AI, privilege, and discovery in view of 'Heppner' and 'Morgan'

By Nirav N. Desai, Esq., and David W. Haars, Esq., Sterne, Kessler, Goldstein & Fox PLLC

MAY 27, 2026

As parties and attorneys increasingly use generative artificial intelligence for litigation preparation, federal courts are beginning to wrestle with the significant practical impacts to litigation discovery.

Two recent decisions, *United States v. Heppner*, No. 25 CR. 503 (JSR), 2026 WL 436479 (S.D.N.Y. Feb. 17, 2026) and *Morgan v. V2X, Inc.*, No. 25–CV–01991–SKC–MDB, 2026 WL 864223, (D. Colo. Mar. 30, 2026), offer a study in contrasts, revealing that the discovery landscape around AI-generated materials depends in large part on the facts and context of the case. For legal counsel managing AI use across an organization, these cases merit close attention.

The *Heppner* decision

Heppner arose in a criminal prosecution in the Southern District of New York. Bradley Heppner, a corporate executive charged with securities fraud and related offenses, independently used the publicly available AI tool Claude to prepare reports that outlined defense strategy and what he might argue with respect to the facts and the law against anticipated charges.

Importantly, Heppner undertook these activities without direction from counsel. When FBI agents executed a search warrant at Heppner's home during his arrest, they seized approximately 31 documents memorializing these AI interactions (the "AI Documents").

Heppner's counsel asserted both attorney-client privilege and work product protection over the AI Documents, arguing that Heppner had input information received from counsel, created the documents for the purpose of speaking with counsel, and subsequently shared the contents with his attorneys. The government moved for a ruling that neither attorney-client privilege nor the work product doctrine applied. Judge Jed S. Rakoff agreed with the government.

Central to the court's conclusion regarding attorney-client privilege were two features: the absence of attorney involvement or direction in generating the AI Documents, and the lack of a reasonable expectation of confidentiality of a publicly available AI platform.

In the court's view, materials generated unilaterally by a represented defendant using a publicly available AI platform could not be considered communications with an attorney or for the purpose of obtaining legal advice because Claude is not an attorney and expressly disclaims providing legal advice. Moreover, because Claude's privacy policy permits the collection and retention of user inputs and AI outputs, use of that data for training, and disclosure to third parties, the AI Documents were not privileged because they lack confidentiality.

Taken together, Heppner *and* Morgan *show that disputes over AI-assisted litigation preparation are deeply fact and context specific.*

On work product, the court found two independent grounds for denial. Again, the AI Documents were not prepared "by or at the behest of counsel," and the documents did not reflect counsel's strategy at the time of their creation. While counsel acknowledged the AI Documents may have "affected" strategy going forward, they did not "reflect" counsel's thinking when Heppner generated them, and so did not implicate the core concern of the work product doctrine, shielding counsel's mental processes.

Notably, the court left open a narrow path. Had counsel directed Heppner to use Claude, the AI platform "might arguably be said to have functioned in a manner akin to a highly trained professional" acting as counsel's agent, a distinction that may prove significant in future cases.

The *Morgan* decision

Morgan, by contrast, involved a civil employment discrimination action brought by a *pro se* plaintiff in the District of Colorado. The defendant sought both disclosure of the AI tool the plaintiff had used and amendment of the protective order to restrict AI-based handling of confidential discovery.

The court granted the request in part, compelling disclosure of the identity of the AI tool and imposing new AI-specific protective order protections, but it did so with markedly different analysis than *Heppner*.

The court first resolved a threshold question: Does the workproduct doctrine apply at all to *pro se* litigants? The court answered yes. Unlike in *Heppner's* criminal posture, work product protections in *Morgan* were governed by Federal Rules of Civil Procedure Rule 26(b)(3). Because this Rule protects materials prepared by a "party," not merely by counsel, the court held that the doctrine applies to *pro se* litigants who necessarily serve dual roles of party and advocate.

Both decisions suggest that whether AI use was directed by counsel may be outcome-determinative.

The court then quickly disposed of the narrow discovery dispute, requiring the plaintiff to identify the AI tool that he used. The court explained that identifying the tool does not itself disclose legal theories or strategy, and that the defendant's request was reasonable given the legitimate need to assess whether its confidential information may have been exposed to a public AI platform.

The most noteworthy portion of *Morgan*, though, stems not from its narrow discovery ruling but from its broader discussion of confidentiality and privacy. The court questioned whether use of a public AI platform necessarily destroys confidentiality for privilege purposes, comparing AI to other cloud-based services like document storage or email platforms having similar privacy policies and wondering: "Does that mean that anyone with a Gmail account has forfeited all rights to confidentiality and privacy?"

The court analogized to Fourth Amendment jurisprudence, noting that the Supreme Court has held that possession of information by a third-party intermediary does not necessarily extinguish a reasonable expectation of privacy in that information. While acknowledging that Fourth Amendment doctrine does not govern discovery disputes, the court found the underlying principle instructive: Disclosure to an intermediary that is functionally and practically necessary in modern life does not automatically eliminate all expectations of privacy.

At the same time, the court was careful not to minimize the distinct risks AI systems pose. Unlike traditional email or document-storage services, the court noted that many AI platforms are expressly designed to train their models on user data. And as acknowledged in *Heppner*, the court noted that privacy policies often permit broad retention and disclosure to third parties.

Against that backdrop, the court concluded that even if AI use does not automatically waive attorney-client or workproduct protections, it nonetheless creates heightened risks to confidential discovery that courts cannot ignore.

These concerns ultimately shaped the court's amendment of the protective order. The court adopted a new provision prohibiting submission of confidential discovery to AI systems unless the provider is contractually barred from using the data for training or improvement, and from disclosing the data to third parties beyond what is necessary to provide the service. The court acknowledged that these requirements would likely disadvantage *pro se* litigants who may only have access to consumergrade AI tools but concluded that such limits were necessary to preserve the integrity of the discovery process.

Takeaways

Taken together, *Heppner* and *Morgan* show that disputes over AI-assisted litigation preparation are deeply fact and context specific. So far, courts have resisted categorical rules. Instead, they are asking familiar questions in a new technological setting: whose mental processes are at issue, what risks of adversary access exist, and how should discovery rules adapt to tools that are powerful, increasingly ubiquitous, and imperfectly understood. Practitioners can take several lessons from these cases:

1. AI conversations with public platforms may be discoverable, and internal policies should reflect that reality.

If employees use publicly available AI tools to think through legal problems, draft strategy memos, or explore litigation positions, those conversations may not be protected from discovery.

2. Attorney direction matters — and should be documented.

Both decisions suggest that whether AI use was directed by counsel may be outcome-determinative. Companies should work with litigation counsel to establish clear protocols for AI use in connection with anticipated or pending disputes, and to document that direction when it occurs.

3. Protective orders may now specifically restrict AI use, and platform selection has legal consequences.

Morgan signals that courts are prepared to impose AI-specific restrictions in protective orders — and that those restrictions may effectively bar the use of mainstream, consumer-grade AI tools for processing confidential discovery. The contractual terms of your AI provider — whether the platform trains on user data, or permits deletion, or third-party disclosures are restricted — may determine whether your use of the tool is permissible under a protective order. Enterprise-tier AI

accounts with appropriate data-handling commitments offer a materially different risk profile than consumer products.

4. The law is still developing.

Heppner and *Morgan* are not in direct conflict, but their holdings are not without tension. As one example, *Heppner* treated the AI platform's privacy policy as essentially dispositive on the question of confidentiality. *Morgan* pushed back on this framing, reasoning that routing information through a third-party system does not forfeit all privacy. As courts continue to grapple with new discovery issues arising from increasing use of AI tools, the only certainty is that context, not technology alone, will continue to shape the next generation of discovery and work product disputes.

About the authors

Nirav N. Desai (L) is a director in **Sterne, Kessler, Goldstein & Fox's** trial and appellate and electronics practice group**s**, where he focuses on patent litigation and global IP strategy. He has served in district court litigation matters as well as before the U.S. Patent and Trademark Office's Patent Trial and Appeal Board (PTAB), and the U.S. International Trade Commission (USITC) in Section 337 actions. He can be reached at ndesai@sternekessler.com. **David W. Haars** (R) is a director in the firm's electronics practice group. His practice focuses on post-grant proceedings before the U.S. Patent and Trademark Office's Patent Trial and Appeal Board (PTAB), as well as U.S. district court litigation and International Trade Commission (ITC) litigation. He can be reached at dhaars@sternekessler.com. The authors are based in Washington, D.C.

This article was first published on Reuters Legal News and Westlaw Today on May 27, 2026.