LOS ANGELES & SAN FRANCISCO

Daily Journal

www.dailyjournal.com

FRIDAY, AUGUST 22, 2025

# AI chats aren't privileged: What California lawyers need to do now

**Sharing sensitive info with AI like ChatGPT isn't privileged -- lawyers and clients risk discovery and breaching confidentiality under California law.**

By Usama Kahf


Shutterstock

The head of OpenAI just set off alarm bells when he said that people commonly share deeply personal information with certain AI tools (like its own ChatGPT) on topics like health, relationships, or legal issues – assumedly not realizing these conversations are not protected under any sort of privilege. Sam Altman's wake-up call should especially jolt members of the bar since many might not consider the fact that conversations with ChatGPT and similar tools are not protected by attorney-client privilege, and providers could be compelled to produce chat logs in litigation. That means anything clients – or lawyers – type into an AI chat may be discoverable.

### Where privilege stops — and why AI falls outside it

Let's first examine Altman's statement about AI discourse falling outside privilege boundaries to see why that applies to attorney/client communications.

- Cal. Evid. Code § 954 protects confidential communications between a client and lawyer. By definition, chats with a third-party AI provider are not client-lawyer communications, so § 954 doesn't attach. (Justia Law)
- CCP § 2018.030 applies absolute work product protection for writings reflecting an attorney's impressions and qualified protection for other work product. Uploading facts, strategy, or drafts to an AI platform risks waiver or at least factual discovery of inputs/outputs.
- Bus. & Prof. Code § 6068(e)(1) actually creates a duty of confidentiality for California attorneys that's broader than privilege and requires lawyers to "maintain inviolate" client confidences. Putting client data into a tool that lacks enforceable confidentiality safeguards can violate this duty, even if privilege technically remains.

### What current law does (and doesn't) cover

Are there other arguments attorneys and clients make regarding the confidentiality of the data they input into AI systems? Two key ones would currently fail.

- California's main privacy law (the CCPA) regulates businesses' collection, use, and disclosure of personal information. The CCPA requires contracts obligating "service providers" to comply with comparable protections, but this does not create a privilege; at best it imposes data-handling duties and consumer rights (access, deletion, etc.).
- The California State Bar's Generative AI Practical Guidance warns lawyers not to input confidential client information into AI tools lacking adequate confidentiality and security. Likewise, ABA Formal Op. 512 stresses confidentiality in client communications when using GenAI.

### Practical risks for lawyers and clients

So what are risks if you ignore these alarm bells and use generative AI tools for confidential communications?

First and foremost - subpoena or discovery exposure. If a provider stores chats, litigants can seek them, and no privilege would bar production. Altman acknowledged this risk explicitly.

Next, you could see erosion of waiver and work-product. Pasting strategy memos or client facts into a third-party system can support arguments that confidentiality was not preserved, narrowing § 2018.030

protections. How long communications will be retained is the great unknown. Provider terms and retention policies vary, and CCPA contractual guardrails don't guarantee deletion or non-disclosure in response to lawful process.

Finally, clients attempting self-help with AI could cause you headaches. Clients may feed sensitive details to chatbots and later assume those exchanges are protected, only to find they are discoverable. Altman's remarks highlight how common this is becoming.

### 5 things California lawyers should do today

Given this new fear, what should you do? Here are five key action steps to consider.

**1. Lock down your stack.** Seek enterprise tools offering no-retention or short-retention modes, contractual promises of no training on inputs, data localization/encryption, audit logs, and subpoena notice.

**2. Treat AI vendors like agents — or don't use them for client data.** If you must use AI with client matters, use tools under written engagement that binds the vendor to confidentiality/security equivalent to your obligations under § 6068(e). Otherwise, do not input client-identifying facts or strategy.

**3. Update governance.** Adopt an AI use policy for your firm covering use cases, tool approvals, red-teaming, logging, and a ban on uploading privileged content absent client approval.

**4. Manage clients' expectations.** Add engagement-letter language and client alerts explaining that chats with consumer AI are not privileged and may be discoverable. Just like a patient may seek a "second opinion" through WebMD or other online medical portals, don't be surprised to learn that your clients are turning to ChatGPT to get advice about their case. Warn them directly not to do this.

**5. Preserve privilege when drafting.** Keep true mental-impressions writings offline or within secure firm systems. If using AI for style/editing, strip client identifiers and sensitive facts, and document your sanitization step.

### Sensible reform options for California

No article on this cutting-edge topic would be complete without a quick word on where this area of law should evolve. In a perfect world, the next few years would see the following reforms take hold:

- **Statutory safe harbor for lawyer-directed AI.** The Evidence Code should be amended to recognize privileged status for communications with AI when used by or at the direction of counsel under confidentiality-preserving conditions (e.g., where AI vendor would be an agent of the attorney bound by written terms mirroring § 6068(e) and non-waiver provisions).
- **CCPA enhancements for AI providers.** The state's privacy law should be enhanced to require AI vendors processing Californians' data to offer deletion-by-default for inputs, timely subpoena notice to enterprise customers, and explicit non-training options. All this can be done without suggesting any privilege is created.

**Usama Kahf** *is a partner at Fisher & Phillips LLP.*