LOS ANGELES & SAN FRANCISCO
Daily Journal
www.dailyjournal.com
MONDAY, FEBRUARY 23, 2026

Consumer AI tools put confidentiality at risk

The use of consumer generative artificial intelligence tools such as ChatGPT and Claude may not be protected by attorney-client privilege or the work-product doctrine.

By Marc D. Alexander

SHUTTERSTOCK

On Feb. 10, 2026, Federal District Judge Jed Rakoff of the Southern District of New York ruled from the bench in *United States v. Heppner*, No. 25-cr-00503-JSR (S.D.N.Y. Feb. 10, 2026), that documents created by a client using a consumer-plan generative artificial intelligence tool, which the client later sent to his attorney, were not protected by the attorney-client privilege. The AI tool used was Claude. (*See* case report in Debevoise & Plimpton Data Blog, Feb. 11, 2026.)

Based on the court's ruling, communications with an AI tool are not communications with an attorney and therefore do not fall within the scope of the attorney-client privilege. Moreover, sharing non-privileged materials with an attorney does not retroactively transform them into privileged communications.

Also, the work-product doctrine did not apply under the circumstances presented, as the materials were created independently and were not prepared by or at the direction of counsel. Further, it appears confidentiality was lost by using the third-party AI. And it is difficult to see how documents created by a client using AI reflected the mental impressions, conclusions or opinions of law of an attorney.

Although the holding appears to be fact-specific, it raises broader concerns at a time when individuals increasingly rely on generative AI tools to research legal issues, create documents and analyze disputes before consulting counsel. A modest monthly subscription fee to a consumer generative AI service does not, by itself, ensure confidentiality or privilege protection. To the contrary, it may destroy confidentiality and privilege protection.

The same issues may arise with other consumer generative AI platforms, including ChatGPT. Consumer-level subscriptions are generally available to any individual with an email address and a payment method; unlike enterprise plans, consumer subscriptions typically do not allow negotiation of confidentiality provisions.

Depending on the applicable terms of service–which may change over time–consumer platforms may store user inputs for a defined period, and human review of stored content may occur for safety or system-improvement purposes. Policies regarding training use, retention periods, and access controls vary by provider and by subscription tier. By contrast, business and enterprise plans often provide enhanced contractual protections, including commitments not to use customer content for model training and stronger data segregation practices.

These distinctions are legally significant because attorney-client privilege depends upon a confidential communication between attorney and client made for the purpose of seeking or providing legal advice. Disclosure of confidential information to a third party can, in some circumstances, undermine the confidentiality necessary to sustain the privilege. Whether a particular AI interaction constitutes disclosure to a third party, and whether any such disclosure results in waiver, will likely depend on the specific facts, the governing jurisdiction and the provider's contractual terms.

The issue is particularly relevant in California in light of proposed Senate Bill 574, which addresses the use of generative AI in the practice of law. Among other provisions,

SB 574 proposes adding the following duty to the Business and Professions Code:

"6068.1. (a) It is the duty of an attorney using generative artificial intelligence to practice law to ensure all the following: (1) Confidential, personal identifying, or other nonpublic information is not entered into a public generative artificial intelligence system."

The proposed statute defines "generative artificial intelligence system" as:

"[A]n artificial intelligence system that can generate derived synthetic content, including text, images, video, and audio that emulates the structure and characteristics of the system's training data."

The term "public," however, is not expressly defined. Whether consumer generative AI platforms qualify as "public" systems under the proposed law would likely depend on how courts interpret the level of access, confidentiality safeguards and contractual protections provided by the service. Given the broad availability of consumer generative AI subscriptions and the limited ability of individual users to negotiate confidentiality protections, consumer-tier services could plausibly be characterized as "public" within the meaning of the statute.

Traditional legal research platforms such as Lexis and Westlaw historically have not operated as open generative AI systems. They provide curated databases of cases, statutes, regulations, treatises and commentary within subscription-based, contractually governed environments. While many such platforms now incorporate generative AI features, those tools operate within structured research ecosystems that typically provide contractual confidentiality protections and do not function as publicly accessible generative systems in the same manner as open consumer AI tools.

Courts often evaluate whether the client had a reasonable expectation of confidentiality and whether disclosure was consistent with maintaining privilege. How courts will treat disclosures to generative AI systems remains an evolving question.

Given the unsettled legal landscape, attorneys and clients should consider the following:

Legal research or investigative work performed using generative AI should, where possible, be conducted by counsel or at counsel's direction. Documenting that AI-assisted research was undertaken as part of counsel's litigation strategy may strengthen a claim of work-product protection.

When using consumer generative AI platforms, avoid entering confidential or personally identifying information. SB 574 defines such information broadly to include driver's license numbers, dates of birth, Social Security numbers, criminal identification numbers, addresses, telephone numbers, medical or psychiatric information, financial information, account numbers, and "[a]ny other content sealed by court order or deemed confidential by court rule or statute."

Anonymization and redaction should be the default practice when using consumer generative AI tools.

Higher-tier business or enterprise subscriptions may offer enhanced contractual protections, including commitments not to use customer content for training, defined data-retention limits, access controls, audit rights and written confidentiality provisions. These contractual safeguards may reduce the risk that AI use will compromise privilege or confidentiality obligations.

Attorneys, mediators, and arbitrators should consider adopting tailored disclosures regarding their use of generative AI tools and obtaining informed consent where appropriate. The California State Bar Committee on Professional Responsibility and Conduct has offered practical guidance on this point: "The lawyer should consider disclosure to their client that they intend to use generative AI in the representation, including how the technology will be used, and the benefits and risks of such use."

The State Bar Committee also offers guidance regarding the confidentiality of generative AI: "A lawyer must not input any confidential information of the client into any generative AI solution that lacks adequate confidentiality and security protections. A lawyer must anonymize client information and avoid entering details that can be used to identify the client. ... A lawyer or law firm should consult with IT professionals or cybersecurity experts to ensure that any AI system in which a lawyer would input confidential client information adheres to stringent security, confidentiality, and data retention protocols." Ethical duties of competence and confidentiality–including those reflected in ABA Model Rule 1.6 and corresponding state rules–require attorneys to understand the technology they use and to take reasonable steps to safeguard client information.

Generative AI tools are rapidly transforming legal research, drafting and analysis. However, the convenience and accessibility of consumer AI platforms do not eliminate longstanding requirements of confidentiality and privilege. Communications with public generative AI systems may not themselves be privileged, and disclosures of confidential information may jeopardize attorney-client privilege or work-product protection, depending on the circumstances.

As courts, legislatures, and ethics authorities continue to address these issues, attorneys should proceed cautiously, evaluate the terms of service governing any AI platform they use, and implement safeguards designed to preserve confidentiality and privilege. The law in this area is still developing, and new problems–and solutions–are likely to emerge.

Marc D. Alexander *is an attorney and an ADR neutral with Alternative Resolution Centers. He can be reached at AlexanderDisputeResolution@gmail.com.*