Drawing the Discovery Line: Courts Nationwide Begin Defining Privilege, Work Product, and Expert Discovery in the Era of Generative AI

Author: Abigail C. Kanavos, Associate, Cozen O'Connor
Contributor: Joel Maxwell, Member, Cozen O'Connor

Over the last few years, generative artificial intelligence has rapidly evolved from novelty to a standard litigation tool. It should come as no surprise that there is a growing body of decisions and case law from courts grappling with the impact of AI use on litigation, including specifically with respect to the discoverability of AI prompts. In this article, we discuss several recent rulings which suggest that the issue will turn less on the technology itself than on who used it, why it was used, and the role the AI interaction played in the litigation. In other words, courts nationwide are confronting the inevitable next question: when lawyers, litigants, and experts use generative AI in connection with litigation, who gets to see the prompts and output?

In a series of decisions issued in 2026, courts have begun to supply an answer, although hardly a uniform one. In *Conservation Law Foundation, Inc. v. Shell Oil Company*, a federal magistrate judge in Connecticut ordered disclosure of prompts used by an expert to assist in selecting documents for her analysis, characterizing that process as part of the expert's methodology. Civil No. 3:21-cv-00933 (VDO) (D. Conn.). In *United States v. Heppner*, the Southern District of New York declined to protect a criminal defendant's interactions with a consumer AI platform. (No. 25-cr-00503-JSR, S.D.N.Y. 2026). By contrast, *Warner v. Gilbarco, Inc.* (No. 2:24-cv-12333, at 11 (E.D. Mich Feb. 10, 2026)) and *Morgan v. V2X, Inc.*(D. Colo. Mar. 30, 2026), recognized work-product protection for AI-assisted litigation preparation by *pro se* civil litigants. And most recently, in *Assini v. Hayward*, 2026 N.Y. Slip Op. 26086 (Sup. Ct. Nassau Cnty. June 4, 2026), a New York state court adopted the reasoning of *Morgan* and quashed a subpoena directed to OpenAI seeking a litigant's ChatGPT records.

Taken together, these decisions suggest that courts are unlikely to adopt a categorical rule that AI prompts are either discoverable or protected. Instead, these courts have applied familiar doctrines such as attorney work-product, confidentiality, waiver, expert discovery, and proportionality to an unfamiliar technological setting.

For insurers and policyholders involved in complex litigation, the consequences extend well beyond discovery disputes over ChatGPT. AI-assisted claim analysis, expert work, litigation strategy, document review, damages modeling, underwriting analysis, and coverage evaluation may all generate a new category of electronically stored information, the discoverability of which will depend heavily on how the technology is used and the governing case law in the jurisdiction

Recent Cases

Conservation Law Foundation v. Shell Oil

On May 18, 2026, Magistrate Judge Thomas O. Farrish of the U.S. District Court for the District of Connecticut entered what appears to be the first federal order requiring disclosure of generative AI prompts as part of an expert's methodology. *Conservation Law Foundation, Inc. v. Shell Oil Company, et al.*, No. 3:21-cv-00933 (D. Conn.), ECF No. 970. In that case, the Conservation Law Foundation ("CLF") retained Dr. Naomi Oreskes as an expert concerning the history of climate-change science. In preparing her report, Dr. Oreskes and her research assistant used a commercially available generative AI tool to assist in filtering defendants' document productions

Copyright 2026 Cozen O'Connor. All Rights Reserved.

and identifying potentially relevant materials. Accordingly, Dr. Oreskes disclosed the use of AI in her expert report, including certain search terms and frequencies.

In response, the defense sought additional information concerning the prompts and outputs used during the document-selection process. After the parties were unable to resolve the dispute, the defense filed a motion to compel. Magistrate Judge Farrish granted the motion, reasoning that an expert's methodology is a proper subject of discovery and that the process by which Dr. Oreskes reduced defendants' document production to the subset ultimately used in her work formed part of that methodology.

The distinction is significant. The court did not treat the prompts merely as informal research notes or communications with a research assistant as suggested by parties in prior cases. Instead, it viewed them as part of the process by which the expert selected the universe of information on which her opinions would be provided.

CLF objected under 28 U.S.C. § 636(b)(1)(A) and Rule 72(a), arguing, among other things, that the ruling improperly applies Rule 26(b)'s general discovery standard rather than the more specific disclosure requirements applicable to experts under Rule 26(a)(2)(B)(ii). CLF also argued that no formal Rule 33 or Rule 34 discovery requests were served, that the parties' Rule 29 agreement protects the materials, and that the dispute overlaps with the defense's pending *Daubert* challenge. Ultimately, the district court stayed the production order pending resolution of the defense's objections, which have not yet been heard.

The practical significance of the district court's decision is already apparent. If prompts used to select documents constitute expert "methodology," AI-assisted expert workflows may become fertile territory for discovery. Parties could seek prompt histories, iterative refinements, outputs, system instructions, model information, uploaded source materials, and other records reflecting how an expert used AI to move from raw information to a final opinion. Moreover, for AI novices, failing to consider the potential discoverability of this information could be the difference between testifying at trial or having their opinions stricken by the court.

U.S. v. Heppner

The emerging cases concerning litigants' own AI use presents a slightly different issue. In *United States v. Heppner*, No. 25-cr-00503-JSR (S.D.N.Y. Feb. 17, 2026), Judge Jed S. Rakoff considered whether materials generated by a criminal defendant using the consumer version of Claude (Anthropic) were protected by attorney-client privilege or the work-product doctrine. In the widely-discussed decision, the defendant used Claude independently to analyze factual and legal issues relating to his criminal case and to develop potential defense strategies. Although he ultimately provided AI-generated documents to his attorneys, the court concluded that the materials were protected by neither attorney-client privilege nor the work product doctrine.

The court's reasoning focused heavily on confidentiality and the circumstances surrounding the AI use. Claude was a third-party platform rather than an attorney; the defendant had used it on his own initiative rather than at counsel's direction; and the applicable privacy terms permitted Anthropic to collect and use information submitted to the platform and, under certain circumstances, disclose that information to governmental authorities. Stated differently, the court's decision treated the consumer AI platform less like a private research tool and more like an independent third party with whom the defendant voluntarily shared information about his case. For this reason, *Heppner* represents the most restrictive approach among these early decisions.

Warner and Morgan: AI as a Tool for Creating Work Product

On the same day of the initial *Heppner* ruling, Magistrate Judge Anthony P. Patti of the Eastern District of Michigan reached a markedly different conclusion in *Warner v. Gilbarco, Inc.*, No. 2:24-cv-12333 (E.D. Mich. Feb. 10, 2026). There, a *pro se* civil plaintiff used ChatGPT in connection with her litigation and the defense argued they were entitled to any AI-related materials given their understanding that any work-product protection had been waived by using ChatGPT.

The court disagreed. It recognized that the work-product doctrine is not limited to materials created by attorneys and found that a *pro se* litigant may invoke the protection for materials prepared in anticipation of litigation. Critically, the court also rejected the proposition that using ChatGPT necessarily waived work-product protection.

Copyright 2026 Cozen O'Connor. All Rights Reserved.

In so ruling, the court held that work-product waiver ordinarily requires disclosure to an adversary or conduct substantially increasing the likelihood that the protected material will reach an adversary.

In other words, the court's ruling in *Warner* frames the inquiry differently from *Heppner*. Instead of asking primarily whether an AI provider is a "third party," the analysis asks whether the material itself constitutes litigation work product and, if so, whether use of the AI system resulted in the type of disclosure that traditionally constitutes waiver in similar circumstances without AI.

The District of Colorado developed that reasoning further in *Morgan v. V2X, Inc.*, No. 25-cv-01991-SKC-MDB (D. Colo. Mar. 30, 2026). Notably, *Morgan* expressly confronted *Heppner* and concluded that, in the context of a *pro se* litigant preparing for civil litigation, AI interactions can closely resemble confidential, iterative litigation work product. The court also rejected the proposition that the data-collection practices of publicly available AI systems necessarily "eliminate all expectations of privacy or automatically waive protections."

Thus, the distinction noted by the Warner and Morgan courts is increasingly important. If every disclosure to an AI provider automatically constituted work-product waiver, use of virtually any cloud-based litigation technology could raise analogous questions. In recognition of this concern, these courts instead focused on the purpose and character of the material and the traditional principles governing waiver. An important question to ask is whether the courts in Warner and Morgan would have decided differently if the party using a cloud-based AI was a *pro se* litigant? Are courts protecting *pro se* litigants where they otherwise would not?

Assini v. Hayward

The developing split reached New York state court in *Assini v. Hayward*, 2026 N.Y. Slip Op. 26086 (Sup. Ct. Nassau Cnty. June 4, 2026). In a commercial dispute, plaintiffs served a subpoena on OpenAI seeking prompts, inputs, uploaded materials, outputs, and other ChatGPT information associated with *pro se* defendant John Recchio. The subpoena sought materials concerning the parties, claims and defenses, court filings, motions, and litigation-related correspondence. However, Recchio promptly moved to quash, arguing that his ChatGPT interactions reflected private legal research, draft filings, litigation preparation, mental impressions, and strategy protected from disclosure in routine discovery.

Justice Rhonda E. Fischer agreed. In rendering its decision, the court considered *Heppner* but found the reasoning of *Morgan* more persuasive. In particular, the court embraced the view that a *pro se* litigant's use of AI to prepare litigation materials can resemble the type of confidential and strategy-driven iterative work that the work-product doctrine traditionally protects. As a result, the court quashed the subpoenas directed to OpenAI.

In evaluating the effect of the decision, it is first important to note that *Assini* appears to be among the first New York state-court decisions squarely confronting work-product protection in the context of generative AI. Second, it adds momentum to an emerging line of civil cases declining to treat the involvement of a commercial AI provider as an automatic waiver of work-product protection. The court nevertheless made clear that protection from discovery does not amount to judicial approval of unrestricted AI use. Justice Fischer specifically directed Recchio to New York's rules governing the use of artificial intelligence and warned that failure to comply with applicable requirements could result in sanctions.

In sum, *Assini* illustrates two principles that should not be conflated: AI-assisted litigation preparation may qualify for protection from discovery, while the manner in which AI is used in litigation remains subject to independent duties of accuracy, candor, confidentiality, and compliance with court rules for *pro se* litigants and represented parties alike.

Copyright 2026 Cozen O'Connor. All Rights Reserved.

A Framework Beginning to Emerge

These decisions initially appear difficult to reconcile. Viewed together, however, they suggest the beginnings of a workable framework.

The relevant question is increasingly not simply: *Was generative AI used?* but rather: *Who used it? For what purpose? What was disclosed? What role did the AI interaction play in the litigation? And what expectation of confidentiality was objectively reasonable under the circumstances?*

Stated differently, where an expert uses AI as part of the process for selecting or analyzing information underlying an opinion, as in *Conservation Law Foundation*, prompts may be characterized as part of the expert's methodology and therefore the prompts are potentially discoverable.

Where a represented party independently uses a consumer AI platform for legal analysis without counsel's direction and under terms permitting significant third-party use or disclosure, *Heppner* suggests that claims of privilege and work-product protection may face substantial scrutiny.

Where a *pro se* civil litigant uses AI to develop legal theories, drafts, and litigation strategy, *Warner*, *Morgan*, and now *Assini* suggest a materially stronger basis for work-product protection.

Where experts, represented parties and *pro se* civil litigants assert an expectation of confidentiality in connection with their use of a consumer AI platform, *Heppner* further suggests the terms and conditions (including any privacy policy) of the consumer AI platform will weigh heavily on a determination as to the reasonableness of that expectation.

While the resulting doctrinal map is still being drawn, one proposition is becoming increasingly apparent: the mere involvement of generative AI does not determine the privilege or discovery analysis. The technology is not the doctrine. Its function in the particular workflow is.

The decisions in *Heppner*, *Warner*, *Morgan*, *Assini*, and *Conservation Law Foundation* do not yet provide definitive answers. They do, however, reveal the contours of the reasoning the courts will apply in ruling on future discovery disputes.

In light of these decisions, litigants and their counsel are encouraged to ask not just whether generative AI may be used, but how it will be used, under whose direction, on what platform, subject to what terms and conditions, with what information, and with what consequences if the interaction later becomes Exhibit A in a discovery motion.

Author:



Abigail C. Kanavos | Associate | akanavos@cozen.com
Abigail represents international and domestic insurers in complex insurance coverage matters, arbitration, and litigation involving directors & officers liability, employment practices liability, professional liability, and commercial general liability claims. Abigail handles coverage litigation in state and federal courts across the country. She also assists commercial carriers in responding to first-party and third-party cyber insurance claims involving ransomware, fraudulent transfers, and social engineering, as well as overseeing vendor retention for appropriate breach notification and response efforts.

Contributor:



Joel Maxwell | Member | jmaxwell@cozen.com
Joel helps clients navigate a wide variety of matters, from construction defect and construction accidents to cyber breach-related claims, including class actions and actions to recoup client's funds after breaches, foodborne illness matters, motor vehicle accidents, and general liability matters. He has deep experience in all aspects of trial preparation and trying cases to verdict. He also skillfully negotiates settlement agreements when doing so is in the best interest of his client. He also acts as defense liaison counsel in one of the largest cyber breach cases pending in the county, involving over 13 million people.

Copyright 2026 Cozen O'Connor. All Rights Reserved.