ORANGE COUNTY
LAWYER®
OCBA

Orange County Bar Association
Charitable Fund

125 YEARS
Est. 1901

RAISE YOUR GLASS

A WINE TASTING & CHARITY AWARDS BENEFIT

CELEBRATING 2026'S HONOREES

Nikki Presley Miliband
Individual Honoree

JAMS
Community Partner

PLUS:

Moral Rights Requiring Art Preservation
Civility in the Practice of Law
Lessons From Trail Running

www.ocbar.org AN OFFICIAL PUBLICATION OF THE ORANGE COUNTY BAR ASSOCIATION August 2026 Vol. 68 No. 8 $6.00

A GUIDE TO AI CLIENT SLOP

by AHMED SHAIKH



A client sends his lawyer an email about a complex legal strategy they have been working on. The email sounds nothing like the client. It asks several numbered questions that seem intelligent, and even well-considered, but something is off. It seems too well organized, with extensive bullet points, the rule of three deployed all over the place, odd transitions different from how humans talk, and lots of em dashes—as if the client were Emily Dickinson, J.K. Rowling, or Shakespeare—or something. Except, instead of literature, it is a far too lengthy letter to a lawyer. The letter wanders into irrelevant side quests, posed with the confidence of someone deeply knowledgeable about conversations the client never had with the lawyer. It is as if the client, whom the lawyer had figured for a street-smart business owner who likes to get to the point, suddenly turned into someone who has read the entire Western canon of classic literature and children's books and derived all the wrong lessons. The email has, in current parlance, an uncanny-valley quality. The client has fed his lawyer AI slop.

Lawyers charging by the hour may find this to be a further billing opportunity, though annoying. There is evidence that AI is adding to the volume of attorney work and billings without much apparent value added. *See* Elizabeth Bratton, *Clients' Barrage of AI-Generated Queries Risks Pushing Up Lawyers' Fees*, Fin. Times (Apr. 10, 2026).

Attorneys charging flat fees will regret doing so. Judges who see the results of artificial intelligence in their courtrooms are also annoyed enough to issue sanctions against lawyers and lecture pro se litigants on wasting their time. Khorri Atkinson, *Pro Se Workers' AI Assists Lead Courts to Warnings, Sanctions*, Bloomberg Law (Apr. 10, 2026), https://news.bloomberglaw.com/daily-labor-report/pro-se-workers-ai-assists-lead-courts-to-warnings-sanctions; *see also Noland v. Land of the Free, L.P.*, 114 Cal. App. 5th 426, 446 (2025). It seems everyone agrees: AI is a gigantic waste of time and extremely annoying.

Yet, we still use it.

The Jailhouse Snitch Who Lies to You

While attorney use of artificial intelligence tools has gotten much attention from ethics writers and bar associations, only recently has attention turned to the problem of client-produced AI slop either finding its way to their lawyers, or clients using AI as a thinking and research tool to wade through legal problems. State Bar of Cal., Standing Comm. on Pro. Resp. & Conduct, *Practical Guidance for the*

Use of Generative Artificial Intelligence in the Practice of Law (Nov. 2023), https://www.calbar.ca.gov/sites/default/files/portals/0/documents/ethics/Generative-AI-Practical-Guidance.pdf.

Many AI users hear from social media influencers on how to prompt AI to act as if it were a $2,000-an-hour lawyer. This article addresses two major problems here: hallucinations and discovery. AI is the legal equivalent of a jailhouse snitch that lies constantly, and AI users cannot help but bare their souls to this snitch because it's so mesmerizingly useful. Both problems are separable. The first is that it's a "snitch," in that there is often no expectation of privacy and that attorney-client privilege and confidentiality that may have been preserved, won't be. This is the dangerous part of AI. The other is that it lies, or what the industry calls "hallucinates." This is true for all AI models, and the newer "reasoning models" appear to be even more prolific as fabulists. *See* Cade Metz & Karen Weise, *A.I. Is Getting More Powerful, but Its Hallucinations Are Getting Worse*, N.Y. Times (May 5, 2025), https://www.nytimes.com/2025/05/05/technology/ai-hallucinations-chatgpt-google.html. The hallucination issue causes many sanctions orders against lawyers, but it is not only their problem. AI products are built for engagement, and there is widespread public health and product-liability concern about the effects of long, validating AI conversations on vulnerable users. This includes pending lawsuits arising from teen suicides connected to chatbot use. *See* Rhitu Chatterjee, *Their Teenage Sons Died by Suicide. Now, They Are Sounding an Alarm About AI Chatbots*, NPR (Sept. 19, 2025), https://www.npr.org/sections/shots-health-news/2025/09/19/nx-s1-5545749/ai-chatbots-safety-openai-meta-characterai-teens-suicide. These problems are part of the baggage that comes with AI but are beyond the scope of this article. For our purposes, the relevant point is narrower: AI lies to clients, who forward these lies to their lawyers. Lawyers use AI that lies to them, and they must spend time vetting everything AI tells them to keep themselves out of trouble.

The legal profession now has a Rule 1.1 duty to be competent while seemingly being asked to operate in a technology-infused funhouse of mirrors.

There are virtually no use cases where a client should chew the fat, or in this case, water and electricity, about sensitive legal topics with artificial intelligence that sends out information to a third party (common with ChatGPT, Claude, Grok, and Gemini). There is no chatbot privilege for the person with legal anxiety, as much as people would like there to be one. Recently, a federal judge found a defendant could not shield his conversations with the chatbot Claude from being used as evidence against him. *United States v. Heppner*, 820 F. Supp. 3d 292, 296 (S.D.N.Y. 2026).

The commentary around *Heppner* seems to make the issue look more novel than it was. It did address Anthropic's terms of service and issues like the expectation of privacy. But the defendant in the case was not talking to a lawyer or at the direction of a lawyer. The implications of the case may have been overdramatized in some of the legal press because we are talking about AI, and the result is meaningfully bad for a criminal defendant. Weak privacy policies have been a hallmark of cloud-based consumer software for two decades. However, the Stored Communications Act, 18 U.S.C. §§ 2701-2712, means attorneys doing work that is privileged and confidential reside in a different universe from ordinary consumers when it comes to such services. Attorneys must take reasonable steps to ensure confidentiality of the documents. However, using a cloud-based service that constantly scans your documents, or a chatbot that reads and responds to queries, does not automatically defeat or confer attorney-client privilege. There is no attorney-client privilege when there is no attorney or client on which to confer the privilege. This part is not terribly technical or novel. Responsible attorney use of these tools, however, is beyond the scope of this article.

Google Docs, which started in 2006 after Google acquired Writely, has always given Google the ability to review whatever is written on its platform. Consumer-tier Google Docs is paid for by an exchange of data so the user can be advertised to. The enterprise Workspace tiers have different terms but still scan all content for safety and abuse compliance. Software products from Google Docs to Adobe Photoshop now embed AI tools that send user content to external data centers for processing. None of those data centers is licensed to practice law. You cannot expect to use these services to commit crimes and have an expectation of privacy. If the police execute a search warrant at a defendant's home and find a handwritten memo confessing to securities fraud, the memo will be admissible against him at trial.

Chatbots are software, and in many respects, are no different from the cloud computing that has been with us for two decades. The government did not get the Claude chats from Anthropic; it got them from the defendant through a search warrant. He eventually shared the chats with counsel, who argued the logs should be attorney-client privileged. Counsel, however, had nothing to do with directing the client to use the chatbot. Heppner had reached for Claude on his own, treating the AI as a knowledgeable confidant. This was a mistake.

Client Chat Logs Are Inevitable

The client has not been sued yet. He has not been threatened with a lawsuit. It's just that something happened that makes him wonder if he might be, or if the thing that happened could lead to a lawsuit. His first inclination is to run the scenario through an AI chatbot. This conversation is logged and can go on for days. It may be insightful. It may even recommend he talk to a lawyer about asset protection. Like many areas of law, asset protection involves peril. Not only can the strategy fail, but it can backfire spectacularly and result in civil and criminal liability. The problem is that the attorney can do nothing to stop it. It is virtually inevitable a computer-literate client is going to use AI. It is also inevitable that litigators will seek out these AI chat logs.

Whom You Chat With Matters

A criminal defense attorney usually knows to tell his client to not talk about it to anyone. Operational discipline in this field is the stuff of legend, as are the clients who routinely ignore this advice. Lawyers in other contexts don't seem to do this as often as they should; however, legitimate third-party disclosures are more common in non-criminal contexts. Regardless of the context, it is about as bad an idea to tell your bartender you may have been drunk during a surgery last week as it is to say the same thing to a chatbot, regardless of what the privacy settings are, or write it on a notepad. It's all a jailhouse snitch.

Documents Prepared for Counsel

In general, documents a client provides in preparation for discussions with counsel are attorney-client privileged. The biggest risk is not that the documents will come out because Google or another company will be subpoenaed. It will be that the client will need to provide the documents because there was a warrant or a subpoena. The goal here is not to "hide" the existence of documents or pretend they don't exist. Rather, it is to preserve the attorney-client privilege under California Evidence Code section 954. For an increasing number of people, artificial intelligence is the default method of documenting things and understanding the world. A client may be interested in contextualizing her experiences, making sure she asks her lawyers the right

questions, and being as helpful to the cause of whatever she is being represented in. Perhaps artificial intelligence is a way for her to do this. This may horrify her lawyer. But it's the idea many people will have. Indeed, lawyers will do this themselves with a growing catalogue of AI products developed just for this kind of thing.

Step 1: Get a File Folder Marked "Attorney-Client Communications"

The client should have a password-protected computer nobody has access to. All documents intended to be communicated with the attorney should be placed in this folder. That includes privileged and non-privileged information. Anything that can be used to devise a strategy. In the event the client's documents are subject to discovery, the attorney will create a privilege log describing the documents, but would not provide the documents themselves. If the logs are legitimately privileged or not may be subject to litigation. However, the documents will not be automatically vacuumed up in the discovery process.

In many instances, potential clients won't know to do all of this before speaking with their lawyer. Many people will go straight to the AI chatbot to understand the lawyer letter, threatening email, or lawsuit.

If there is discovery against the client, the chatlog may be obtained from the client directly, or from the AI provider, which under federal law is a "remote computing service." 18 U.S.C. § 2711.

Heppner involved getting electronic records from the criminal defendant directly through a search warrant. Even having a chatlog from a source like ChatGPT saved in a designated folder shared with counsel in advance of any request for the records may help preserve privilege if the facts are better than *Heppner*. Privilege may be waived by sharing it with a third party, like an AI provider, but the mere fact that an outside technology provider has access does not by itself defeat privilege.

How a third party may obtain the documents depends on if the matter is civil or criminal. In civil cases, the Stored Communications Act, 18 U.S.C. § 2702, generally prohibits the AI provider from producing stored content in response to a subpoena from an opposing litigant. The civil discovery path runs through the client. The opposing party serves a Rule 34 production request (in federal practice) calling for the chatlog in the client's possession, custody, or control, and the client has the standard preservation and production obligations. In state practice, this is a demand for production under Code of Civil Procedure Section 2031.010(e). In criminal cases, the framework is different: under Section 2703 of the Stored Communications Act, the government can compel provider production directly with a warrant or court order, and providers regularly comply. Either way, if someone in litigation against your client with a reason to get the chatlog wants it, she can likely do so; the difference is whether the reach goes through the client or through the provider. Governments have it a bit easier.

In all cases, disclosure of chatlogs can be regulated through the discovery process. Counsel won't necessarily share the documents immediately, but would have the ability to share a privilege log describing the documents being withheld. Like any other claims of privilege, this can be subject to dispute by whoever wants the chatlog.

Step 2: Use an Open Weight Small Language Model

Most AI chatbots used by clients—like ChatGPT and Claude—are large language models. However, those clients interested in security can download open weight small language models. Many people already have a version of this since Apple and Microsoft both have tools that will help with writing or math without sending information to a third party, though—depending on how these tools are set up—information leaks can certainly happen. There are also small language model chatbots available, such as LM Studio. There is no sharing data with a third party, unless you want to, and the reasoning all happens on the computer itself. These small language models are more limited than the large language models, thus the name. They are, after all, processing information on a laptop and not a $50 billion data center.

It's also possible to use the large language model, like the paid versions of Claude or Gemini, with appropriate privacy settings, such as the data may not be used for training AI models. However, for the client to do this, it must be in service of the legal representation, at the direction of counsel in a manner that preserves privilege under the evidence code. This is true for both large language and small language models. The only difference is that you are taking more steps against sharing information with third parties with small language models. Because it is generative AI, an attorney may not want her client to engage in any of this kind of activity.

The best advice, if the client is willing to take it, is to stop using AI to generate correspondence. It's fine using it as an editor to clarify the client's communication with an on-device tool, but mucking up the correspondence likely does not serve the representation well, and may only contribute to legal fees that don't add much value to the client.

Step 3: If You Can't Stop the Slop, Bill for It

Of course, this does not address the main emotional and potentially business problem lawyers have with AI-generated emails and correspondence—those incessant em dashes and massive streams of useless information. This is a costly problem for lawyers and jurists.

The solution of moving correspondence from large to small language models, or taking other steps to preserve privilege, may stop the "snitch" problem, but it does nothing to stop the lying-to-you-all-the-time problem.

Attorneys, in their fee agreements and in attorney newsletter communications, should educate clients about the use of "Communications with Counsel" folders and avoiding the use of AI that go to third parties. But discouraging a stream of nonsense from clients can only come when there are financial incentives for clients to stop. For attorneys who charge hourly, this is relatively easy. Flat-fee and contingent-fee arrangements shift costs to lawyers and law firms, and that may warrant billing adjustments.

An "Artificial Intelligence Usage Policy" should govern the relationship between attorneys and their clients. These would naturally look different for lawyers who represent corporate entities in mergers and acquisitions than it would for lawyers who represent surgeons who may have been drunk on the job. Clients want competent and cost-effective legal services. Removing a few mirrors from the AI funhouse may help us get there.

Ahmed Shaikh *is a Certified Specialist in Taxation Law, and a Certified Specialist in Estate Planning, Trust and Probate Law by the State Bar of California. He practices Estate, Tax, Asset Protection and Charitable Planning. He is a Partner at Bohm Wildish and Matsen in Costa Mesa and his email is ashaikh@bohmwildish.com.*

This article first appeared in Orange County Lawyer*, August 2026 (Vol. 68 No. 8), p. 28.* *The views expressed herein are those of the author. They do not necessarily represent the views of* Orange County Lawyer *magazine, the Orange County Bar Association, the Orange County Bar Association Charitable Fund, or their staffs, contributors, or advertisers. All legal and other issues must be independently researched.*