

Leveraging Artificial Intelligence to Streamline Factual Investigation and Case Development

A Practical and Defensible Framework for Modern Litigation

I. Introduction: Why AI in Litigation Is No Longer a Novelty

Artificial intelligence is no longer a speculative or emerging concept in litigation. Courts across jurisdictions have now spent more than a decade addressing technology-assisted review, analytics-driven discovery workflows, and the proportional use of advanced tools to manage ever-expanding volumes of electronically stored information (“ESI”). What has changed in recent years is not whether technology may be used, but how it is deployed, explained, and defended.

Modern litigation increasingly presents a familiar problem: massive data sets, compressed timelines, escalating costs, and heightened judicial expectations regarding efficiency and cooperation. Against that backdrop, AI-enabled tools—ranging from traditional analytics to large language models—are being incorporated into litigation workflows. The central question for practitioners is no longer whether these tools are permissible, but whether their use is reasonable, transparent, and defensible under the Federal Rules of Civil Procedure and analogous state rules.

This shift mirrors broader changes in judicial management of civil litigation. Judges increasingly expect parties to articulate *how* discovery decisions were made, *why* particular methodologies were selected, and *what steps were taken to ensure reliability*. As discovery disputes have moved away from debates over theoretical completeness and toward proportionality-driven outcomes, courts have shown a growing willingness to approve innovative workflows—provided counsel can explain and defend them.

Importantly, AI adoption does not occur in a doctrinal vacuum. The Federal Rules of Civil Procedure are technology-neutral by design. Rule 1’s mandate of “the just, speedy, and inexpensive determination of every action” and Rule 26’s proportionality framework implicitly encourage the use of tools that advance those goals. AI therefore fits naturally within existing procedural doctrine rather than requiring new rules or judicial reinvention.

This manuscript examines AI in litigation not as a disruptive force, but as part of a continuum of discovery practices already accepted by courts. When properly implemented, AI tools shift attorneys away from time-intensive, low-value document review tasks and toward higher-order legal analysis, strategic decision-making, and early case assessment—without undermining professional judgment or ethical obligations.

II. From Analytics to GenAI: A Continuum Rather Than a Leap

Courts have long recognized that manual, linear document review is neither the gold standard nor a required one. Early judicial acceptance of analytics and technology-assisted review (“TAR”) rested on a practical reality: humans are inconsistent, expensive, and error-prone when reviewing large document sets, particularly under time pressure.

Discovery analytics generally fall into two broad categories. **Conceptual analytics** include clustering, concept searching, and visualization tools designed to identify thematic relationships among documents. **Structured analytics** include email threading, near-duplicate detection, language identification, and automated redactions. These tools do not decide legal issues; they assist counsel in organizing and prioritizing information.

Courts evaluating early analytics tools quickly recognized that these technologies did not replace legal judgment; instead, they improved consistency and efficiency at scale. Traditional keyword searching, while familiar, often suffered from over-inclusiveness, under-inclusiveness, and inconsistent application across reviewers. Analytics tools addressed these weaknesses by identifying patterns that humans frequently miss, particularly in large, unstructured datasets.

TAR workflows evolved from early “seed set” models (often referred to as TAR 1.0) to continuous active learning (“CAL” or TAR 2.0), in which models are trained dynamically as reviewers code documents. Courts evaluating these technologies focused on recall, precision, validation, and proportionality—not on technical novelty.¹

The evolution from static seed-set training to continuous active learning marked an important doctrinal moment. CAL systems allowed review strategies to adapt in real time as new information emerged, aligning discovery practice more closely with Rule 26’s proportionality mandate. Rather than front-loading costs to achieve theoretical completeness, parties could assess marginal utility as review progressed and make informed decisions about when continued review no longer justified its expense.

More recent generative AI tools, including large language models, should be understood as an extension of these principles rather than a departure from them. While large language models introduce new capabilities—such as semantic summarization, chronology building, and multi-issue reasoning—they do not alter the fundamental legal inquiry. Courts remain concerned with whether discovery efforts are reasonable in scope, transparent in execution, and proportionate to the needs of the case. GenAI does not displace these standards; it operates within them.

¹ *Recall* measures the **thoroughness** of the search. It asks: Of all the relevant documents that exist in the entire collection, what percentage did the AI actually find? *Precision* measures the **efficiency** of the search. It asks: Of all the documents the AI flagged as “relevant,” how many are actually relevant? *Validation* is the **quality control** process used to prove that the TAR model actually worked. *Proportionality* is what courts consider under Fed R. Civ. P. 26(b)(1) and state equivalents that balances the need for information against the cost of getting it.

III. Judicial Acceptance: What Courts Actually Care About

A. Reasonableness as the Governing Standard

Judicial acceptance of advanced analytics is well established. In *Da Silva Moore v. Publicis Groupe*, one of the earliest decisions addressing predictive coding, the court emphasized that computerized review methods can be “at least as accurate, if not more so, than manual review,” and that the Federal Rules “do not require perfection.”² This recognition has become foundational to modern discovery jurisprudence.

Across jurisdictions, courts have repeatedly emphasized that discovery obligations are governed by reasonableness rather than perfection. No review methodology—human or machine—can guarantee error-free results, particularly when millions of documents are involved. Courts therefore evaluate whether a party conducted a reasonable inquiry, not whether it captured every potentially responsive document.

Subsequent decisions have reinforced several consistent themes. Producing parties are generally entitled to choose their discovery methodology, provided it is reasonable. Validation and quality control matter, particularly where large volumes of data are involved. Transparency facilitates cooperation, but does not require disclosure of proprietary algorithms or internal tooling. Discovery disputes are evaluated holistically, focusing on process rather than outcomes.

For example, in *Global Aerospace Inc. v. Landow Aviation, LP*, the Virginia Circuit Court observed that analytics “is capable of locating upwards of seventy-five percent of the potentially relevant documents . . . at a fraction of the cost and in a fraction of the time of linear review.”³

And, in *Federal Housing Finance Agency*, the Southern District of New York permitted the Defendants to use predictive coding over plaintiffs objections.⁴ The Court observed: “the literature that the Court reviewed . . . indicate[s] that predictive coding has a better track record in the production of responsive documents than human review.” Further, the Court explained “[N]o one could or should expect perfection from the discovery process. All that can be legitimately expected is a good faith . . . commitment to produce . . . responsive documents.”

The Western District of Pennsylvania provided a detailed decision describing how courts evaluate TAR workflows. The Special Master in the *In re Diisocyanates Antitrust Litigation* emphasized that “transparency transcends cooperation” and requires disclosure sufficient to make meet-and-confer discussions meaningful.⁵ At the same time, the court rejected rigid or formulaic stopping rules and reaffirmed that the legal standard remains reasonableness—not exhaustive completeness.

² *DaSilva Moore v. Publicis Groupe*, 1:2011cv01279 (S.D.N.Y. April 25, 2012).

³ *Global Aerospace Inc. v. Landow Aviation, L.P.*, 2012 WL1431215, No. CL 61040 (Va. Cir. Ct. Apr. 23, 2012).

⁴ *Federal Housing Finance Agency v. HSBC North America Holdings, Inc.*, 2014 WL584300 (S.D.N.Y. Feb. 14, 2014)

⁵ *In re Diisocyanates Antitrust Litig.*, No. MC 18-1001, 2021 WL 4295729, at *7 (W.D. Pa. Aug. 23, 2021), report and recommendation adopted, No. MC 18-1001, 2021 WL 4295719 (W.D. Pa. Sept. 21, 2021).

Significantly, the *Diisocyanates* court examined not only the existence of a TAR protocol, but also how that protocol accounted for validation, recall estimation, and iterative testing. The court’s analysis reflects an understanding that discovery methodologies must be evaluated in context, with attention to the nature of the case, the volume of data, and the significance of remaining unreviewed material.

For example, a review process that continues to surface duplicative invoices or low-level administrative communications may reasonably conclude earlier than one that continues to identify novel or highly probative documents. Courts have consistently recognized that such qualitative assessments are a legitimate component of proportionality analysis.

Similarly, in *Huntsman v. Southwest Airlines Co.*, the court approved a hybrid approach using keyword searches in tandem with technology-assisted review, holding that such a workflow satisfied the obligation to conduct a reasonable inquiry under Rule 26.⁶ The decision underscores that proportionality, phased discovery, and defensible methodology are often more important than any particular technical choice.

IV. Agentic AI as Workflow Orchestration, Not Autonomous Judgment

Recent discussion of “agentic AI” has generated concern that AI systems are making independent decisions or replacing attorney judgment. In litigation practice, however, agentic AI is best understood as workflow orchestration rather than autonomy.

Much of the concern surrounding agentic AI arises from terminology rather than functionality. In litigation contexts, “agents” do not act independently or exercise discretion. Instead, they perform narrowly defined tasks within parameters established, and quality-checked, by counsel. Each task—whether identifying documents tied to a witness, assembling chronologies, or flagging inconsistencies—is initiated, reviewed, and validated by human attorneys.

Unlike traditional TAR workflows, agentic AI tools do not require seed sets or wholesale retraining when new issues emerge. Counsel may instead investigate new questions incrementally and surgically, reducing cost while preserving proportionality. Importantly, this does not eliminate the need for methodology; it changes its form.

This incremental approach aligns closely with modern discovery principles. Rather than committing to expansive, front-loaded review protocols, counsel may adjust scope based on newly learned facts. From a proportionality standpoint, this allows discovery to remain responsive to the actual needs of the case rather than hypothetical worst-case scenarios.

⁶ “Southwest’s approach to using keyword searches and technology-assisted review in tandem does not offend the court’s expectation that the parties conduct a reasonable inquiry as required by the rules.” *Huntsman v. Sw. Airlines Co.*, No. 19-CV-00083-PJH, 2021 WL 3504154, at *3 (N.D. Cal. Aug. 10, 2021).

Critically, agentic AI does not diminish the lawyer’s supervisory role. To the contrary, it heightens the importance of thoughtful workflow design, ongoing oversight, and documentation. Courts evaluating such systems are likely to focus less on nomenclature and more on whether attorneys maintained control, exercised judgment, and could explain their process.

V. Practical Litigation Applications in Complex Matters

Practical adoption of AI in litigation succeeds when it is tied to discrete, outcome-oriented objectives rather than abstract technological ambition. Courts are generally receptive to tools that reduce burden, clarify facts, and streamline dispute resolution—particularly where those tools support early, informed decision-making.

AI-assisted early case assessment allows parties to identify key documents and factual themes at the outset of litigation. Rapid insight into the evidentiary landscape enables counsel to evaluate exposure, defenses, and settlement posture before discovery costs escalate. In many cases, this early clarity reduces the likelihood of protracted discovery disputes driven by incomplete or speculative understandings of the facts.

Deposition preparation illustrates another high-value application. AI tools can aggregate every reference to a witness across emails, memoranda, contracts, and prior testimony, surfacing patterns and inconsistencies that might otherwise remain buried. Counsel can then focus deposition time on resolving factual ambiguities and testing credibility, rather than locating foundational documents.

Chronology building similarly benefits from automation. AI-assisted “living chronologies” update continuously as new data is ingested, helping counsel understand how events unfold over time. This capability is particularly valuable in complex commercial matters involving extended timelines and multiple custodians.

Privilege log review and contractual analysis also benefit from AI-assisted pattern recognition. Tools can flag anomalies—such as non-lawyers included in communications claimed as privileged—or identify non-standard contractual provisions across large document sets. These applications do not replace legal analysis; they accelerate it.

VI. The Enduring Human Element: Ethics, Supervision, and Accountability

Despite increasing automation, the attorney’s role remains central. Rules governing discovery certification, supervision of non-lawyers, confidentiality, and data security apply with full force to AI-assisted workflows.

Certification obligations under Rule 26(g) remain unchanged. Attorneys must certify that discovery responses are complete and correct to the best of their knowledge, formed after a reasonable inquiry. The use of AI does not dilute this responsibility; it reshapes what constitutes a reasonable inquiry in data-intensive cases.

Professional responsibility rules governing supervision apply by analogy to AI systems. Attorneys must understand the capabilities and limitations of the tools they employ and must remain actively involved in decision-making. Courts are unlikely to fault the use of advanced technology where counsel can demonstrate meaningful oversight, quality control, and accountability.

Data security and confidentiality likewise remain critical considerations. AI tools must be deployed in secure environments with appropriate safeguards to protect privileged and sensitive information. When these conditions are met, AI does not undermine ethical obligations; it reinforces them by enabling more accurate, timely, and proportionate discovery.

VII. Conclusion: From Reviewer to Orchestrator

The evolution of AI in litigation reflects a broader shift in the lawyer's role. Attorneys are no longer expected to personally review every document to demonstrate diligence. Instead, they are expected to design, supervise, and defend discovery processes that are reasonable, transparent, and proportional.

As courts continue to refine discovery expectations in an era of expanding data, AI-assisted workflows are likely to become not merely acceptable, but expected in appropriate cases. The critical distinction will not be whether AI was used, but whether its use was thoughtful, defensible, and clearly explained.

The future of litigation belongs not to the fastest reviewers, but to lawyers who can orchestrate facts, technology, and professional judgment into a coherent and defensible strategy.