

May 4, 2026

State Bar of California
180 Howard Street
San Francisco, CA 94105
communications@calbar.ca.gov

Subject: Public Comment on Proposed Amendments to the California Rules of Professional Conduct Related to Artificial Intelligence

Dear Members of the Committee on Professional Responsibility and Conduct:

Midpage is an AI-native legal research and drafting platform used by more than 300 law firms, from boutique practices to AmLaw 100 firms. Our products give lawyers AI tools to search, analyze, and draft from real cases, statutes, and regulations, with verifiable links back to source. We submit this comment because building AI tools that lawyers can use consistent with the duty of candor is the problem we work on every day.

Support for the Rulemaking

We support the State Bar’s effort to incorporate the *2023 Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law* (the “Practical Guidance”) into the California Rules of Professional Conduct. Our comment addresses the framing of proposed Comment [3] to Rule 3.3 in two respects:

First, the comment should remain technology-neutral and focus on the lawyer’s diligence, not the source of the citation.

Second, the comment should avoid inadvertently discouraging adoption of the AI tools that, as empirical research now shows, can actually help lawyers meet their duties of candor and competence.

The California Supreme Court’s August 22, 2025 directive asked the State Bar to consider incorporating Practical Guidance principles into the Rules and to address agentic AI tools. Our comment is consistent with that directive. AI-related guidance belongs within the Bar’s framework, but the most durable home for technology-specific guidance is the Practical Guidance itself, which can be updated as the technology evolves. The rule comments, which are difficult to amend once adopted, should stay technology-neutral.

We offer these observations in service of the same purposes that animate the Rules: to protect the integrity of the legal system and to promote the administration of justice and confidence in the legal profession.¹

¹ Cal. Rules Prof. Conduct, rule 1.0(a).

Concerns Regarding the Framing of Proposed Comment [3]

We support Proposed Comment [3]’s clarification that the duty of candor “includes the obligation to verify the accuracy and existence of cited authorities, including ensuring no cited authority is fabricated, misstated, or taken out of context, before submission to a tribunal.” That obligation, as we discuss below, is already part of the duty of candor.

The difficulty is the comment’s closing clause, which singles out “authorities generated or assisted by artificial intelligence or other technological tools.” This framing risks two unintended consequences:

First, by singling out AI, the comment invites a negative inference: that citations from non-AI sources carry a lesser obligation. They do not. The duty applies uniformly to every citation, from every source.

Second, foregrounding the technology shifts attention from the question that actually controls—did the lawyer exercise reasonable diligence and personally verify the authority?—to a secondary question about which tool was used.

The Lawyer’s Verification Duty Under Rule 3.3

The duty of candor has always centered on the lawyer, not on the source of the citation. A lawyer who knowingly cites a non-existent case, relies on an overruled decision, or misquotes authority violates Rule 3.3(a)(1) and (a)(2). Good faith reliance on a citation that turns out to be wrong does not insulate the attorney from responsibility for the filing.

The duty of candor has never been delegable to a research tool. It does not become delegable now simply because the tool is more sophisticated. As Judge Castel explained in *Mata v. Avianca, Inc.*, “existing rules impose a gatekeeping role on attorneys to ensure the accuracy of their filings.”² The California Court of Appeal articulated the same principle in *Noland v. Land of the Free, L.P.*: “no brief, pleading, motion, or any other paper filed in any court should contain any citations—*whether provided by generative AI or any other source*—that the attorney responsible for submitting the pleading has not personally read and verified.”³

That phrase—*or any other source*—captures both the recent AI sanctions cases and the larger body of pre-AI cases sanctioning counsel for citing authorities that didn’t stand for the proposition asserted, were overruled, or were imperfectly quoted from secondary sources. The verification obligation is a feature of the lawyer’s practice, not of any tool.⁴

² *Mata v. Avianca, Inc.*, 678 F. Supp. 3d 443, 461 (S.D.N.Y. 2023).

³ *Noland v. Land of the Free, L.P.* (2025) 114 Cal. App. 5th 426, 431 (emphasis added).

⁴ See Michael D. Murray, *Algorithmic Ethics in an Era of Agentic AI Advocacy*, 16 St. Mary’s J. on Legal Malpractice & Ethics 279, at 312 (2026) (“[T]he lawyer, not the AI, is

The verification duty exists because no research method is infallible. No methodology—digest searches, Boolean queries, natural-language search, or AI-assisted retrieval—surfaces every relevant authority or guarantees accurate characterization in subsequent drafting. Accordingly, “the duties imposed by Rule 11 require that attorneys read, and thereby confirm the existence and validity of, the legal authorities on which they rely [as there is] no other way to ensure that the arguments made based on those authorities are warranted by existing law or otherwise legally tenable.”⁵ Rule 3.3 reflects the same principle: verification is the duty of the lawyer whose reading and judgment stand behind the filing.

Material Differences Among AI Tools

We are not asking the Bar to ignore AI. We are asking it to approach AI with nuance informed by how these tools actually work. The proposed framing treats “AI” as a single category, but the category contains tools that behave differently from one another, and those differences matter for the verification obligation the rule is trying to reinforce.

The hallucination problem that has prompted this rulemaking is principally a feature of *ungrounded* use, i.e., situations in which a generative model produces text from learned statistical patterns without retrieving specific source documents. *Grounded* systems, also called retrieval-augmented generation, or “RAG,” work differently. They first retrieve source material from a defined corpus, then constrain the model’s output to that material. This substantially reduces the risk of fabrication, though it does not eliminate it.

Two features make grounded tools verifiable in practice: source linkage to documents the lawyer can access in their original form, and citation traceability that allows each authority to be verified with reasonable effort. These are exactly the kinds of distinctions that belong in the Practical Guidance, not in a rule comment. Embedding them in the rule as currently proposed would flatten meaningful differences among today’s tools and freeze that frame even as the technology changes; locating them in the Practical Guidance lets the Bar update its guidance as the tools evolve, while the rule remains anchored to the lawyer’s verification duty.

The architecture of an AI tool affects how often it fabricates citations, and the empirical evidence is now supporting that intuition. A recent randomized controlled trial of next-generation legal AI tools found that work produced with a retrieval-augmented research tool contained substantially fewer hallucinated citations than work produced with an ungrounded

absolutely and unequivocally responsible for the accuracy and veracity of every statement made in a court filing.”).

⁵ *Park v. Kim*, 91 F.4th 610, 615 (2d Cir. 2024) (internal quotations omitted).

reasoning model.⁶ Both tools also produced significant productivity gains across most tested tasks.⁷ And courts that have sanctioned lawyers for AI misuse have, in those same opinions, recognized that AI, used responsibly, can benefit legal practice and access to justice.⁸ The emerging consensus, reflected in ABA Formal Opinion 512, is becoming clearer: AI is permissible; the lawyer must verify; the depth of verification depends on the tool and the task.⁹

A rule comment that signals AI use as the locus of disciplinary concern, whether by intention or by negative inference, risks dampening adoption of the very category of tools that have the potential to help lawyers fulfill their duties of candor under Rule 3.3 and competence under Rule 1.1. No party to this rulemaking would intend that result, but it is the predictable consequence of a framing that targets AI broadly rather than ungrounded AI use specifically. The stakes will only grow as AI evolves toward agentic configurations, where the lawyer’s verification work will depend directly on the auditability and traceability of the system’s decision trail.

To address these concerns, we respectfully suggest revising Comment [3] in the following terms:

“A lawyer’s duty of candor towards the tribunal includes the obligation to verify the accuracy and existence of any cited authority before submission to the tribunal, and to ensure that no cited authority is fabricated, misstated, or taken out of context. This obligation applies regardless of how the authority was identified, and is not lessened by the use of any technological or research tool.”

⁶ Daniel Schwarcz, Sam Manning, J.J. Prescott, Patrick Barry, David R. Cleveland & Beverly Rich, *AI-Powered Lawyering: AI Reasoning Models, Retrieval Augmented Generation, and the Future of Legal Practice*, J.L. & Empirical Analysis (2026).

⁷ *Id.*

⁸ See, e.g., *Wadsworth v. Walmart, Inc.*, 348 F.R.D. 489, 492-93 (D. Wyo. 2025) (“When done right, AI can be incredibly beneficial for attorneys and the public.”); *State v. Coleman*, 2026-Ohio-965, ¶ 132 (“AI tools, when used responsibly and under appropriate supervision, have the potential to improve access to justice, increase efficiency, and enhance the quality of legal services.”).

⁹ ABA Formal Op. 512 at 3–4 (a lawyer’s “reliance on, or submission of, a GAI tool’s output—without an appropriate degree of independent verification or review of its output—could violate the duty to provide competent representation as required by Model Rule 1.1,” and “[t]he appropriate amount of independent verification or review required to satisfy Rule 1.1 will necessarily depend on the GAI tool and the specific task that it performs as part of the lawyer’s representation of a client”).

This formulation preserves the protective purpose of the proposed comment, aligns with the emerging consensus regarding AI use, and refocuses the rule on the conduct of the attorney.

Conclusion

We respectfully urge the Bar to revise proposed Comment [3] in two ways: adopt a technology-neutral verification standard focused on the lawyer’s diligence, and move technology-specific guidance, including the distinction between grounded and ungrounded AI tools, into the Practical Guidance, where it can evolve with the technology. We believe these refinements would advance the very purposes the Rules are intended to serve. Midpage stands ready to serve as a resource in this further work.

Respectfully submitted,

Nitesh Daryanani

Nitesh Daryanani
Counsel, Midpage AI, Inc.
nitesh.daryanani@midpage.ai