AI IN LITIGATION PRACTICE

Colin S. Levy

Contents

1. Introduction
2. The Discovery Reckoning
3. Drafting Briefs and Motions Without Falling Into Mata
4. Deposition Preparation and Witness Work
5. Trial Preparation and Presentation
6. Damages, Settlement, and ADR
7. Authenticating and Challenging AI-Generated Evidence
8. The Compliance Layer: Standing Orders and Verification
9. Privilege and Confidentiality in Litigation Practice
10. Vendors, Tools, and the Litigation Support Ecosystem
11. Ethics, Supervision, and the Duty of Technological Competence
12. Building a Litigation AI Stack That Survives Scrutiny
13. What Comes Next

1. Introduction

A litigator's working life has changed in ways that are easy to underestimate. Five years ago, the workflow of a commercial dispute looked roughly the same as it had for two decades. Documents were loaded into a review platform and coded by issue. Briefs were drafted in Word. Deposition outlines were built from a notepad. Exhibits were prepared in binders. Today, an attorney handling that same matter may run privilege screens through a generative AI layer, draft an opening section of a summary judgment brief by feeding case law into Westlaw CoCounsel, summarize a thousand-page deposition with Steno's Transcript Genius, and ask a language model to surface inconsistencies between two witnesses' prior testimony.

The pace of that shift has outrun the pace of guidance. Bar associations have issued opinions, judges have issued orders, and the sanctions docket has accumulated cases the names of which now circulate as cautionary tales: *Mata v. Avianca*, *Park v. Kim*, and a growing number of others tracked by the RAILS AI Use in Courts project and similar databases. What has not yet emerged is a practical, integrated treatment of how AI changes the actual work of litigation, from the first preservation letter to the post-trial motion. This guide is an attempt to fill that gap.[1]

The companion volume to this one, *AI in the Courtroom*, addressed the bench. It examined how judges and court systems are confronting AI, what policies have emerged, and what the judiciary's engagement reveals about institutional adaptation. This guide turns the chair around. The questions here are practitioner questions. How should a litigator approach predictive coding in 2026 when generative AI has begun to displace older TAR workflows? What does Rule 11 require when a brief was drafted with AI assistance? When opposing counsel introduces a video, what does the deepfake objection look like, and how do you raise or defeat it? How do supervisory obligations under Model Rule 5.1 attach when an associate uses ChatGPT on a privileged document?

The guide is written for litigators and the people who support them. Senior partners managing complex matters. Associates building briefs and outlines. Paralegals running document review. Litigation support professionals selecting tools. It does not assume technical expertise, but it does assume familiarity with the basic mechanics of civil and criminal practice. Readers without that background may want to start with *AI for Lawyers* or *AI for Legal Teams*, also in this series.

A final note about scope. This guide covers civil and criminal litigation in U.S. courts, with attention to the federal-state distinctions that affect AI policy. International arbitration, regulatory enforcement, and pre-litigation investigation each carry their own considerations and are addressed only where they intersect with traditional litigation practice. Where relevant, the guide cross-references *AI in the Courtroom* for material covered there in greater depth.

2. The Discovery Reckoning

For most commercial litigators, document review is where AI first arrived and where it now lives most consistently. The history matters because it shapes how courts treat the next wave.

Predictive coding entered the litigation mainstream after Magistrate Judge Andrew Peck's opinion in *Da Silva Moore v. Publicis Groupe* in 2012, which approved the use of technology-assisted review under the Federal Rules. Judge Peck's later opinion in *Rio Tinto v. Vale*, observing that predictive coding had become "black letter law," gave practitioners cover to deploy continuous active learning workflows without negotiating their use ex ante in every case.[2] By 2020, TAR was routine on matters of any size, and the question for a producing party was no longer whether to use predictive coding but how to defend the protocol if challenged.

Generative AI has unsettled that settled state. The new generation of review tools relies less on classifier-based ranking and more on language models that can read documents, summarize them, identify privileged content, and code by issue based on natural-language prompts. Vendors like Reveal, Relativity (with its aiR product line), Everlaw, and Lighthouse have integrated generative capabilities that operate alongside or in place of older predictive coding.[3] The question for litigators is no longer whether AI improves review at the margins but whether the workflows you used last year still satisfy current proportionality and defensibility standards.

Defensibility and Rule 26(g)

Federal Rule of Civil Procedure 26(g) requires that every discovery response be signed by counsel, and that signature certifies the response is "complete and correct as of the time it is made" after a "reasonable inquiry." The sanction provisions of Rule 26(g)(3) are mandatory: if a certification violates the rule without substantial justification, the court "must impose an appropriate sanction." When a generative AI tool is making coding decisions on potentially privileged or responsive material, that signature is doing work it has not historically done.

Three practical questions follow. First, what testing has been conducted on the model's accuracy for the specific data set in this case? Vendors will provide marketing claims about general performance, but defensibility under Rule 26(g) requires case-specific validation. Second, what is the human review layer? A pure AI workflow with no human sampling is unlikely to survive scrutiny if the production turns out to be incomplete. Third, what is the documented error rate, and how was it measured?

Privilege Review and the Recall Problem

Privilege review is where generative AI has produced both the most interesting results and the most concerning ones. The performance variance across vendors is wide. Some tools, in the published demonstrations and in customer pilots, perform comparably to or better than experienced contract-attorney review at the same cost. Others, on the same corpus, miss material amounts of privileged content. The variance has been documented anecdotally by litigation support professionals and at industry conferences, but rigorous independent benchmarking remains thin on the ground.[4]

The variance matters because privilege errors are asymmetric. Producing a privileged document is a serious problem; over-withholding is recoverable through clawback or motion practice. A producing party that relies on a tool with poor recall has produced, at scale, material that should have been withheld. The Rule 502(d) order in the case may protect against subject-matter waiver, but it does not undo the disclosure of damaging communications to opposing counsel.

The takeaway for litigators is to treat AI privilege review as a layer rather than a substitute. Run the AI screen, then re-review the production documents the tool flagged as a close call, then sample the documents the tool cleared as non-privileged. The cost is meaningful but it is the cost of defensibility.

Sanctions Risk Under Rule 37(e)

Rule 37(e) governs sanctions for failure to preserve electronically stored information. Generative AI introduces a new wrinkle. When an attorney uses an AI tool in the ordinary course of producing or modifying documents that may later be discoverable, the resulting outputs may themselves be ESI subject to preservation. This is most pointed in two contexts. First, when a client uses AI tools to draft documents in the regular course of business, those drafts and their underlying prompts may be discoverable in subsequent litigation. Second, when counsel uses AI to prepare litigation materials, the question of whether the prompts and outputs are work product turns on the analysis covered in Chapter 9.

Producing parties should address AI use in their ESI protocols. Model ESI protocols and the working commentary from organizations like the Sedona Conference and EDRM have begun addressing AI tools, and the negotiation of protocol language around AI disclosure, prompt preservation, and training data is now a routine feature of complex matters.[5]

Cross-Border Discovery and the Data Localization Problem

For multinational matters, generative AI has complicated cross-border discovery in specific ways. Most enterprise generative AI tools route data through cloud infrastructure that may not satisfy GDPR's requirements for international transfer, and the EU AI Act's high-risk classification of AI used

in litigation raises additional compliance questions. Producing parties handling EU data should confirm with the AI vendor where data is processed, what residency commitments are in place, and whether the vendor maintains conformity assessments for high-risk use cases.

Chinese data is harder still. The Personal Information Protection Law and the Data Security Law each impose restrictions on the export of certain categories of data, and the use of foreign AI tools to process Chinese-origin documents has been the subject of enforcement actions. Litigators handling Chinese discovery should plan for on-premise or domestically hosted AI solutions, or accept the slower pace of traditional review.[6]

3. Drafting Briefs and Motions Without Falling Into Mata

The hallucination cases have not slowed down. The public sanctions docket has continued to grow, and the trackers maintained by RAILS, the LK Law database referenced in the prior guide, and the Bolch Judicial Institute now collectively document well over a hundred reported instances of AI-generated false citations or quotations in court filings.[7] A naive reading of these cases is that the lawyers involved were lazy or technologically incompetent. A closer look shows that many of the sanctioned attorneys were experienced, that they did some verification, and that the verification they did was inadequate to the failure mode.

Why Spot-Checking Fails

Generative AI hallucinations are unevenly distributed in a way that defeats casual verification. A typical AI-drafted brief contains a mix of real, well-cited authorities and fabricated ones. The fabricated cases are often plausible. The case names track actual parties, the citations follow standard reporter format, the holdings sound consistent with the area of law. An attorney who spot-checks three of the ten citations and finds them real may reasonably conclude that the brief is reliable. The fabricated citations are in the seven citations that were not checked.

Schwartz's deeper failure in *Mata* was not that he used ChatGPT. It was that he asked ChatGPT to verify ChatGPT, and accepted the answer. The model confirmed its own fabrications because confirmation was statistically the more likely response to the prompt. Models in 2026 are better calibrated, but the failure mode persists.

The Verification Stack

A defensible verification workflow for AI-assisted brief writing has three layers. First, every citation must be confirmed in a primary source: Westlaw, Lexis, the court's docket, or the published reporter. The AI tool is not a primary source, no matter what features the vendor markets. Second, every quotation must be verified against the underlying opinion, not merely against a search result that contains the quoted text. AI tools have been documented to fabricate quotations from real cases, which means that finding the case is not enough. Third, every proposition must be checked against the authority cited for it. Real cases can be cited for propositions they do not support, and an AI-drafted brief is particularly prone to this because the model is generating language that "fits" rather than language that accurately summarizes a holding.

The most efficient verification stack uses the citation-verification features now built into Westlaw CoCounsel and Lexis+ AI Protege, supplemented by manual review of the underlying authority for any proposition that drives the argument. The "Quick Check Judicial" feature in CoCounsel and "Brief Analysis" in Lexis Protege are not substitutes for human review, but they catch the easy cases (citations that do not exist, quotations that do not appear in the cited authority) and free attorney time for the harder ones (citations that exist but do not support the proposition).[8]

Drafting Workflows

Different phases of brief writing have different AI risk profiles. Initial outline generation is relatively safe. An AI tool that produces a structural outline of an argument is not generating citations and is therefore less prone to producing fabricated authority. First-draft prose for descriptive sections (statements of fact, procedural history) is also relatively low-risk, because the underlying material is the record, not the law.

The high-risk phase is the legal analysis. When an AI tool drafts a discussion of the applicable standard, it is reaching for authority, and that is where hallucination occurs. Two workflows reduce the risk. The first is what some firms call "argument-then-cite": the lawyer writes the argument in plain prose, identifies the propositions that need authority, and uses the AI tool only to suggest cases that support those propositions. The cases are then verified before being included. The second is "draft-then-strip": the AI generates a draft, the attorney removes every citation, and citations are added back manually with verified authority. Both workflows treat the AI as a structuring aid rather than a research tool.

The Sanctions Docket

The patterns visible in the sanctions cases are instructive. Solo and small-firm practitioners are over-represented relative to their share of federal civil filings, which suggests that lack of institutional verification infrastructure is a contributing factor. Pro se litigants account for a large minority of the cases, and several courts have responded by issuing standing orders that apply specifically to unrepresented parties. Sanctions amounts have ranged from a few hundred dollars in some district court matters to the $5,000 imposed in *Mata*, with referral to disciplinary authorities in cases like *Park v. Kim*.[9]

The sanctions docket extends beyond fabricated citations. Courts have begun sanctioning attorneys for AI-generated misstatements of holdings, fabricated procedural history, and reliance on AI summaries that omitted dispositive case facts. The verification duty extends to the substance of the cited authority, not merely to its existence.

4. Deposition Preparation and Witness Work

The deposition is one of the more underexplored frontiers for AI-assisted litigation work. For decades, deposition preparation followed a familiar arc. An associate read the relevant documents, drafted an outline, the partner refined it, and the deposition proceeded. The transcript came back days or weeks later, and the team produced a summary by hand. Generative AI compresses each step.

Outline Generation

A first-draft deposition outline can be produced by an AI tool given a corpus of relevant documents and a list of topics. The output is rarely usable as-is, but it is often usable as a structural starting point. An associate working from an AI-generated outline can spend her time refining the questions, sequencing them strategically, and adding the document references rather than starting from a blank page.

The risk in this workflow is the same as in brief drafting. The AI may suggest factual premises that are not supported by the underlying record. An outline that asks "is it correct that you sent an email on June 15 confirming the agreement?" is only useful if such an email exists. Practitioners should treat AI-generated questions as proposals to be verified against the documents, not as established facts.

Real-Time Transcript Analysis

A more substantial change is happening in the deposition itself. Tools like Steno's Transcript Genius, Trustpoint.One's Trial Director with AI transcript integration, and several smaller vendors offer real-time transcript analysis that can flag inconsistencies between current testimony and prior statements, surface relevant documents, and highlight responsive answers that contradict allegations in the pleadings. Some of these tools run on a laptop at counsel table; others are designed for second-chair use.[10]

The practical value depends on how well the tool is loaded with the relevant materials. A real-time transcript tool that has access to the deponent's prior depositions, sworn declarations, and key emails can produce useful flags. A tool that has access only to the current transcript adds little. The investment in setting up the tool before the deposition is usually a few hours of paralegal time, which is recovered many times over during the deposition itself.

Post-Deposition Summaries

After the deposition, AI tools can produce topic-organized summaries, exhibit-by-exhibit indexes, and impeachment-ready charts of inconsistencies. The Indiana Supreme Court pilot referenced in the prior guide, which compressed mental health appeal timelines from 143 days to 31 days, is the most visible example of what AI transcription and summarization can do at scale. For private litigators, the analog is summary turnaround. What used to be a multi-day project for an associate can now be a one-day project with substantial AI assistance.

The verification problem reasserts itself here. AI summaries of deposition testimony have been documented to misattribute statements, conflate witnesses, and characterize testimony in ways that subtly favor the side that prompted the summary. A summary used for internal strategy can tolerate this margin of error. A summary used in a brief or a settlement memorandum cannot.

Witness Preparation

The use of AI in witness preparation is more cautious. Some practitioners have experimented with using language models to simulate cross-examination, generating questions from an opposing party's perspective. The resulting practice sessions can be useful, particularly for inexperienced witnesses. The model is reasonably good at generating plausible cross questions, even if it cannot match the rhythm or instinct of an experienced opposing counsel.

The line that should not be crossed is preparation that becomes coaching. Feeding an AI tool the witness's prior statements and asking it to "help the witness explain inconsistencies" is mechanically similar to the kind of testimony preparation that has long drawn judicial criticism under Model Rule 3.4 and corresponding state rules.[11] The fact that the coordination was conducted with an AI tool does not change the ethical analysis.

5. Trial Preparation and Presentation

Trial preparation has historically been the phase of litigation least affected by AI. Briefs are written, documents are reviewed, witnesses are prepared, and then the team retreats into a war room for two weeks of binder assembly and demonstrative editing. That pattern is changing.

Demonstratives and Visual Aids

Generative AI image and video tools have made it possible to produce demonstrative exhibits more rapidly than the in-house graphics teams of the past. A timeline that took a graphics designer a week to produce can now be drafted in an hour and refined to final quality in a day. Animations of complex events (an accident reconstruction, a sequence of trades, a chain of corporate transactions) can be generated from text descriptions and refined iteratively.

The evidentiary risk is real. Demonstratives are not evidence in most courts; they are admitted under FRE 611(a) as aids to understanding testimony. But the line between a demonstrative aid and substantive evidence has always been thinner than the rule suggests, and AI-generated visuals make it thinner still. A plaintiff's animation of an accident is more persuasive when it looks photorealistic, and the sponsoring witness must have a sufficient basis to authenticate it under FRE 901. If the animation was generated from a written description by a model that filled in details the witness did not provide, the foundation is weaker than it appears.

The practical guidance is to disclose the AI involvement in the production of demonstratives early, document the source material on which the AI worked, and prepare a sponsoring witness who can testify credibly that the demonstrative is a fair and accurate representation of the underlying facts. Opposing counsel will object, and a court that has been told about the AI involvement upfront is more likely to allow the demonstrative than a court that learns about it from a Daubert motion.

Jury Selection

Jury consultants have used statistical models for decades. Generative AI changes the calculus in two ways. First, the speed of analysis is faster. A model can summarize the social media presence of a venire in minutes rather than the hours required by a traditional consultant team. Second, the depth of analysis is greater. Language models can identify rhetorical patterns, expressed political views, and indicators of attitudes toward parties, lawyers, or industries that traditional consultant intake would miss.

The ethical and legal limits are significant. Some jurisdictions have begun to address whether AI-driven juror profiling raises issues under state privacy laws or under judicial codes governing

investigation of jurors. The general framework that has emerged from prior bar opinions on social media research, such as ABA Formal Opinion 466 (2014), maps reasonably well onto AI-assisted analysis.[12] Publicly available information remains fair game, the use of AI to compile and analyze that information is generally permissible, and any direct contact with prospective jurors (including through social media platforms) crosses the line into impermissible communication.

Real-Time Research at Counsel Table

Trial judges have not, as a class, welcomed laptops at counsel table running AI tools that respond in real time to testimony. The combination of judicial concern about distractions in court and the possibility that the AI might generate hallucinations during live proceedings has produced a cautious posture. Some courts have issued orders restricting real-time AI use at counsel table without prior approval, and trial-specific disclosure requirements for AI tools have begun appearing in local rules and standing orders. Practitioners should check the current rules of any court before relying on AI tools at trial.

Within those constraints, AI-assisted research during recess is increasingly common. A break in the trial day is now an opportunity to refresh research on a question that arose unexpectedly, draft language for a motion in limine, or prepare a witness for the next morning's testimony. The work is still done by lawyers, but the research that supports it is faster.

6. Damages, Settlement, and ADR

Damages and settlement work is one of the better fits for AI tools, because the analytical questions are quantitative and the risk of hallucination is lower. A model that draws on a database of jury verdicts is operating on structured data, not generating prose, and the failure modes are different.

Damages Modeling

Tools like Bloomberg Law's litigation analytics and Lex Machina's judge analytics have used machine learning for years. The newer generation of damages tools incorporates generative AI to translate model outputs into narrative damages analyses suitable for client presentations or settlement memoranda. A plaintiff's lawyer can produce a damages spreadsheet, run it through a tool that compares it against verdicts and settlements in similar cases, and receive a probability-weighted range that incorporates the relevant variables.

The output is only as good as the underlying data, and the underlying data is partial. Public verdict and settlement databases miss confidential settlements, omit contributory factors, and over-represent reported decisions. A tool that estimates likely recovery from these data sources is making predictions on a biased sample. Practitioners should treat the outputs as one input into the negotiation rather than as a defensible quantitative answer.

Settlement Analytics

Settlement analytics tools attempt to predict the probability of settlement at various points in a case, the likely settlement range, and the strategic considerations that affect each. Several vendors now market modules of this kind, often integrated with case management platforms. Performance claims tend to be aggressive and are rarely backed by independent validation, which means a litigator should approach the predictive outputs with the same skepticism appropriate to any model trained on incomplete data.[13]

A more conservative use of these tools is descriptive rather than predictive: using them to compile a comprehensive view of how similar cases have resolved, with the lawyer drawing strategic inferences from the data. This is closer to how Lex Machina has historically been used and is less subject to the failure modes that affect prediction.

ADR and Mediation

For mediation, AI tools are useful in two specific ways. The first is preparation. An AI tool can generate a comprehensive memorandum that lays out the legal posture, the damages range, the pro-

cedural status, and the likely outcomes if the case does not settle. The second is real-time analysis. During a mediation, an AI tool can rapidly assess a new offer against the case's value range and flag the strategic implications.

The confidentiality concerns are pointed. Mediation communications are typically protected by statute or rule, and uploading mediation materials to a consumer AI platform exposes those communications to the platform's data practices. The *Heppner* analysis from the prior guide applies here. The protection that mediation confidentiality provides depends on the materials remaining confidential, and a consumer AI tool may not satisfy that condition.

7. Authenticating and Challenging AI-Generated Evidence

The deepfake question, deferred in the prior guide for a fuller treatment here, is now the most active area of evidentiary AI law. Two cases have reached substantial appellate analysis, and several more are pending.

The State of the Doctrine

Federal Rule of Evidence 901(a) requires that a proponent produce "evidence sufficient to support a finding that the item is what the proponent claims it is." For decades, this standard was easy to satisfy for digital evidence: a witness with knowledge testified that a video was an accurate recording, the chain of custody was traced, and the evidence was admitted. The proliferation of generative AI image, audio, and video tools has changed the calculus.

The Washington courts addressed a deepfake-style objection in *State v. Puloka* in 2024, in a context that became a reference point for how trial courts handle AI-authenticity challenges. The court declined to require the proponent of video evidence to "disprove" AI generation as a prerequisite to admission, treating traditional authentication methods (witness testimony, chain of custody, metadata consistency) as sufficient where they were available.[14] Practitioners should treat the precise contours of the holding as fact-bound and verify the current state of the case before relying on it in briefing.

The Tesla Autopilot litigation has surfaced related questions about AI-altered or AI-enhanced video evidence. The pretrial record in those cases reflects a general expectation that parties offering video evidence will be prepared to address whether the recording has been altered, enhanced, or otherwise processed.[15] The trajectory of the case law suggests that foundational expectations for video evidence are tightening, even where no formal rule amendment has occurred.

The Practitioner's Foundation

A litigator introducing video, image, or audio evidence in 2026 should plan for a deepfake objection regardless of whether one is anticipated. The foundational testimony should address the source of the recording, the chain of custody, any post-recording processing or enhancement, the metadata characteristics that are consistent with an unaltered original, and (where available) the verification methods used to confirm authenticity. The witness need not be a forensic expert in most cases, but a forensic expert may be necessary if the opposing party has produced affirmative evidence of manipulation.

The Practitioner's Challenge

Challenging an opposing party's AI-generated evidence is harder than authenticating one's own. The producing party usually has access to the source material; the challenging party usually does not. Three strategies have emerged.

The first is metadata analysis. Most AI-generated content has subtle metadata characteristics that differ from camera-recorded content. A forensic expert can identify these characteristics, though the analysis is increasingly difficult as generative tools improve.

The second is content analysis. AI-generated images and videos still produce occasional artifacts: inconsistent shadows, unnatural physics, errors in fine detail (text in the background, fingers, reflections). A trained expert can identify these, but the cost is significant and the analysis may not be conclusive.

The third is source-document discovery. If the opposing party generated the evidence with a specific tool, the prompts and intermediate outputs may be discoverable. The Sedona Conference 2025 ESI protocol update includes specific provisions for the discovery of AI prompts and outputs as ESI, which gives challenging parties a procedural lever.

Statutory Responses

Louisiana's Act 250, referenced in the prior guide, requires attorneys to exercise reasonable diligence to determine whether evidence has been AI-generated. Several other states are considering similar legislation, and California has been particularly active on the watermarking and provenance side, with measures that impose obligations on generative AI providers rather than on litigants directly.[16] The federal landscape remains undeveloped. Discussions about whether the Federal Rules of Evidence should be amended to address AI-generated content are ongoing in academic and rules-committee settings, but no amendment has been adopted as of the date of this guide.

8. The Compliance Layer: Standing Orders and Verification

The patchwork of standing orders, local rules, and state directives that the prior guide described creates a compliance burden that did not exist three years ago. Litigators should not rely on memory or general counsel guidance when filing in an unfamiliar court. The rules change, and the penalties for noncompliance are not limited to in-court embarrassment.

Tracking the Patchwork

Several services now track AI-related court orders. The RAILS AI Use in Courts Tracker, maintained by a consortium of legal organizations, is the most comprehensive public resource. The major legal research providers have begun integrating compliance feeds into their case management products. Westlaw's Practical Law and Lexis Practical Guidance both maintain updated trackers of state and federal AI orders. For larger firms, internal knowledge management systems that flag AI requirements for specific judges and districts have become standard.

The minimum due diligence before filing is to check the assigned judge's standing orders, the district's local rules, and any state-level directive that may apply. The check should be repeated at the start of every new matter, because the rules continue to change.

The Mata Doctrine in Practice

The verification duty articulated in *Mata v. Avianca* has been reinforced in a steady stream of subsequent cases. The Second Circuit's treatment of the issue in *Park v. Kim*, in which counsel cited nonexistent authority generated by ChatGPT, made clear that the duty to verify is not satisfied by reliance on the tool's output, regardless of the lawyer's subjective belief in its reliability.[17] The composite picture across the sanctions cases is that an attorney who uses AI to draft any portion of a filing has the same verification duty as if she had drafted it herself, and that the duty extends to factual assertions as well as to citations.

The ABA's Standing Committee on Ethics and Professional Responsibility issued Formal Opinion 512 in July 2024, addressing the use of generative AI in legal practice. The opinion emphasizes that competence (Rule 1.1), confidentiality (Rule 1.6), supervision (Rules 5.1 and 5.3), and candor (Rule 3.3) all attach to AI use.[18] The opinion is non-binding but has been cited approvingly in subsequent state ethics opinions and in court orders addressing AI sanctions.

Disclosure and Certification

The disclosure question divides the federal courts. Some judges require an affirmative certification that AI was either not used or was used and verified. Others require disclosure only when AI was used. Still others require no disclosure but expect that any AI use will be subject to the same accuracy obligations as traditional drafting. The practical posture for a litigator is to assume disclosure is required and to disclose. The downside of unnecessary disclosure is small; the downside of a sanctions hearing for failure to disclose is large.

Several firms have adopted internal policies that require an AI-use log for every brief, identifying which sections were AI-assisted and what verification was performed. The log is privileged work product but is available if a sanctions question arises. This is a sensible defensive posture.

9. Privilege and Confidentiality in Litigation Practice

The privilege analysis in *United States v. Heppner*, treated more abstractly in the prior guide, has practical implications for litigators that deserve closer attention.

The Consumer-Enterprise Distinction

Judge Rakoff's central holding was that consumer AI tools, used by clients independently of counsel, do not generate privileged communications. The reasoning extends to a broader principle: the privilege protections available for AI-mediated communications track the conditions under which the communications occur. Consumer tools without confidentiality commitments fail the test. Enterprise tools with contractual data isolation, audit trails, and explicit confidentiality provisions are better positioned, though no court has yet held them protected.

For litigators advising clients on AI use during ongoing or anticipated litigation, the practical advice is direct. Clients should be told that their use of consumer AI tools to discuss legal matters may produce documents that will be discoverable. If the client wants to use AI to think through legal questions, the use should be routed through counsel, on an enterprise platform that the firm has vetted for privilege purposes.

Work Product and the Counsel-Direction Test

The work product doctrine, codified in FRCP 26(b)(3), protects materials prepared in anticipation of litigation by or for a party or that party's representative. The *Heppner* court declined to extend work product to materials Heppner prepared independently before retaining counsel, but flagged that materials prepared at counsel's direction would be analyzed differently.

The litigator's task is to ensure that AI use occurs under counsel's direction in a way that satisfies the doctrinal test. Three elements help. First, an explicit engagement with the AI tool framed as work prepared in anticipation of litigation. Second, a record showing that counsel directed the AI use, in the same way that *Kovel* agents are directed by counsel. Third, a confidentiality framework around the AI tool that prevents disclosure to third parties.

Inadvertent Disclosure

The Federal Rules and most state analogs include rules governing inadvertent disclosure of privileged material. The classic FRCP 26(b)(5)(B) procedure allows a producing party to claw back inad-

vertently produced privileged material, subject to the receiving party's right to challenge. The privilege analysis under Federal Rule of Evidence 502(b) considers whether the disclosure was inadvertent, whether the holder took reasonable steps to prevent disclosure, and whether the holder promptly took reasonable steps to rectify it.

When an attorney uploads a privileged document to a consumer AI tool, the disclosure analysis becomes more complicated. The "reasonable steps to prevent disclosure" prong asks whether the producing party's privilege review process was reasonable. A process that includes uploading documents to a tool with broad data-use rights is unlikely to satisfy that standard, and the asymmetry between the small effort required to use an enterprise tool and the substantial risk of using a consumer one will weigh heavily on a court evaluating the reasonableness of the firm's procedures.[19] Litigators should expect that the first wave of post-*Heppner* case law will further sharpen the doctrinal exposure for firms that have not formalized their AI use policies.

10. Vendors, Tools, and the Litigation Support Ecosystem

The litigation technology vendor landscape changed substantially in 2024 and 2025. Several major acquisitions consolidated the market, and new entrants displaced legacy products in specific niches. Litigators selecting tools should understand both the competitive dynamics and the diligence questions that distinguish responsible vendors from problematic ones.

Vendor Diligence

The minimum diligence questions for any AI vendor are well-established by now. Where is the data hosted? What contractual commitments exist around data use, training, and retention? What audit logs are available? What certifications has the vendor obtained? SOC 2 Type II is the baseline. FedRAMP and HITRUST are common in regulated industries. What is the vendor's incident history?

The newer questions are about model behavior. What language models underlie the tool, and have they been fine-tuned for the specific use case? What testing has been conducted on accuracy, and how were the test sets constructed? What is the vendor's policy on model updates that may change behavior? Can the customer freeze the model version for the duration of a matter, or is the customer subject to whatever update the vendor pushes?

The Tool Categories

Litigation-specific AI tools fall into several categories. Document review platforms (Relativity aiR, Reveal, Everlaw, Lighthouse) handle the bulk processing of discovery materials. Legal research platforms (Westlaw CoCounsel, Lexis+ AI Protege, Harvey, vLex's Vincent AI) handle research and brief drafting. Transcript analysis tools (Steno's Transcript Genius, Veritext's AI suite) handle deposition and hearing materials. Trial presentation tools (Trial Director with AI integration, OnCue) handle exhibits and demonstratives. Analytics tools (Lex Machina, Bloomberg Law's litigation analytics, Trellis) handle case research and prediction.

The boundaries between categories are blurring. The major legal research providers are extending into document review, the document review providers are adding research features, and the analytics providers are adding generative drafting capabilities. For most firms, the question is no longer which tool to use for which task but how to integrate tools across the litigation workflow without creating confidentiality gaps or duplicative costs.

Validation

A vendor's accuracy claims should be tested against the firm's own data. Most vendors will provide a pilot program that allows a litigation team to evaluate the tool on a representative case. The validation should include both accuracy metrics (precision, recall, F1) and qualitative assessment by experienced reviewers. A tool that scores well on synthetic data but poorly on the firm's actual cases is not worth the contract price.

11. Ethics, Supervision, and the Duty of Technological Competence

Comment 8 to Model Rule 1.1 has been on the books since 2012, requiring lawyers to "keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology." For most of its existence, the comment was honored more in citation than in observance. The advent of generative AI has changed that.

The Competence Floor

The technology competence comment has now been adopted in some form by a substantial majority of state bars. The state ethics opinions issued since 2023 have begun to articulate what competence means in the AI context. The Florida Bar's January 2024 opinion held that competence requires lawyers to understand how their AI tools work, what their failure modes are, and how to verify their outputs. The California State Bar's November 2023 guidance took a similar position and added that lawyers who do not have the time or capacity to verify AI outputs should not use the tools.[20]

These opinions create a meaningful floor. A litigator using an AI tool without understanding how it can fail is not competent in the AI use, and competence is a continuing obligation, not a static one.

Supervision Under Rules 5.1 and 5.3

Rule 5.1 governs supervisory responsibility over subordinate lawyers. Rule 5.3 governs supervisory responsibility over non-lawyer assistants. Both rules require partners and managing lawyers to make "reasonable efforts" to ensure that the lawyers and staff they supervise comply with the Rules of Professional Conduct.

The application to AI is direct. A partner who allows an associate to use an AI tool without verification training has not made reasonable efforts. A managing lawyer who allows a paralegal to upload privileged documents to a consumer AI tool has not made reasonable efforts. The supervisory duty extends to ensuring that the firm's AI policies are understood, followed, and enforced.

Confidentiality Under Rule 1.6

Rule 1.6 prohibits a lawyer from revealing information relating to the representation of a client without informed consent or implied authorization. The application to AI tools depends on the data practices of the tool. Consumer tools that train on user inputs likely violate the rule when used for client work.

Enterprise tools with contractual data isolation may not violate the rule, though informed client consent remains the cleaner path.

Several state ethics opinions have begun requiring informed client consent for any AI use that involves the disclosure of client confidential information to a third-party platform. The trend is toward broader consent requirements, and prudent practitioners should assume that explicit client consent will be required for substantial AI use within a few years.

Fees Under Rule 1.5

The fee implications of AI use are now a serious topic in bar discussions, in client billing guidelines, and in the practice management literature. Rule 1.5 requires that fees be reasonable. When AI compresses what would have been ten hours of associate time into one hour, the reasonableness of billing the client for the larger figure becomes a question that thoughtful counsel should be prepared to answer. The state and local bar opinions on AI in legal practice have begun touching this terrain, generally counseling transparency with clients and disfavoring billing structures that obscure the role AI played in producing the work.[21]

The market is moving in this direction independently of ethics opinions. Sophisticated clients increasingly require AI-use disclosures in their billing guidelines and resist paying for time on tasks that have been substantially automated. Litigators should expect that fee arrangements will become a routine subject of negotiation in matters where AI is used substantially.

12. Building a Litigation AI Stack That Survives Scrutiny

A coherent AI stack for a litigation practice is not a matter of selecting the most advanced tool in each category. It is a matter of integration, documentation, and habit. The firms that have done this well share several characteristics.

Centralized Approval

The firms that have avoided sanctions and ethics complaints around AI use have done so by centralizing the approval process. A single committee, typically including a partner with technology experience, the firm's general counsel, and the firm's information security officer, approves which tools may be used and on what kinds of matters. The committee maintains a current list, monitors changes, and handles exceptions.

This is more cumbersome than individual lawyer choice, but it produces three benefits. The firm has a defensible record of having vetted its tools. The firm avoids the situation in which two lawyers in adjacent offices use different AI tools with different confidentiality profiles on the same matter. And the firm's negotiating leverage with vendors is greater when it speaks with a single voice.

Documented Workflows

The firms that survive scrutiny have written workflows for common AI tasks. The workflow for AI-assisted brief drafting specifies the verification stack, the documentation requirements, and the disclosure obligations. The workflow for document review specifies the validation requirements and the review layer. The workflow for client confidential matter discussions specifies which tools may be used and what client consent is required.

Written workflows are not a guarantee against errors, but they are a strong defense if errors occur. A sanctions hearing in which the firm produces a documented workflow that the responsible attorney followed is a different proceeding from one in which no workflow exists.

Training and Audit

Continuing legal education in AI use should be a routine requirement, not a one-time onboarding. The technology changes quickly enough that quarterly training is justified for the lawyers and staff who use it most heavily. Annual training is the minimum for the rest of the firm.

Audit complements training. The firms that take AI seriously sample AI-assisted work product to verify that the workflows are being followed and that the verification steps are producing the intended results. The audit can be conducted by the firm's general counsel office or by an outside vendor. What matters is that it occurs.

Client Communication

Clients increasingly expect their litigators to disclose AI use proactively. Some require explicit consent before the firm may use AI on their matters. Others require disclosure in the engagement letter and in periodic status reports. Litigators should treat client AI policies as part of the engagement and should refresh client communications regularly as the firm's tools evolve.

The Cost-Benefit Reality

AI investment in a litigation practice does not produce uniform returns. Some matters benefit substantially: complex document review, large-volume deposition work, voluminous brief writing. Other matters benefit marginally: small disputes, single-document cases, settlement-driven work. Firms that have honestly assessed where AI provides leverage have outperformed firms that have applied AI uniformly.

The honest assessment requires data. A firm that tracks AI-assisted hours alongside outcomes, fees, and client satisfaction has the foundation for ongoing optimization. A firm that assumes AI is uniformly valuable will likely overinvest in low-return applications and underinvest in high-return ones.

13. What Comes Next

The next two years of litigation AI development are unlikely to be quiet. Several developments are already on the horizon.

Agentic tools at counsel table. Judge Schlegel's prediction of agentic AI in chambers, referenced in the prior guide, has a parallel in the litigation context. Multi-step AI agents that can independently gather facts, draft motions, and prepare exhibits are being piloted in 2026 by several large firms. The trial-side analog is the AI agent that assists at counsel table, monitors testimony, prepares cross-examination questions in real time, and surfaces impeachment material as it becomes relevant. The promise is significant; the risks are substantial. Real-time AI use during trial proceedings raises the same evidentiary and ethical questions as offline use, but with less time for verification.

The Federal Rules response. The Federal Rules of Civil Procedure and Federal Rules of Evidence are both subjects of active discussion in the rulemaking community as commentators, judges, and practitioners consider whether amendments are needed to address AI-specific issues. Rules 26(g), 37(e), and FRE 901 are the most frequently mentioned candidates. Whether and when formal amendments will emerge through the standard rulemaking process is uncertain.[22] Litigators should track the proposals as they develop and participate in any comment processes that open. The shape of any future rules will affect daily practice for years to come.

Privilege doctrine develops. The *Heppner* opinion will be tested in subsequent cases, and the precise scope of the privilege analysis will be refined through litigation. Several pending matters in the Second, Ninth, and Eleventh Circuits raise related questions, and the first circuit-level guidance is expected within the next year. The questions to be resolved include whether enterprise AI tools, properly configured, can support privilege; what role contractual confidentiality plays in the analysis; whether the *Kovel* doctrine extends fully to AI agents directed by counsel; and whether attorney work product doctrine has the same scope when the work was AI-assisted.

The insurance market adapts. Professional liability insurers have begun treating AI-related risk as a distinct underwriting consideration. Some carriers have added AI-specific questions to their applications; others are reportedly developing exclusions or riders. Litigators selecting professional liability coverage should review the policy language carefully and confirm that the coverage matches the firm's actual practice, particularly as AI use continues to expand into substantive legal work.

The technology continues to improve. The most important development is the one that is hardest to forecast. Generative AI in 2026 is meaningfully different from generative AI in 2024, and the pace of improvement does not appear to be slowing. The verification problems that drove the

early sanctions cases are receding as model capabilities improve, but the new failure modes that come with greater capability are still being mapped.

The constant across these developments is the same as it has been since *Mata*: the lawyer is responsible for the work that bears her name. The tool that produced the text is not relevant to that responsibility. The litigator who internalizes that principle, who builds her practice around verification habits and clear workflows, who treats AI as a powerful but error-prone collaborator rather than as an oracle, will be in a stronger position than her competitors as the technology continues to evolve.

This guide is part of the AI and Law Series by Colin Levy. Other titles include AI for Lawyers, AI for Legal Teams, AI Agents Data Handling and Cybersecurity Guide, AI in the Courtroom, and Law School in the Age of AI.

Endnotes

1. The cases referenced are *Mata v. Avianca, Inc.*, 678 F. Supp. 3d 443 (S.D.N.Y. 2023), and *Park v. Kim*, 91 F.4th 610 (2d Cir. 2024) (per curiam). The cumulative sanctions docket is tracked in part by the RAILS AI Use in Courts Tracker, available at https://rails.legal/resources/resource-ai-orders/, and by parallel projects discussed in the prior guide, *AI in the Courtroom* (2026).

2. *Da Silva Moore v. Publicis Groupe*, 287 F.R.D. 182 (S.D.N.Y. 2012); *Rio Tinto PLC v. Vale S.A.*, 306 F.R.D. 125 (S.D.N.Y. 2015) (Peck, M.J.). For background on the empirical case for technology-assisted review, 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 (2011).

3. Vendor materials and product documentation for Relativity aiR, Reveal, Everlaw, and Lighthouse. Summaries and analysis of the generative AI shift in document review can be found through EDRM publications and the major legal technology trade press (LawNext, Legaltech News, Above the Law's tech coverage).

4. The author has drawn on industry conference materials, vendor demonstrations, and discussions with litigation support professionals. Independent, peer-reviewed benchmarking of generative AI privilege review tools remains limited as of the date of this guide. Practitioners considering tools should request case-specific validation results rather than relying on vendor marketing claims.

5. The Sedona Conference and EDRM both maintain working groups addressing the impact of generative AI on discovery practice. Practitioners should consult the most recent publications from those organizations, as well as model ESI orders from their jurisdiction's bench-bar committees, for current guidance.

6. For background on cross-border discovery considerations, see The Sedona Conference, "International Principles on Discovery, Disclosure & Data Protection" (current edition); for the EU AI Act, see Regulation (EU) 2024/1689, classifying AI in the administration of justice as high-risk and imposing related obligations.

7. The figure draws on the RAILS AI Use in Courts Tracker, the LK Law database referenced in *AI in the Courtroom*, and the discussion of the broader sanctions trend in the legal trade press.

8. Vendor product documentation for Westlaw CoCounsel and Lexis+ AI Protege. The Quick Check Judicial and Brief Analysis features are described in the providers' published materials and have been the subject of trade press review.

9. *Park v. Kim*, 91 F.4th 610 (2d Cir. 2024). For a survey of the developing sanctions docket, see the resource trackers cited in note 1 and the broader treatment in *AI in the Courtroom*, Chapter 2.

10. Vendor product documentation for Steno Transcript Genius and other transcript-analysis platforms. The pace of product development in this category is rapid, and capabilities described here may have changed by the time this guide is read.

11. For background on the limits of permissible witness preparation, see Restatement (Third) of the Law Governing Lawyers §116 cmt. b; ABA Model Rule 3.4 and accompanying commentary. The application to AI-assisted preparation is the author's own analysis.

12. The Florida Bar, Ethics Opinion 24-1 (Jan. 2024), "Use of Generative Artificial Intelligence in the Practice of Law"; Cal. State Bar Standing Committee on Professional Responsibility and Conduct, "Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law" (Nov. 2023). For background on the ethics of jury investigation generally, see ABA Formal Opinion 466 (2014).

13. Discussion of settlement analytics tools draws on vendor product documentation and trade press coverage. Independent validation studies have been limited; practitioners should treat predictive accuracy claims with appropriate skepticism.

14. *State v. Puloka*, Wash. Ct. App. (2024), addressing the proper authentication standard for video evidence in light of AI-generation concerns. Practitioners citing the case in briefing should pull the current version of the opinion and any subsequent procedural history.

15. The pretrial record in the consolidated Tesla Autopilot litigation includes treatment of foundational expectations for video evidence. The case management orders in those proceedings have been the subject of trade press coverage and academic commentary.

16. Cal. Assemb. Bill 3211 (2024), addressing watermarking and provenance obligations for generative AI providers; La. Act 250 (2024), addressing attorney diligence requirements. Other states are pursuing related measures, and the legislative landscape is changing rapidly.

17. *Park v. Kim*, 91 F.4th 610 (2d Cir. 2024). The verification duty has been reinforced in numerous district court orders since *Mata*; readers should consult the current sanctions trackers cited in note 1 for an updated picture.

18. ABA Standing Committee on Ethics and Professional Responsibility, Formal Opinion 512, "Generative Artificial Intelligence Tools" (July 29, 2024).

19. This paragraph is the author's own analysis of how Rule 502(b) is likely to apply to AI-mediated inadvertent disclosure. As of the date of this guide, the author is not aware of a published federal opinion squarely addressing the question; practitioners should verify before relying on the analysis in briefing.

20. The Florida Bar, Ethics Op. 24-1 (Jan. 2024); Cal. State Bar, "Practical Guidance" (Nov. 2023). Other state and local bars have issued related opinions; readers should consult the current opinions of their own bar.

21. State and local bar opinions on AI in legal practice have begun touching on billing practices, with the general trend favoring transparency with clients about AI use and disfavoring billing structures that obscure that use. See ABA Formal Opinion 512 (2024) for related discussion of fee implications.

22. Discussion of possible Federal Rules amendments in response to AI is ongoing in academic, judicial, and rules-committee settings. As of the date of this guide, no formal rules amendment specific to generative AI has been adopted, and the timing of any future amendment through the standard rulemaking process is uncertain.