

Publications

February 25, 2026 • Updates

Court Rules on How Client Use of AI for Legal Strategy is Not Protected

Key Takeaways

- A federal court ruled that a defendant's independent use of a publicly available AI platform was not protected by attorney-client privilege or the work product doctrine.
- The decision highlights how confidentiality and counsel involvement remain central to privilege analysis, even as AI tools evolve.
- Organizations should take a proactive approach to AI governance and risk mitigation, especially with the current technological shift from generative AI to Agentic AI.

A new federal court ruling highlights why defendants should think twice before independently consulting AI tools. Last week, a federal judge in the Southern District of New York issued a 12-page decision in *USA vs. Heppner*, a case involving allegations of fraud in which the defendant independently used generative AI to prepare elements of his defense.

While *USA vs. Heppner* is in its infancy, and the ruling did not decide the merits of any allegations or claims, the ruling may have dealt a serious blow to the defendant's ability to defend himself based on how he used AI without involving his counsel. The issue before the court was "a question of first impression nationwide."

Specifically, the issue involves "whether, when a user communicates with a publicly available AI platform in connection with a pending criminal investigation, are the AI user's communications protected by attorney-client privilege or the work product doctrine?" The decision offers meaningful insight into how clients and attorneys should intentionally assess the risks involved in utilizing AI.

The *Heppner* Ruling: A First-of-Its-Kind Decision

On November 5, 2025, Heppner was arrested. He was released on bond five days later with a trial date set for April 6, 2026. In connection with his arrest, the FBI executed a search warrant at Heppner's home and seized documents and electronic devices, including documents memorializing Heppner's interactions with a generative AI platform.

Related People

- Romaine C. Marshall
- Matt A. Todd
- Jennifer Bauer
- Bryce H. Bailey

Related Capabilities

- Artificial Intelligence & Machine Learning
- Privacy & Cybersecurity

According to Heppner and his counsel, of the documents that were seized, 31 of them were protected by attorney-client privilege (ACP) and the work product doctrine (WPD) because they were communications Heppner had with generative AI platform Claude for and about his case and in preparation for discussions with his counsel. The court disagreed.

Why Heppner's Communications With AI Were Not Protected

ACP protects "communications (1) between a client and his or her attorney (2) that are intended to be, and in fact were, kept confidential, (3) for the purpose of obtaining or providing legal advice." The court ruled that Heppner's communications failed to satisfy these prerequisites for three reasons:

1. The AI documents were not communication between Heppner and his counsel.

Instead, the documents were from an AI platform acting in a non-attorney role, that lacks the "trusting human relationship...with a licensed professional who owes fiduciary duties and is subject to discipline." It is this trust that lies at the heart of the ACP.

2. The communications memorialized in the AI documents were not confidential.

Claude's privacy policy specifies that it retains interactions (including user inputs and system outputs) to train its AI and reserves the right to disclose such data to "third parties," including "governmental regulatory authorities." Thus, Heppner could not have a reasonable expectation of privacy or confidentiality.

3. Heppner did not communicate with Claude to obtain legal advice.

That advice can only be provided by an attorney, which Claude specifically disclaims being in writing, or via an attorney's agent (at the direction of counsel). Also, Claude could not have been acting as the attorney's agent because counsel never directed Heppner to use the platform for these purposes.

As to the WPD, this protection exists to "shelter[] the mental processes of the attorney, providing a privileged area within which he can analyze and prepare his client's case," but since the AI documents were not "prepared by or at the behest of counsel," they could not and did not reflect counsel's strategy at the time they were created by Heppner.

Understanding the Implications for AI Use by Clients and Attorneys

Embedded within the court's decision in *Heppner* are several caveats that could heavily influence when and how clients and attorneys should use AI and when such AI communications will be protected under the ACP and WPD.

First, *Heppner* articulated how AI could act as an agent of the attorney and thereby have its work product protected. Specifically, the court said if "counsel directed Heppner to use Claude, Claude might arguably be said to have functioned in a manner akin to a highly trained professional who may act as a lawyer's agent within the protection of the attorney-client privilege."

While this could apply to generative AI platforms like Claude, it also introduces the possibility of agentic AI serving a similar function in the future. Generally considered a step up from generative AI, agentic AI encompasses advanced technologies that can plan and execute complex tasks with little to no oversight.

As a leading AI Governance scholar and practitioner explained, agentic AI "chains

decisions together autonomously. You set a goal, and the system determines the path — finding clauses, drafting language and routing for signature — without waiting for approval at each step. That autonomy changes where liability lives.”¹

Of course, that shift in risk and liability makes it even more important to carefully consider when and how to deploy agentic AI, particularly when privilege is potentially involved. To be clear, AI — agentic, generative or otherwise — can only act as an attorney’s agent when an attorney is involved in the interaction.

Second, disclosures to AI can sever protection under the ACP whenever an AI platform’s terms, conditions and/or privacy notice dictate that inputs (and any corresponding outputs from those interactions) are not confidential. The same is true for AI platforms that may use data to train and reserve the right to disclose data to third parties for a variety of reasons.

This poses challenges for attorneys from both an ethics and client relationship perspective. But it suggests that private, closed environments with AI tools supported by enterprise-specific licenses that prohibit training and third-party disclosures are AI tools that attorneys and clients can feasibly use should they wish to preserve confidentiality and privilege.

What’s Next for AI Governance?

As the court artfully concluded: “Generative artificial intelligence presents a new frontier in the ongoing dialogue between technology and the law... [b]ut AI’s novelty does not mean that its use is not subject to longstanding legal principles, such as those governing the attorney-client privilege and the work product doctrine.”

Heppner clarifies how the ACP and WPD apply, and whether and how current use cases should be reevaluated. *Heppner* also provides another example of how AI Governance — the combination of principles, laws and policies that relate to AI’s development and deployment — continues to evolve to include doing the following:

- Evaluate the terms, conditions and privacy protections associated with your AI vendors, and advise on contract negotiations and third-party due diligence.
- Ensure your legal team is involved in any communications and work product development that should be privileged and ensure that any AI-supported workflow is structured in a way that preserves privilege where appropriate.
- Develop and deploy an AI policy, impact assessments and risk management program that systemically addresses risks like these.

We have routinely recommended holistic and specific approaches, with an emphasis on data collection practices and tailored AI risk mitigations, and recently we recommended beginning with a basic five-factor test to better understand your use case and tailor your risk mitigation strategy accordingly. For questions regarding AI Governance, please contact the authors.²

[1] How Agentic AI breaks the accountability model we’ve relied on (February 23, 2026). See *generally* <https://kenpriore.com/>.

[2] In addition, if you haven’t already, consider also the guidance last week in *The Hidden Risks of AI Notetakers: What Organizations Need to Evaluate Before Deployment* by Pavel (Pasha) Sternberg and Caitlin Smith.

