "do-over" should systemic problems emerge in the production.

For parties electing to take a conservative approach to disclosure and collaboration, validation documentation and validation metrics become even more important as part of risk mitigation strategy. The following factors inform a producing party's decisions concerning what information to disclose and what information to protect. The goal is

transparency about methodology and results without revealing work product that necessarily merits protection.

1. What to Disclose

When defending a production, the producing party may elect to disclose information sufficient to demonstrate that the review was reasonable.

Workflow description (high level): Describe the review methods used: attorney review, paralegal review, machine learning, generative AI, or some hybrid combination. Explain the general approach without revealing specific strategies.

Validation methodology: Describe how the review was validated. Include the sample selection method (random sampling from the entire population), the sample size, and who conducted the validation review. Explain that validation reviewers applied the responsiveness criteria to a sample of documents and that results were used to estimate recall, precision, and elusion.

Metrics with confidence bounds: Report the validation results with confidence intervals. For example: "Based on a random sample of 1,200 documents from the reviewed population, we estimate precision at 85% (with a 95% confidence level and a confidence interval of 82% to 88%) and recall at 78% (with a 95% confidence level and a confidence interval of 74% to 82%)." Including confidence intervals demonstrates statistical rigor and honest acknowledgment of uncertainty.

Remediation, if any: If initial validation revealed problems that were corrected, describe the remediation process and subsequent revalidation. This demonstrates that the legal team took quality seriously and addressed issues when discovered.

2. What to Protect

Certain information generated by the documentation of a review workflow may become the subject of disputes over claims of work product protection, even when validation results are disclosed.

Training decisions and rationales: For machine learning and generative AI, the specific examples used to train the model, the reasoning behind prompt design, and the coding rationale may be the subject of a claim for protection.[63]

Detailed error analysis: Internal analysis of why certain documents were missed or miscoded, and what those errors reveal about case weaknesses, may be the subject of a claim for protection. While the validation results themselves (aggregate metrics) are disclosable, the

63. *See* Hickman v. Taylor, 329 U.S. 495, 508 (1947); FED. R. CIV. P. 26(b)(3).

strategic assessments of those results are usually claimed as work product.

The principle is that the producing party demonstrates the reasonableness of its process through methodology and aggregate results, not by disclosing the substance of attorney decision-making.

3. When Expert Consultation is Warranted

In high-stakes matters or when facing sophisticated challenges, expert consultation may be appropriate. Experts in information retrieval, statistics, or e-discovery can provide independent validation, explain complex statistical concepts to courts, and testify about industry standards. Expert involvement is not required for routine matters, but it can strengthen defensibility when the stakes justify the cost.

G. *Summary of Recommendations*

This Section summarizes the key recommendations from Part II. These are best practices, not requirements. Different matters warrant different approaches, and legal teams should adapt these recommendations to their specific circumstances. However, these principles provide a foundation for defensible, statistically rigorous document review validation.

1. Define Completion Criteria Before Starting Review

Do not wait until the end of review to decide what "done" means. Establish target metrics, acceptable confidence levels, and validation procedures at the outset. This validation plan should be documented. Having clear criteria prevents disputes about whether the review was adequate and provides a framework for consistent decision-making throughout the process. Provide a mechanism for cooperation if the completion criteria are not attained or needs to be modified.

2. Choose Metrics Appropriate to Case Needs

Recall, precision, and elusion are all relevant, but their relative importance varies by matter. High-stakes litigation involving core evidence may warrant different targets than a routine matter with lower stakes. The proportionality principle should apply to validation efforts as being part of the scope of discovery. There is no "one size fits all" numerical value of a validation metric.

3. Understand How Metrics Relate and Select an Appropriate Validation Approach

Recall, precision, and elusion each measure different aspects of review performance. When using elusion-based estimation, understand

that elusion is not the inverse of recall—a 3% elusion rate can correspond to recall of 55%, 75%, or 95% depending on richness and the relative sizes of the responsive and nonresponsive sets. Select a validation approach (control set, elusion sampling, full-population sampling, or a combination) that fits the workflow and provides the evidence needed for defensibility.

4. Design Validation Samples with Appropriate Statistical Rigor

Sample sizes should be sufficient to provide meaningful precision at stated confidence levels. For low-proportion estimates like elusion, standard sample sizes may be inadequate. Use larger samples or the Rule of Three approach when validating that elusion is below a threshold.[64] Document the sampling methodology.

5. Use the Simplest Method Mix That Meets Your Outcome Requirements

Complex workflows with multiple review methods can optimize cost and quality, but complexity adds risk. Simpler approaches are easier to explain, validate, and defend. Add complexity only when it provides clear benefits.

6. For Multimethod Reviews, Validate with a Single Sample Across the Full Corpus

When different portions of the document population are reviewed using different methods, the recommended default is a single random sample from the entire population. This unified validation answers the question courts and opposing parties care about: how well did the overall production perform? Method-specific validation can be added if specific questions need to be answered, but unified validation should be the primary approach.

7. Document Workflow and Validation Sufficient to Explain, Not Over-disclose

Create contemporaneous records of the review methodology, validation procedures, and results. Be prepared to disclose methodology and aggregate results to demonstrate reasonableness and defensibility. Make deliberate work product assessments of training decisions, prompts, and strategic analysis. The goal is transparency about process and results without revealing attorney judgments that may be subject to protections.

64. *See* discussion *supra* Part II.C.3.

8. Test Early to Catch Problems Before the End of Review

Do not wait until review is complete to discover that the approach is not working. For machine learning and generative AI, monitor performance during the review process. For human review, conduct ongoing quality control sampling. Early detection of problems allows course correction while there is still time and budget to address them.

9. For Generative AI, Build Defensibility Through Validation and Documentation

The validation framework that supports traditional TAR workflows applies equally to those using generative AI tools. Focus on demonstrating that the review achieved acceptable performance levels, regardless of the technology that was used. Defensibility will be established on a case-specific basis through testing, statistical validation, human oversight, and transparent documentation.

10. Recognize That No Review is Perfect, and Perfection is Not the Standard

Research consistently shows that human reviewers disagree on borderline documents and that defensible methodologies do not necessarily achieve 100% recall. The legal standard is reasonableness, not perfection. A well-validated production with documented methodology and honestly reported results is defensible even if some responsive documents were missed. The question is whether the search was reasonable given the circumstances, and validation provides the evidence to answer that question.

CONCLUSION

This Article has sought to address a gap in the literature and jurisprudence regarding AI-assisted document review: the absence of clear, practical guidance on how to measure and defend the completeness and accuracy of a document production. The framework presented in Part II, built upon the foundation established in Part I, provides legal teams with the tools they need to demonstrate, with statistical rigor, that their review process was reasonable and their production defensible.

The core insight is straightforward: the defensibility of a production depends not on which technology was used, but on whether its performance can be measured and explained. A producing party that can demonstrate, through sound statistical methods, that its review achieved acceptable recall and precision levels is well-positioned to satisfy its discovery obligations, regardless of whether attorneys, paralegals, machine learning, or large language models conducted that review.

For generative AI in particular, the absence of judicial precedent specifically addressing its use does not preclude defensible deployment. As the *TAR 1 Reference Model* establishes, the conceptual steps of TAR workflows remain constant regardless of the underlying algorithm.[65] The validation framework that has supported traditional TAR workflows for over a decade applies equally when generative AI tools are used: systematic testing, statistical sampling, human oversight, and transparent documentation of methodology and results. Legal teams that master these principles will be equipped to adopt new technologies with confidence, knowing that their defensibility rests on demonstrated performance rather than judicial approval of specific tools.

As AI continues to evolve, the technologies available for document review will change. The fundamental principles of statistical validation, however, will remain constant, and the legal teams that master these principles will be equipped to defend whatever methods they choose to employ.

65. *See TAR 1 Reference Model*, *supra* note 25.

APPENDIX A

This appendix presents a detailed case scenario demonstrating the principles discussed in Part II. The scenario involves a hybrid workflow using multiple review methods and unified end-of-review validation.

Case Context

A large corporate litigation involves the collection of 1,000,000 documents from multiple custodians, spanning emails, contracts, chat messages, and financial records. The producing party has a court-ordered deadline to review and produce all responsive documents within 90 days. The overall richness of the population is estimated to be 25%. Given the volume, budget constraints, and tight timeline, the legal team implements a hybrid review workflow that leverages different review methods for different document populations.

Example Three-Method Workflow

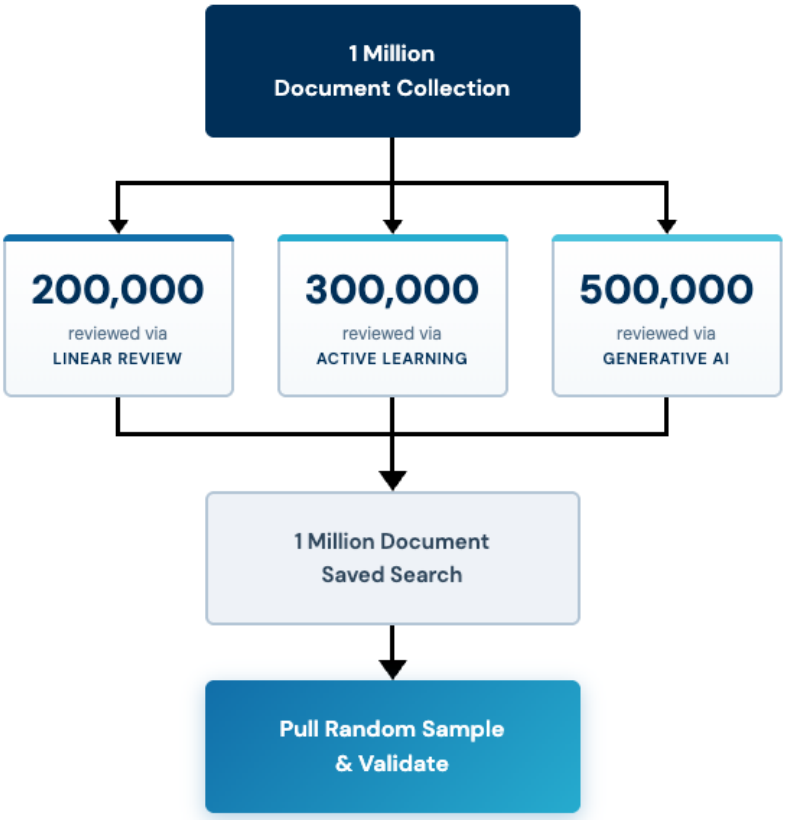

Method 1: Linear Review (200,000 Documents): The legal team assigns a population of 200,000 documents, organized by custodian group or data source, to human reviewers working through saved search queues. Reviewers code documents for responsiveness and flag potential privilege issues.

Method 2: Active Learning (300,000 Documents): A second population of 300,000 documents is reviewed using active learning. Attorneys train the model through iterative queues, with the system continually reprioritizing documents based on reviewer coding decisions. This approach surfaces likely responsive documents earlier in the review process.

Method 3: Generative AI Classification (500,000 Documents): The remaining 500,000 documents are classified using generative AI based on iteratively refined prompts developed by subject matter experts. Before full deployment, the prompts were validated on a random sample to confirm acceptable performance levels.

Validation Approach

Following the recommendation in Section E, the legal team conducts unified end-of-review validation with a single random sample across the entire 1,000,000-document population.

Step 1: Create a saved search encompassing all reviewed documents across all three methods.

Step 2: Given the estimated 25% richness, a random sample of 1,600 documents from the full population is expected to yield 400 responsive documents, which should be sufficient to estimate recall with a 95% confidence interval of less than 10%. Given the population proportions, this sample will naturally include documents from each review method in rough proportion to their share of the population.

Step 3: Qualified reviewers (senior attorneys not involved in the original review) independently assess each sampled document for responsiveness, without knowing how the document was originally classified.

Step 4: Calculate metrics based on validation results.

Sample Validation Results

Suppose the validation sample yields the following results: Total sample of 1,600 documents; documents originally classified as responsive is 448; documents originally classified as non-responsive is 1,152; responsive sample found 350 responsive documents (precision =

78%); non-responsive sample found 50 responsive documents (elusion = 4.3%).

Of the 400 responsive documents in the sample, 350 were correctly identified by the review process and 50 were missed. Recall estimated directly from the sample equals 350 divided by 400, which is 87.5%.

Using the formula from Section B: Total responsive in population (estimated) equals 25% multiplied by 1,000,000, which is 250,000. Documents classified as non-responsive are approximately 720,000. Responsive documents in the discard set (estimated) equals 4.3% multiplied by 720,000, which is approximately 31,000. Responsive documents correctly identified equals 250,000 minus 31,000, which is 219,000. Recall equals 219,000 divided by 250,000, which is 87.6%—consistent with the sample-based estimate.

Precision: 78% ± 3.8% (based on 448 documents from the responsive-classified set)

Elusion: 4.3% ± 1.2% (based on 1,152 documents from the non-responsive-classified set)

Richness: 25% ± 2.1% (based on 1,600 documents from the full population)

Recall: 87.5% ± 3.3% (based on 400 responsive documents in the sample)

Validation summary: The unified validation estimates precision at 78% ± 3.8% (n=448) and recall at 87.5% ± 3.3% (n=400), with elusion at 4.3% ± 1.2%. These results indicate that the hybrid workflow performed well across all three methods.

Defensibility Analysis

The producing party can demonstrate defensibility as follows. First, statistical rigor: validation used a random sample of 1,600 documents to calculate several key metrics with defined confidence intervals. Second, unified assessment: the single-sample validation approach answers the question courts care about—overall production quality, not method-by-method performance. Third, documentation: the workflow and validation were documented contemporaneously. Fourth, the generative AI component specifically: prompts were developed iteratively and validated before full deployment, human oversight was maintained through sample review, and performance was confirmed through statistical testing.

The court is likely to accept this hybrid methodology because the producing party can demonstrate that the review process was reasonable

and defensible, maintaining reasonable assurance of completeness and accuracy through validated results.

APPENDIX B: GLOSSARY OF TERMS

Active Learning: A machine learning approach where the model identifies which documents would be most informative to code next, focusing human review effort on documents that will most improve model performance.

Confidence Interval: A range of values that, with stated probability (typically 95%), contains the true population value. A 95% confidence interval means that if sampling were repeated many times, 95% of the resulting intervals would contain the true value.

Confidence Level: The probability that a confidence interval contains the true population value. Common levels are 90%, 95%, and 99%.

Continuous Learning (TAR 2.0): A technology-assisted review protocol in which the machine learning model continuously trains on reviewer coding decisions and dynamically re-ranks all documents, prioritizing the review queue with those predicted most likely to be responsive or to improve model accuracy.

Control Set: In TAR 1.0 workflow, a random sample of documents coded independently of the training set, used to test and estimate model performance.

Elusion: The proportion of true responsive documents in the discard set (the set predicted to be non-responsive). Elusion measures what proportion of discarded documents were actually responsive and serves as the primary mechanism for validating review performance through sampling.

F1 Score: The harmonic mean of precision and recall, providing a single balanced performance metric. Formula: F1 = 2 x (Precision x Recall) / (Precision + Recall).

Generative AI: Artificial intelligence systems, typically based on large language models, that generate text in response to prompts. In document review, generative AI can classify documents based on natural language instructions describing responsiveness criteria.

Large Language Model (LLM): A deep learning model trained on enormous amounts of text that can generate coherent text in response to prompts. LLMs form the foundation of generative AI document classification.

Margin of Error: The range around a sample estimate within which the true value is likely to fall. A margin of error of plus or minus 5% means the true value is probably within 5 percentage points of the estimate.

Precision: The proportion of documents classified as responsive that are actually responsive. Precision measures how accurate the positive classifications are.

Prevalence: See Richness.

Recall: The proportion of all responsive documents in the population that were identified as responsive. Recall measures how much of the responsive material was found.

Richness: The proportion of responsive documents in a population. Also called prevalence. If a population has 1,000,000 documents and 50,000 are responsive, richness is 5%.

Rule of Three: A statistical principle for low-proportion estimation. When a sample contains zero instances of a characteristic, the 95% upper confidence bound is approximately 3/n, where n is the sample size.

Seed Set: In TAR 1.0 workflows, the set of coded documents used to train the model.

Simple Random Sampling: A sampling method where every document in the population has an equal probability of being selected, with no stratification.

Stratified Sampling: A sampling method where the population is divided into subgroups (strata) and separate random samples are drawn from each stratum.

TAR (Technology-Assisted Review): The use of computer machines to extrapolate human judgments made about one set of documents to a different set of documents not reviewed by humans.

TAR 1 Reference Model: Described in: Tara Emory, Jeremy Pickens & Wilzette Louis, T*AR 1 Reference Model: An Established Framework Unifying Traditional and GenAI Approaches to Technology-Assisted Review*, 25 SEDONA CONF. J. 109; A framework establishing that technology-assisted review follows five steps (Scope, Label Control Set,

Iterate Model, Classify, Validate) regardless of the underlying algorithm, whether discriminative machine learning or generative AI.

Validation: The process of measuring review performance through sampling and assessment, typically including precision, recall, and elusion estimation.

Zero-Shot Learning: The ability of a model to make predictions about a class of interest from a natural language description, without labeled training examples. Generative AI document classification typically operates through zero-shot learning.

PROMPTING CHANGE WITH GENAI IN LARGE-SCALE DOCUMENT REVIEW: A REAL-WORLD STUDY

Robert Keeling,[*] *Ray Mangum*,[**] *Amy Hanke*[***] & *Alyssa Ogden*[****]

INTRODUCTION

The evolution of document review in civil litigation has reached another milestone—one that invites a new partnership between humans and technology that promises gains in accuracy, efficiency, and scalability. Generative artificial intelligence (genAI) is remaking business and professional life at extraordinary speed,[1] and, within the legal sphere, nowhere is its impact more immediate—or the potential changes more dramatic—than document review.

Yet significant questions remain. When does genAI meaningfully improve document review? How should it be integrated into established workflows? How should its outputs be validated? As with prior technologies that are now widely accepted in eDiscovery, genAI's broader adoption will depend on demonstrated reliability and value supported by empirical data.

---

* Partner, Redgrave LLP.
** Partner, Redgrave LLP.
*** Counsel, Redgrave LLP.
**** Counsel, Redgrave LLP.

1. *See, e.g.*, Charlie Campbell et al., *The Architects of AI Are TIME's 2025 Person of the Year*, TIME MAG. (Dec. 11, 2025), https://time.com/7339685/person-of-the-year-2025-ai-architects/ [https://perma.cc/GX7T-MCLX] (noting the rapid advancement of generative AI technology).

Document review in civil litigation has undergone significant transformations before.[2] In the 1990s and earlier, review was largely manual: teams of attorneys examined thousands (and sometimes millions) of paper documents in warehouses over months. The early 2000s ushered in "e-discovery," as digital document collections and keyword searching became central to review and production. In the 2010s, technology-assisted review (TAR) (a.k.a. "predictive coding") introduced more sophisticated tools based on machine-learning—along with much debate about acceptable use, validation, and defensibility. With each evolution came novel issues, skepticism, and—ultimately—broad adoption.[3]

Today, the claim is that genAI can perform large-scale document review more effectively, economically, and defensibly than earlier approaches. The profession needs more empirical data to test that claim—to evaluate whether genAI truly delivers on its potential across varying types of legal matters and document sets. Among the key questions:

- Does genAI accurately and consistently predict relevant vs. non-relevant documents better than earlier technologies?
- Does genAI outperform or at least match human reviewers?
- Is genAI a defensible review methodology under legal discovery standards?
- Are there use cases where genAI does not perform well?
- What are the actual time and cost savings?
- How can lawyers best integrate genAI into existing workflows?
- What risks, if any, arise from using genAI in document review?

In this Article, we seek to answer some of these questions with empirical data, providing insight into the use of genAI to assist with document review in large-scale legal and regulatory matters.

For this case study, we worked with a client company and an eDiscovery vendor to select a 1,600-document sample from the client's prior, real-life legal matter. We then compared the relevance and issue coding of human reviewers to the coding decisions made by a market-leading genAI tool, Relativity's aiR for Review. The results varied depending on document type, but overall, the relevance review produced a high recall rate of 83.9% and a precision rate of 84.7%. The issue-coding results were more mixed, with accuracy often depending on the

2. *See, e.g.*, Maura R. Grossman & Gordon V. Cormack, *Technology-Assisted Review in E-Discovery Can Be More Effective and More Efficient than Exhaustive Manual Review*, 17 RICH. J.L. & TECH. 11 (2011); George L. Paul & Jason R. Baron, *Information Inflation: Can the Legal System Adapt?*, 13 RICH. J.L. & TECH. 10 (2007).

3. *See generally* Charlie Hernandez, *Tech Shifts & the Law*, 48-AUG L.A. LAW 26 (July/Aug. 2025).

nuances of the specific issue. The testing also revealed certain limitations of the tool. Based on these results, this Article concludes with recommendations for potential genAI use cases within a typical document review workflow and discusses the potential cost savings, including the factors that may affect that calculus.

Part I of this Article describes the evolution of document review methodologies from manual review to TAR and concludes with the rise of genAI and its early impact on the legal profession. In Part II, we discuss the need for empirical data to evaluate the utility and defensibility of genAI tools in document review, and we describe our proof-of-concept study. In Part III, we offer analysis of the study's results and provide our observations, learnings, and recommendations on how lawyers and their clients can best leverage genAI to enhance efficiency, improve accuracy, and make informed decisions about integrating those technologies into their existing document review workflows.

I. THE EVOLUTION OF TECHNOLOGY IN LAW MARCHES ON WITH GENAI

The evolution of technology in law provides a useful starting point for understanding both the hesitation and the promise surrounding state-of-the-art genAI tools in legal document review. Each technological advancement in the progression—from digital keyword searching to TAR—has sparked a mix of optimism and skepticism within the legal community. Technology adoption in law has historically been slow and steady, and the past serves as a helpful guide to what the inevitable adoption and integration of genAI in eDiscovery may look like.[4]

The origins of discovery in civil litigation began with manual, paper-based document review—an arduous and labor-intensive endeavor that was often inefficient, costly, and highly prone to human error.[5] This approach, now largely obsolete, is chiefly remembered by attorneys who practiced before the digital transformation of the twenty-first century.[6]

---

4. *See* Hernandez, *supra* note 3, at 28 (tracing the historical patterns of technology adoption in law and concluding that it demonstrates a predictable cycle of "initial fear and resistance, followed by increasing client pressure, eventual regulatory guidance, and ultimately, widespread acceptance").

5. *See* Dana A. Remus, *The Uncertain Promise of Predictive Coding*, 99 IOWA L. REV. 1691, 1702–03 (2014) (citing Maura R. Grossman & Gordon V. Cormack, *Technology-Assisted Review in E-Discovery Can Be More Effective and More Efficient Than Exhaustive Manual Review*, 17 RICH. J.L. & TECH. 1, 3 (2011)); Maura R. Grossman & Gordon V. Cormack, *Inconsistent Responsiveness Determination in Document Review: Difference of Opinion or Human Error?*, 32 PACE L. REV. 267 (2012)); The Sedona Conference, *The Sedona Conference Best Practices Commentary on the Use of Search and Information Retrieval Methods in E-Discovery*, 8 SEDONA CONF. J. 189, 199 (2007).

6. *See* Bennett B. Borden & Jason R. Baron, *Finding the Signal in the Noise: Information Governance, Analytics, and the Future of Legal Practice*, 20 RICH. J.L. & TECH. 7, 3 (2014) (describing "the beginning" as manual review, a process of "legions of lawyers with hundreds if

By the early 2000s, the digital transformation was well underway—organizations across industries had adopted electronic solutions for record-keeping and data storage (e.g., cloud computing), communication (e.g., email), and transacting business (e.g., electronic contracts and e-signatures).[7] The result was an explosion of electronically stored information (ESI) that made continued reliance on traditional manual document review impractical. Despite uncertainty and concerns about the "black box" of technology, this unprecedented growth in ESI compelled the legal community to evolve its thinking and practices.[8]

Regulatory and judicial guidance lagged behind the rapid shift to digital business processes. Between 2003–2004, the landmark decisions in *Zubulake v. UBS Warburg LLC* established the foundation of modern-day ESI preservation obligations by recognizing that parties have certain duties to preserve and produce relevant ESI.[9] The court also addressed concerns that the overwhelming volume of ESI could render litigation cost-prohibitive by offering guidance on proportionality and cost-sharing principles for handling large-scale electronic data production.[10] Shortly thereafter, the 2006 amendments to the *Federal Rules of Civil Procedure* formally recognized ESI as a discoverable category of evidence and established rules specifically addressing electronic discovery and some of the challenges it raised.[11] Moreover, in 2012, the American Bar Association revised Model Rule 1.1 related to the duty of competence to expressly recognize the role of eDiscovery, requiring that lawyers stay

---

not thousands of boxes in warehouses, reviewing folders and pages one-by-one in an effort to find the relevant needles in the haystack").

7. Hernandez, *supra* note 3, at 28–31.

8. Hernandez, *supra* note 3, at 31 (noting that "law firms were forced to play ball" following the "rapid digitization" of the early 2000s); Kate Bauer, *Technology-Assisted Review: Overcoming the Judicial Double-Standard*, RICH. J.L. & TECH. BLOG (Jan. 24, 2018), https://jolt.richmond.edu/2018/01/24/technology-assisted-review-overcoming-the-judicial-double-standard/ [https://perma.cc/T3AB-8RGY] (calling for greater acceptance of TAR over manual review in light of "increasing document volumes and research on the shortcomings of traditional review methods"); Paul E. Burns & Mindy M. Morton, *Technology-Assisted Review: The Judicial Pioneers*, THE SEDONA CONF. INST. (Mar. 2014), https://www.americanbar.org/content/dam/aba/publications/litigation_committees/commercial/materials/technology-assisted-review-the-judicial-pioneers.pdf [https://perma.cc/3KWN-ZMHL] (noting plaintiffs' objections in the *Da Silva Moore* case that they do not understand the "black box" of predictive coding and "there is no way to be certain if MSL's method is reliable").

9. Zubulake v. UBS Warburg LLC, 229 F.R.D. 422, 431 (S.D.N.Y. 2004); *but see* Robert Keeling, *Sometimes, Old Rules Know Best: Returning to Common Law Conceptions of the Duty to Preserve in the Digital Information Age*, 67 CATH. U. L. REV. 67, 102 (2018) (indicating that there is no duty under traditional common law to preserve until a lawsuit is filed or imminent).

10. *See* Zubulake v. UBS Warburg LLC, 217 F.R.D. 309, 317–18 (S.D.N.Y. 2003).

11. *See* Burke T. Ward et al., *Electronic Discovery: Rules for a Digital Age*, 18 B.U. J. SCI. & TECH. L. 150, 179 (2012).

apprised of "the benefits and risks associated with relevant technology."[12] At this point, using electronic solutions such as keyword searches and applying metadata filters to a document set had become the norm.[13]

But as lawyers and their teams of document reviewers struggled to keep up with the ever-growing volumes of ESI, new technology solutions using predictive coding emerged.[14] Computer-assisted review based on supervised machine learning—at the time often called "predictive coding" and today more commonly referred to as technology-assisted review (TAR)—is a process whereby machine learning algorithms are trained by human reviewers to predict the relevance of documents within large data sets and then applied to entire document populations to classify documents, improving both speed and accuracy over manual review and keyword searching.[15]

As before, the practical application of TAR outpaced the development of judicial and regulatory guidance. In 2012, the landmark case of *Da Silva Moore v. Publicis Groupe* became the first federal court decision to approve the use of TAR in document review.[16] In support of his holding that a producing party could use TAR in appropriate cases, Magistrate Judge Peck explained the challenges of manual review and keyword searching, examined the empirical data in support of integrating TAR, and concluded that TAR was better than the alternatives in that case.[17] The defensible use of TAR in document review was fortified in 2015 by *Rio Tinto PLC v. Vale S.A.*, which confirmed as black letter law that TAR is an acceptable discovery tool and that a producing party does not need prior approval from the opposing party or the court to implement it.[18]

---

12. MODEL RULES OF PRO. CONDUCT r. 1.1 cmt. 8 (A.B.A. 2020); A.B.A. Comm'n on Ethics 20/20, *Resolution and Report on Technology and Confidentiality 105A* 3 (Aug. 2012); *see also* Lori D. Johnson, *Navigating Technology Competence in Transactional Practice*, 65 VILL. L. REV. 159, 168 (2020) (explaining the significance of the ABA's adoption of amended Model Rule 1.1 regarding the duty of competence to include technological competence).

13. *See* The Sedona Conference, *The Sedona Conference Best Practices Commentary on the Use of Search and Information Retrieval Methods in E-Discovery*, 8 SEDONA CONF. J. 189, 200 (2007) ("By far the most commonly used search methodology today is the use of 'keyword searches' of full text and metadata as a means of filtering data for producing responsive documents in civil discovery.").

14. *See* Robert Keeling et al., *Using Machine Learning on Legal Matters: Paying Attention to the Data Behind the Curtain*, 11 HASTINGS SCI. & TECH. L.J. 9 (2020); *see also* Grossman & Cormack, *supra* note 2, at ¶ 28 (offering evidence that technology-assisted review yields superior results to manual review); *see also* Robert Keeling et al., *Separating the Privileged Wheat from the Chaff–Using Text Analytics and Machine Learning to Protect Attorney-Client Privilege*, 25 RICH. J.L. & TECH. 2, ¶¶ 33–34 (2019).

15. *See* Charles Yablon & Nick Landsman-Roos, *Predictive Coding: Emerging Questions and Concerns*, 64 S.C. L. REV. 633, 634 (2013).

16. 287 F.R.D. 182, 193 (S.D.N.Y. 2012).

17. *Id.* at 189–91. *See also* Bauer, *supra* note 8.

18. Rio Tinto PLC v. Vale S.A., 306 F.R.D. 125, 126–27 (S.D.N.Y. 2015).

Critically, *Rio Tinto* further confirmed that TAR should not be held to a higher standard than other document review methodologies, such as manual review or keyword searching.[19]

Over the next decade TAR became a routine part of large-scale document review workflows, with widespread adoption by law firms and corporate legal departments as both courts and clients recognized its increased efficiency, accuracy, and cost-effectiveness in comparison to prior methods.[20] Courts continued to provide guidance as new questions arose in the context of applying TAR, such as the responding party's discretion to use TAR, the effect of ESI protocols on a party's decision to use TAR, transparency and disclosure requirements, TAR methodologies and workflows, proportionality, and cost-shifting.[21] For example, following the acceptance of TAR, a line of cases held that the producing party typically has discretion to decide whether to use TAR and which TAR methodology to employ, as long as the process is reasonable.[22] Subsequent caselaw grappled with technical aspects of using TAR, as well as the use of TAR in conjunction with other review methodologies.[23]

Each stage of technological evolution in discovery has been marked by an increase in the volume and complexity of data, the straining of existing review methodologies, skepticism about the emerging technology, validation of that technology through experience and studies, acceptance within the judiciary and legal community, and, finally, normalization of the technology into standard workflows.[24] GenAI presents the next inflection point. The unrelenting growth of ESI and the advent of new data sources are fueling this continued evolution. With genAI, companies and their counsel are seeking more sophisticated eDiscovery tools to manage overwhelming volumes of diverse and

---

19. *Id.* at 129.

20. *See generally* The Sedona Conference, *The Sedona Conference TAR Case Law Primer, Second Edition*, 24 SEDONA CONF. J. 1 (2023) [hereinafter *TAR Case Law Primer*].

21. *Id.* (citing cases).

22. *See id.* at 27–29 (citing Livingston v. City of Chicago, No. 16 CV 10156, 2020 WL 5253848 at *3 (N.D. Ill. Sep. 3, 2020); Coventry Cap. US LLC v. EEA Life Settlements Inc., No. 17-Civ. 7417 (VM) (SLC), 2020 WL 7383940, at *4 (S.D.N.Y. Dec. 16, 2020), *objections overruled*, 2021 WL 961750 (S.D.N.Y. Mar. 15, 2021); Lawson v. Spirit AeroSystems, Inc., No. 18-1100-EFM-ADM, 2020 WL 1813395, at *8–9 (D. Kan. Apr. 9, 2020); Kaye v. N.Y.C. Health and Hospitals Corp., No. 18-CV-12137 (JPO) (JLC), 2020 WL 283702 (S.D.N.Y. Jan. 21, 2020)).

23. *TAR Case Law Primer*, *supra* note 20, at 48–62.

24. *See, e.g.*, Hernandez, *supra* note 3, at 31–32 (analyzing the historical trajectory of technology adoption in law: "necessity ultimately drives acceptance, with judicial clarification and regulatory adaptation cementing these innovations into standard legal practice"); Matthew G. Kenney, *The Past, Present and Future of Predictive Coding*, 12 FLA. A&M U.L. REV. 165, 176–78 (2016) (analyzing the "slow adoption rate" of predictive coding in document review by lawyers even after judicial acceptance in *Da Silva Moore*).

complex data under tight deadlines in large-scale litigation and regulatory investigations.

The results of a recent study confirm that in-house legal teams see the potential benefits of genAI and are driving adoption at a rapid pace.[25] According to a 2025 survey by the Association of Corporate Counsel of 657 in-house legal professionals from 30 countries, the shift from "passive planning to active implementation" of genAI is underway and moving swiftly to embrace the integration of genAI into legal matters.[26] The 2025 survey shows that genAI adoption by U.S. legal departments has more than doubled since 2024, with 52% of respondents reporting that they are already using genAI in their legal practice.[27] Only 2% of respondents were neither using nor planning to use genAI in 2025.[28] In terms of its perceived benefits, 91% of respondents stated that efficiency was the primary benefit of incorporating genAI into their legal matters.[29] Notably, however, the study found that in-house legal departments are not yet seeing the cost-savings benefit of genAI from outside counsel.[30]

While genAI promises even greater efficiency and cost savings than the technology solutions that preceded it, there is work to be done within the legal community to ensure its acceptance and successful integration in eDiscovery. Empirical research will be useful in further demonstrating its effectiveness.[31] Legal uncertainty adds another layer of complexity: courts and regulatory bodies are only beginning to address AI-generated outputs. The Sedona Conference and other legal organizations have begun issuing guidance on responsible AI use in eDiscovery, highlighting best practices.[32]

As the legal community evaluates genAI, its acceptance will likely depend on the availability of empirical research, judicial guidance, and standards to address new risks. Unlike prior tools that classified

---

25. ASS'N OF CORP. COUNS., *Generative AI's Growing Strategic Value for Corporate Law Departments – Survey Results* (Oct. 14, 2025), https://www.acc.com/resource-library/generative-ais-growing-strategic-value-corporate-law-departments-survey-results [https://perma.cc/2NNF-KTGZ].

26. *Id.*

27. *Id.*

28. *Id.*

29. *Id.*

30. *Id.*

31. *See* Maura R. Grossman et al., *Does the LLMperor Have New Clothes? Some Thoughts on the Use of LLMs in eDiscovery*, NAT'L L. REV. (Nov. 4, 2024), https://natlawreview.com/article/does-llmperor-have-new-clothes-some-thoughts-use-llms-ediscovery [https://perma.cc/V2QY-EXHY] (concluding that empirical studies are necessary to demonstrate the effectiveness of LLM tools in eDiscovery).

32. *See* The Sedona Conference, *Primer on Generative AI in Discovery, Draft* (2025); The Sedona Conference, *The Sedona Canada Primer on Artificial Intelligence and the Practice of Law*, 26 SEDONA CONF. J. 99 (2025); Judge Xavier Rodriguez, *Artificial Intelligence (AI) and the Practice of Law*, 24 SEDONA CONF. J. 783 (2023).

documents based on large amounts of training data, genAI tools allow lawyers to describe what they are looking for in natural language. Moreover, many genAI tools provide narrative explanations and other outputs to support coding decisions, which allows reviewers to understand and evaluate the basis for the decision. The distinctions between these tools raise new questions about reliability and efficiency that cannot be answered simply by analogy to TAR and the body of evidence supporting its use. To that end, this Article presents a study that provides empirical data to help evaluate the use of genAI tools in large-scale document review, including an evaluation of their accuracy, efficiency, and limitations.

II. DATA-DRIVEN GENAI ADOPTION: INSIGHTS FROM AN INNOVATIVE LEGAL STUDY

A key aspect of the adoption of prior technologies in eDiscovery has been the existence of empirical studies to support the effective and defensible use of the emerging technologies. The primary example is Magistrate Judge Peck's seminal 2012 *Da Silva Moore* decision, which expressly approved of using computer-assisted review methodologies in discovery in appropriate cases and made such acceptable use black letter law.[33] This judicial endorsement was a big step in establishing predictability for litigants and lawyers that the use of TAR and similar technologies would be legally defensible, which helped to solidify integration of the technology as a routine part of eDiscovery plans for large-scale document review.[34]

A significant factor in Judge Peck's endorsement of TAR was the existence of empirical data showing an increase in accuracy and the cost-saving benefits of using TAR over manual human review in voluminous cases.[35] Based on his review of several data-driven studies, Judge Peck concluded: "Computer-assisted review appears to be better than the available alternatives, and thus should be used in appropriate cases."[36]

Now, as market demands and ever-growing data volumes fuel the push for rapid adoption of genAI technology in the law, there is a need for more data-driven research to support the defensibility and judicial acceptance of using genAI tools, as well as to guide the responsible integration of genAI-based technologies into document review

33. Da Silva Moore v. Publicis Groupe, 287 F.R.D. 182, 183 (S.D.N.Y. 2012)*; see also* Rio Tinto PLC v. Vale S.A., 306 F.R.D. 125, 127 (S.D.N.Y. 2015) (citing *Moore* and subsequent caselaw for the proposition that the acceptance of TAR is black letter law).

34. *Moore*, 287 F.R.D. at 193 ("Counsel no longer have to worry about being the 'first' or 'guinea pig' for judicial acceptance of computer-assisted review.").

35. *Id.* at 190–91 (citing studies).

36. *Id.* at 191.

workflows.[37] This type of empirical support will help build the foundation necessary to demonstrate the effectiveness, accuracy, and cost-efficiency of genAI, fostering greater trust among clients, lawyers, and the judiciary, and encouraging broader acceptance in the legal community.

Given the relative infancy of genAI legal tools, there is currently little empirical data evaluating genAI in the law.[38] To help fill this void, the authors[39] conducted a study that incorporated actual client data from a prior legal matter to evaluate the performance of a genAI tool to conduct first-level document review for responsiveness and issue coding in a large-scale document review. The study results highlight the capabilities of genAI, as well as identify some challenges and limitations. The study offers valuable guidance for the legal community by assessing the effectiveness of genAI tools in document review, outlining best practices for their integration into current workflows, and providing insights that inform future use cases and workflow strategies.

A. *The Study Protocol*

In our practice, the authors have leveraged cutting-edge genAI tools to assist with document review in several active client matters. The genAI review tools are powered by large language models and simulate human document review. Use cases are proliferating, but common workflows include relevance review, issue review, privilege review, and key document identification. For each analysis type, the user enters prompts analogous to the review protocol in a traditional review workflow. The tool then analyzes the extracted text of the documents according to the prompt and provides predictions. Many genAI tools also provide the rationale for its score prediction and citations within the document that support the prediction. The results can be used in a variety of ways, ranging from replacing first-level review, identifying documents to prioritize for human review, or conducting quality control of human reviewers.

---

37. *See* Grossman et al., *supra* note 31 (concluding that empirical studies are necessary to demonstrate the effectiveness of LLM tools in eDiscovery).

38. *Id.* ("As far as we are aware, the impact of this phenomenon on eDiscovery search has neither been researched nor reported. . . . No study has yet shown either approach to be superior to state-of-the-art TAR methods."). *See also* Eugene Yang et al., *Beyond the Bar: Generative AI as a Transformative Component in Legal Document Review*, IEEE INT'L CONF. ON BIG DATA (Dec. 2024) (presenting empirical data from a legal matter evaluating an LLM-based document review system, reporting 96% recall and 60% precision without matter-specific tuning).

39. Redgrave LLP is a law firm specializing in information law, including eDiscovery, information governance, AI governance, data privacy, and cybersecurity. *Our Practice*, REDGRAVE LLP, https://www.redgravellp.com/our-practice [https://perma.cc/KVB7-65UZ] (last visited Mar. 1, 2026).

Working with an eDiscovery vendor, the authors initiated a study to assess the real-world performance of genAI tools. The study protocol was designed to evaluate genAI's ability to code documents for relevance and issue coding in an apples-to-apples comparison against human reviewers, considering both accuracy of coding decisions and efficiency from a cost and time perspective.

To facilitate this study, the authors worked with a corporate client to identify a prior closed matter in which the company was responding to a subpoena from a regulatory agency. During the actual review, attorney reviewers coded the document set for responsiveness and applied issue coding that correlated with the specific requests outlined in the subpoena. The document population consisted of custodial data (primarily emails and attachments), as well as documents collected from certain workpaper documents (including Word documents and Excel spreadsheets). Because of the nature of the subject matter of the workpaper documents, cover pages were created during collection of those documents for the purpose of providing context to the attorney reviewers in the underlying legal matter. The workpaper documents and associated cover pages were structured as a family during the human review.

The first step in the study protocol called for segregation from the review population of any documents that were deemed ineligible by the genAI tool. Like search terms and TAR, the genAI tool here works on the extracted text of a document. Thus, documents that do not contain extracted text (e.g., images or container files) and documents where the extracted text is too small or too large are ineligible to be processed through the tool. The study protocol excluded any documents without extracted text or with an extracted text size of under .05 KB or over 150 KB, resulting in approximately 6% of the documents being excluded as ineligible. Within the document population, the extracted text restrictions had little impact on emails and common attachment files, but they did eliminate several other files, including workpapers.

Once ineligible files were removed, the study protocol called for identification of statistical samples of documents previously coded by attorney reviewers, including a sample of 1,600 custodial documents, stratified by custodian and document type. The sample set represented all the possible issue codes, as well as a mix of responsive and nonresponsive coding. A similar, separate sample of 200 previously coded documents was identified for the population of workpapers.

Next, Redgrave attorneys drafted prompts to guide the responsiveness review. These initial prompts were based on the same information provided to attorney reviewers in the original document review protocol

in the underlying matter, with only minor revisions.[40] The study team then worked to iteratively refine the prompt by evaluating its performance on small samples of fifty documents, making revisions based on the results of each round.[41]

B. *GenAI Categorization*

The genAI tool used in this study (Relativity's aiR for Review) reports both a prediction and a numerical score for each analyzed document. The prediction is the relevance, key-document, or issue label aiR recommends; the score reflects how strongly relevant the document is, or how well it matches an issue. More specifically, the genAI tool scores documents from 0 to 4 according to their predicted level of relevance.[42] A score of 4 ("Highly Relevant") indicates direct, strong evidence that the document relates to the case or issue. A score of 3 ("Relevant") indicates the document is predicted relevant with citations to supporting text. A score of 2 ("Borderline Relevant") indicates that aiR found some content that might relate to the case or issue. A score of 1 ("Not Relevant") indicates aiR did not find evidence of relevance. A score of 0 ("Junk" or "Highly Not Relevant") indicates the document contains no useful information, such as system files or empty documents. A score of -1 indicates an error—the document either encountered an error or could not be analyzed. A table from Relativity's website summarizes its methodology:

Score	**Description**
4	**Highly Relevant**: The document is predicted very relevant to the issue. aiR found direct, strong evidence that the content relates to the case or issue. Citations show the relevant text.
3	**Relevant:** The document is predicted relevant to the issue. Citations show the relevant text.
2	**Borderline Relevant:** The document is predicted borderline relevant. aiR found some content that might relate to the case or issue. It usually has citations.

40. The genAI tool's (i.e., aiR for Review) prompt template for Responsiveness Analysis includes text boxes related to the following: matter overview, relevant people and aliases, noteworthy organizations, noteworthy terms, additional context, and responsiveness criteria.

41. During the time of the study, the authors understood that Relativity aiR products were using GPT-4o May 2024. Run time for sets of 50 documents was approximately 3 minutes, for both Relevance and Issue Analysis. Run time for the set of 1,600 was approximately 35 minutes.

42. *See Understanding Document Scores*, RELATIVITY https://help.relativity.com/RelativityOne/Content/Relativity/aiR_for_Review/aiR_for_Review_results.htm#Understanding_document_scores [https://perma.cc/EFW3-9B9Z] (last visited June 13, 2026).

1	**Not Relevant:** The document is predicted not relevant. aiR did not find any evidence that it relates to the case or issue.
0	**Junk**: The document contains no useful information or is considered "junk" data, such as system files, an empty document, or sets of random characters.
-1	**Error**: The document either encountered an error or could not be analyzed. For more information, see How document errors are handled.

At this point, two review project managers (the "conflict reviewers") from the underlying matter reviewed conflicts between the attorney coding and the genAI predictions, providing a final call on responsiveness and an explanation for the reasoning. Where the conflict review resulted in overturning the genAI tool's prediction, the study team discussed both the rationale provided by the genAI tool and the explanation provided by the conflict reviewers to determine whether the prompts should be revised.

In the first set of fifty documents, there were six conflicts, five of which were resolved in favor of the genAI tool. There were also eight documents predicted as borderline responsive, all of which were determined to be nonresponsive. The study team determined that iteration was required to address over-reliance on certain industry-based terms as indicia of responsiveness to a specific request in the subpoena. This was accomplished by revising the syntax of the relevant sentence in the prompts and adding to the relevance criteria that certain industry-specific concepts were not independently relevant.

The revised prompt was then run on a second set of fifty documents. On the second set, there were ten conflicts, four of which were resolved in favor of the genAI tool, and eight borderline documents. Following discussion of the conflicts, the study team determined that the prompt should be revised once more to address over-reliance on references to a single industry-specific keyword as indicia of responsiveness to the subpoena requests.

Once the study team determined that the prompt was sufficient, the team ran the genAI tool on the control set of 1,600 custodial documents. To ensure the reliability of the gold standard determinations, the study team implemented a rigorous validation process. Two attorneys who had led the original document review in the underlying legal matter served as conflict reviewers and reviewed the full 1,600-document sample.[43] These reviewers were given the review protocol and Q&A log from the original matter and were instructed to code each document for relevance without

43. The two conflict reviewers split the 1,600-document set between them without duplication in the review.

deference to either the aiR predictions or the prior attorney coding. Each reviewer provided written notes explaining the basis for coding decisions on each document. The resulting validation effort covered all 1,600 documents in the sample and provided the final human coding used as the "gold standard" for the responsiveness metrics reported below.

In addition to identifying responsive documents, genAI can also be used for issue analysis, categorizing documents based on a set of defined legal issues to allow for more in-depth review and case assessment. To test the efficacy of this use case, the refined prompt was adapted for use in the issue tagging analysis and run on the 1,600 custodial document sample, as well as a sample of fifty of the workpapers.[44] The conflict reviewers then reviewed conflicts from the custodial document sample, but as discussed below, following initial review of the results, the study team decided that a full conflict review of the workpapers for issue analysis was not warranted.

III. STUDY RESULTS SUPPORT INCORPORATION OF GENAI IN RESPONSIVENESS REVIEW

GenAI's application in legal document review is still in its early stages, and important questions remain about its efficacy and optimal use cases. This study provides meaningful data on several of those questions, though additional research will be needed as the technology and its applications continue to evolve. Our findings offer a foundation for informed conversation and a starting point for the continued empirical work that will shape genAI's role in eDiscovery.

A. *Responsiveness Analysis for Custodial Documents*

The results of our study demonstrate that genAI is an effective and reliable tool for use in responsiveness review workflows for electronic documents typically found in custodial collections (e.g., email and common attachments like Microsoft Word documents). With respect to responsiveness review, genAI returned more-than-adequate recall and very good precision, identifying actually responsive documents on par with attorney reviewers while generating fewer false positives than attorney reviewers. It also exceeded the widely accepted standards of 70-75% recall for TAR models.[45]

---

44. The Issue Analysis prompt template is the same as the responsiveness review except that the general "Responsiveness Criteria" text box is replaced with up to ten responsiveness criteria for particular issue tags. Here, we used the ten most common issue tags selected by attorney reviewers in the underlying legal matter.

45. Recall is the percentage of all relevant documents identified by the model, whereas precision is the percentage of documents the model indicated were relevant that actually were relevant. Recall and precision are usually inversely proportionate measures, meaning that as recall increases, precision usually decreases. Keeling, *supra* note 14, at 16–17. A commonly accepted

In the control set of 1,600 custodial documents, genAI predicted that 554 were responsive.[46] The gold standard review, conducted by two attorneys who led the original review and who reviewed the full 1,600-document sample, determined that 559 documents were in fact responsive.

Comparing genAI's predictions to the "gold standard" determinations, genAI achieved strong performance on definitive predictions. Of the 554 documents predicted responsive, 469 (84.7%) were confirmed responsive (true positives), while 85 (15.3%) were not responsive (false positives). Of the 600 documents predicted not responsive, 590 (98.3%) were confirmed not responsive (true negatives), while 10 (1.7%) were responsive (false negatives).

The gold standard reviewers also evaluated the 368 documents predicted as borderline responsive by genAI. Of those, 69 were determined responsive by the gold standard review, while the remaining 299 were determined not responsive. Additionally, genAI classified 62 documents as junk, of which 2 were determined responsive. Of the 16 documents for which genAI returned an error, 9 were determined responsive by the gold standard review.

As a result, excluding borderline documents from the analysis, the calculated recall rate was 83.9% (469/559) and the precision rate was 84.7% (469/554). In other words, genAI correctly identified approximately 84 out of every 100 responsive documents and, of the documents predicted by the genAI tool as responsive, approximately 85 out of 100 were in fact responsive. When borderline documents were included as responsive predictions, recall increased to 96.2% ((469+69)/559) but precision dropped to 58.4% ((469+69)/(554+368)). A table summarizing the results is below:

---

recall rate is 70–75%. Given the tradeoff between precision and recall, an acceptable precision rate typically depends on the needs of the case. *See* Neel Guha et al., *Vulnerabilities in Discovery Tech*, 35 HARV. J.L. & TECH. 581, 599–600 (2022) (citing Maura R. Grossman & Gordon V. Cormack, *Vetting and Validation of AI-Enabled Tools for Electronic Discovery*, *in* LITIGATING ARTIFICIAL INTELLIGENCE 407, 409 (Jill Presser, et al. Beatson & Gerald Chan eds., 2021) (discussing requirements on evaluation protocols)); *see also* Order Regarding Search Methodology for Electronically Stored Information at *6, *In re* Broiler Chicken Antitrust Litig., No. 16-cv-08637, 2018 WL 1146371 (N.D. Ill. Jan. 3, 2018).

46. All such documents had an "aiR Score" of 3. This score represents the genAI tool's confidence level on a scale of 0–4. Thus, all were predicted Relevant but none were predicted Highly Relevant.

Without Borderline		
Recall	83.9%	469/559 responsive documents identified
Precision	84.7%	469/554 predicted responsive documents actually responsive
With Borderline Documents Treated as Responsive		
Recall	96.2%	(469+69)/559 responsive documents identified
Precision	58.4%	(469+69)/(554+368) predicted responsive documents actually responsive

The baseline recall of 83.9% compares favorably to state-of-the-art performance metrics for traditional TAR models where a 70–75% recall rate is widely considered acceptable.[47] Precision of 84.7% for a TAR model would be considered highly effective at 70–75% recall and is especially impressive at 83.9% recall.

These results indicate that genAI can be a highly effective tool for identifying responsive documents within certain limitations of the tool. In addition to the text size limitation, which impacted approximately 6% of the document population, there are also limitations to the type of information genAI can consider. For example, most genAI tools (including the one used in this study) currently do not consider family relationships, the relationship between hyperlinked documents, or document metadata. Like TAR, the tool only reviews extracted text and therefore may not consider context in charts or graphs and may not consider text in images or handwriting if not accurately processed as text. In addition, the results show how including the so-called "borderline" documents can have a significant impact on efficacy and efficiency. Including the borderline documents can increase the number of responsive documents identified, while potentially materially increasing the number of false positives and thus increasing the cost of the review. A full summary of the data can be found below:

47. *See supra* note 45.

Table 1. Results of Custodial Documents Responsiveness Review Analysis

Custodial Documents Responsiveness Review	#	%
Sample population	**1600**	
Definitive predictions	1216	76.0%
Borderline predictions	368	23.0%
Errors[48]	16	1.0%
Definitive Population		
Disagreements between aiR and gold standard	95	7.8%
False positives (predicted R, gold standard NR)	85	15.3% of predicted R
False negatives (predicted NR, gold standard R)	10	1.7% of predicted NR
True positives (predicted R, gold standard R)	469	84.7% of predicted R
True negatives (predicted NR, gold standard NR)	590	98.3% of predicted NR
Junk documents (predicted junk, gold standard R)	2/62	3.2%
Borderline Population		
Determined responsive by gold standard	69/368	18.8%
Determined not responsive by gold standard	299/368	81.3%
Error Population		
Coded responsive by attorney reviewers	9/16	56.3%

48. On occasion, the genAI tool failed to return a result even when the document met the eligibility criteria. During execution, 26 documents errored out on the first run. Error details for 8 documents noted the "document text is too long" although none were over the suggested extracted text size. Details for one document listed "ungrounded citations detected in completion," which indicates that the genAI tool generated a citation that cannot be found in the document's text—suggesting a potential AI hallucination. Finally, 17 documents had unknown errors. Following the analysis, it was determined that running genAI a second time on the errored set eliminated 10 of the unknown errors. A third attempt did not eliminate any other errors.

Determined responsive by gold standard	9/16	56.3%
Determined not responsive by gold standard	7/16	43.8%
Overall Metrics		
Total true positives	469	
Predicted responsive and coded responsive	469	
Identified in conflict review	N/A	
Total false positives	85	
Total false negatives[49]	90	
Recall[50]	**469/559**	**83.9%**
Precision	**469/554**	**84.7%**

B. *Issue Analysis for Custodial Documents*

The results of genAI's issue analysis in the study were mixed and influenced by various factors, including document type, metadata fields, and the substantive issue being evaluated. For issue analysis, genAI was run on the same sample of 1,600 documents for the ten issue tags most often selected in the underlying review. Reviewers applied one or more of the ten issue tags to 468 of the 1,600 documents. As shown in Table 2 below, the rates at which genAI predicted that a document was relevant to an issue, as compared to human reviewers, varied across the issues.

The conflict reviewers reviewed the 380 documents for which genAI's prediction conflicted with attorney coding for one or more issue tags (23.8% of the sample). At a document level, conflicts for 167 documents (43.9%) were resolved fully in favor of genAI. Conflicts were resolved fully in favor of the reviewer for 126 documents (33.2%). For 87 documents (22.9%), there were multiple conflicts regarding multiple issues and the conflict reviewer agreed with some but not all of genAI's predictions.

49. Includes responsive documents predicted borderline, junk, or error.

50. Recall and precision are calculated assuming that borderline documents are excluded from production.

Table 2. Comparison of aiR Predictions to Reviewer Coding for Issue Tagging Analysis (Pre-Conflict Review)

	Issue # 1	**Issue # 2**	**Issue # 3**	**Issue # 4**	**Issue # 5**
aiR	93	247	100	122	23
Reviewer	27	363	90	220	57
Overlap	15	173	52	90	13

	Issue # 6	**Issue # 7**	**Issue # 8**	**Issue # 9**	**Issue # 10**
aiR	24	36	50	46	13
Reviewer	31	42	101	22	51
Overlap	19	23	47	7	9

At the issue level, post-conflict-review recall and precision estimates varied. Table 3 below contains estimates for three of the issues. Note that while recall and precision likely could have been improved for some of the issues through iteration, the results for other issues appear to reflect genAI's current limitations.

For example, Issue #1 called for information that was limited to a specific time period, namely the year 2023. The low precision and high recall rates for genAI's predictions for Issue #1 appear to be caused in large part by its inability to consider document metadata or other indicia of date. As a result, genAI predicted documents from other years were relevant, in addition to correctly identifying the relevant 2023 documents. In other words, it could not distinguish documents by year. By contrast, the documents responsive to Issue #5 (above) were largely spreadsheets and charts. GenAI's relatively low recall rate for Issue # 5 appears to be caused in part by a limited ability to identify the relevant information in this type of document format, for which relevance was more evident in native or image view than in extracted text.

At a high level, the issue analysis showed that where determinations were driven by the text of the document, genAI performed fairly well. Where the issue coding depended on metadata, structured content (e.g., spreadsheets or templates/forms), or documents without sufficient text (e.g., images), genAI's issue analysis performed much worse. It is possible that further iteration of the prompts for some issues could have improved the overall results. Another possible solution would be to apply metadata filters prior to running the documents through genAI where a metadata field is particularly relevant to the determination—such as where dates are critical to the determination.

Table 3. Recall and Precision Estimates for Example Issues

Issue Tag	Issue # 1	Issue # 3	Issue # 8
Predicted responsive by aiR	93	100	50
Tagged responsive by reviewer	27	90	101
True positives	17	64	50
aiR + reviewer agreement	15	52	47
Identified in conflict review	2	12	3
False positives	76	36	0
False negatives	0	8	11
Recall	**100%**	**64%**	**75%**
Precision	**18.2%**	**88%**	**100%**

C. *Analysis of Workpapers*

The workpapers collected in the underlying legal matter consisted largely of a specific type of document that is highly relevant in the client's industry. These documents included a number of characteristics that proved challenging when genAI was applied. After applying the issue-tagging prompts to a sample of the workpaper documents, the study team determined that genAI was not an effective tool for issue coding this particular type of industry-specific document at the time of the study. Accordingly, the analysis was not run on the full 200-document sample.

This assessment was helpful in highlighting several factors that may present limitations when using genAI tools to evaluate certain kinds of workpaper documents. First, because current genAI tools used for eDiscovery operate on the extracted text of a single document, genAI cannot consider family relationships or the context provided by family documents. In this study, this limitation was problematic because of the way data was stored and exported from the client's internal workpapers system. Specifically, during export, a cover page was created and assigned as the parent to the underlying document, which was assigned a child relationship. Because the cover page was processed as a separate document, the genAI tool did not have the benefit of context from the coversheets that were available to human reviewers when evaluating the document.

Similarly, most genAI tools currently cannot consider hyperlinked documents embedded within the underlying document. In this study, many of the workpapers contained hyperlinks to various support

documents, which provided additional context to the human reviewers that genAI could not consider at the time of the study.

Finally, several of the workpapers featured formatting issues that posed challenges for genAI. For example, many documents were templates—both completed and blank versions—which the tool processed as text files without recognizing the template formatting. As a result, genAI seemed to have difficulty distinguishing between actual content and template elements. Additionally, workpapers generally do not convert cleanly to extracted text, making it difficult for the tool to interpret their content accurately. Checklists were especially problematic, as checked boxes are not represented in the extracted text, further complicating genAI's ability to analyze the information accurately.

D. *Analysis of genAI's Cost Effectiveness*

Based on the study results showing the effectiveness of genAI in conducting responsiveness review, integrating genAI into a document review workflow in large-scale matters may represent significant cost savings in appropriate cases—particularly compared to manual review.

The cost savings would, of course, vary based on a number of document review-specific factors. For example, the percentage of the document population that is eligible for genAI will impact the calculus, with higher savings for document sets where the eligibility percentage is higher. Further, companies would see greater cost savings where there are fewer or no restrictions on the type of documents where genAI replaces first-level review—such as sensitive documents or C-suite custodians' documents. Cost savings would also be higher where second-level quality control by human reviewers is limited to a sample rather than the entire document set. Furthermore, the complexity of the document requests and the time spent on prompt iteration will also be a factor in the calculus on cost savings.

In addition, clients and their counsel should consider the potential for costly negotiations related to ESI protocols that call for the use of genAI, particularly in comparison to the well-established acceptance of TAR. Although integration of genAI may offer greater efficiency in conducting document reviews, its status as a novel technology not yet widely accepted by courts—and the potential for disputes with opposing counsel over matters such as transparency and validation—could lead to contentious and costly negotiations, potentially making it more expensive than traditional approaches at this point in time.

E. *Potential Use Cases and Future Workflow Considerations*

This study provides important empirical validation for the accuracy of genAI tools for inclusion in document review workflows. In particular, genAI seems well suited to conduct first-level responsiveness review of

text-based documents, subject to certain limitations. Although genAI tools may serve as a defensible and efficient component of the eDiscovery workflow in appropriate cases, human oversight and quality-control checks remain important to ensure accuracy and reliability. GenAI may also be effective as a means of prioritizing documents for attorney review or as a quality-control measure applied to certain document sets following attorney review. On the quality control front, the study results suggest that the tool may be particularly effective at identifying false positives—documents coded by attorneys as responsive that are not, in fact, responsive—thereby reducing the inadvertent production of nonresponsive documents. This study also provides empirical validation for the use of genAI tools for inclusion in document review workflows.

The study also offers valuable insight into the factors to consider when determining if genAI would be suitable for a particular legal matter and/or use case, as well as developing a proactive workflow that ensures the efficient use of the tool. A few practical considerations include:

- Extracted text: Consider the data set's characteristics to determine if genAI would add efficiency, keeping in mind that genAI tools work best on primarily text-based documents. Exclude from the genAI workstream any documents that do not meet the extracted text size specifications and other technical requirements.
- Family relationships: Factor in the inability of genAI tools to recognize family relationships and how that may impact its efficacy for evaluating certain documents. Consider whether workarounds may be feasible and economical. For example, to analyze the industry-specific workpaper documents in this study, the team could have created a version of extracted text that merged the cover page and document text for purposes of the genAI analysis, allowing genAI—like the human reviewers—to consider the context provided by the cover page when assessing responsiveness.
- Information beyond the four corners of a document: In some cases, information outside the four corners of the extracted text, such as metadata (e.g., date, document title, author, custodian, or file path) should be considered to make an accurate call for a category of documents. Consider excluding these documents from the genAI workstream or addressing them with a tailored workflow (e.g., modify the prompt prior to running that set of documents or run the documents through a metadata filter prior to analyzing with genAI).
- Accuracy of extracted text: The efficacy of genAI tools is dependent on the accuracy of the extracted text. In this study, for instance, documents containing checklists presented a particular challenge because the genAI tools used for this study cannot determine whether a box was checked. If the data set includes document categories that may have

incomplete or inaccurate extracted text, consider reviewing a sample of the extracted text to determine whether those documents should be excluded from the genAI workflow. Document categories susceptible to this limitation may include scanned documents and images, graphs, or charts.

The genAI revolution in legal practice is still in its early stages, and additional empirical research will be important to test, validate, and refine the results reported here across a wider range of matters, document types, workflows, and risk profiles. In particular, future studies should examine how tool performance varies with different prompts and review protocols, how well results generalize beyond a single closed matter, and what quality-control measures best support defensible use under prevailing discovery standards. Even with those open questions, the proof-of-concept findings in this Article provide a meaningful starting point: they suggest that, when appropriately deployed and monitored, genAI can exceed what is possible with manual review and search terms. In some respects, genAI is comparable to or exceeds the efficacy of TAR. In short, genAI may enhance the speed and consistency of large-scale review while preserving (and in some use cases improving) accuracy.

Accordingly, these results do not mark the end of the inquiry; rather, they mark the beginning of a conversation about where genAI belongs in the modern review workflow, what safeguards it requires (if any), and how lawyers can integrate these tools in ways that are aligned with core discovery obligations. Ultimately, genAI's value in eDiscovery will turn less on novelty and more on proof of its ability to be implemented in ways that are proportional, reliable, and fair. The challenge—and the opportunity—is to move beyond hype or hesitation and instead deploy genAI grounded in evidence and defensibility.