

**THE STATE BAR OF CALIFORNIA
STANDING COMMITTEE ON
PROFESSIONAL RESPONSIBILITY AND CONDUCT**

**PRACTICAL GUIDANCE FOR THE USE OF
GENERATIVE ARTIFICIAL INTELLIGENCE IN THE PRACTICE OF LAW**

EXECUTIVE SUMMARY

The Committee on Professional Responsibility and Conduct (COPRAC) issued its original *Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law* in November 2023. Since that time, developments in artificial intelligence (AI) technology have warranted review and revision of the guidance. This 2026 Practical Guidance replaces the 2023 version and, at the request of the California Supreme Court, addresses the unique challenges presented by the use of agentic AI. This update reflects technological advances and further clarifies the intersection between AI in legal practice and the core ethical duties of competence, confidentiality, communication, and candor to the court, among others.

AI tools are increasingly embedded in the everyday practice and administration of law and are becoming ubiquitous across legal services. Many technologies already used by lawyers, such as legal research platforms, document review systems, and practice management tools, incorporate some level of AI, even when not traditionally thought of as AI or labeled as such. This Practical Guidance is intended to address the ethical considerations arising from the expanding use of generative AI-enabled tools by lawyers. It clarifies how existing professional responsibility obligations apply when lawyers use or rely on these technologies in the provision of legal services.

Artificial intelligence refers to computer systems or software that are designed to perform tasks that would ordinarily require human intelligence, such as analyzing information, identifying patterns, making predictions or recommendations, or generating responses, and it is based on data or programmed models rather than fixed instructions alone.

Generative AI

Generative AI is a subset of artificial intelligence designed to create new content, such as text, images, audio, video, computer code, or analytical outputs, in response to user prompts or inputs based on patterns learned from data rather than predefined processes.

Generative AI is a technology that has wide-ranging applications for the practice of law and administrative functions of the legal practice for all licensees, regardless of firm size, and all practice areas. Like any technology, generative AI must be used in a manner that conforms to a

lawyer’s professional responsibility obligations, including those set forth in the California Rules of Professional Conduct and the State Bar Act. A lawyer must understand the risks and benefits of the technology used in connection with providing legal services. How these obligations apply will depend on a host of factors, including the client, the matter, the practice area, the firm size, and the tools themselves, ranging from free and readily available to custom-built, proprietary formats.

Generative AI presents unique challenges: It uses large volumes of data, there are many competing AI models and products, and, even for those who create generative AI products, there is a lack of clarity as to how it works. In addition, generative AI may produce outputs that are probabilistic, not deterministic, and may generate plausible but inaccurate information, including purported citations to legal authority that does not exist (frequently called “hallucinations”). Generative AI also poses the risk that information provided as an input may be used for learning purposes by the generative AI tool and, as a result, become part of, or be used to generate, a response to another user of the generative AI product. In a similar respect, generative AI poses the risk of encouraging greater reliance and trust on its outputs because of its purpose to generate responses and its ability to do so in a manner that projects confidence and effectively emulates human responses. A lawyer should consider these and other risks before using generative AI in providing legal services.

Agentic AI

Generative AI products have increasingly incorporated “agentic” capabilities, which can enable systems to autonomously perform tasks or workflows without human prompting. Agentic systems are built not just with multistep functionality but also can pursue goals, “plan,” select tools, and even execute tasks without immediate human intervention. Examples include pleading revisions across multiple iterations, preparing discovery responses, coordinating review of documents, and autonomously facilitating client intake.

Agentic AI introduces an additional degree of autonomy, including the ability to initiate tasks or interact with external tools, and the capacity to conclude tasks without human review or approval. Yet, this autonomy does not satisfy a lawyer’s duty to exercise independent judgment, nor does it alter a lawyer’s existing ethical obligations. Rather, the ability of agentic AI systems to initiate tasks or interact with external tools increases the risk that a lawyer may rely on automated processes without contemporaneous review. Lawyers must not deploy agentic systems in a manner that allows the system to make substantive legal determinations, communicate legal advice, prepare and file pleadings, or otherwise act in a representative capacity without meaningful lawyer supervision¹ and review. The lawyer’s duty of competence requires understanding the nature and extent of an agentic system’s autonomy, the

¹ The terms “supervise,” “supervision,” and “supervisory” are used in their ordinary sense unless referring specifically to a lawyer’s supervisory responsibilities under rules 5.1, 5.2, or 5.3 of the California Rules of Professional Conduct. COPRAC notes that the duty of competence requires a lawyer to supervise technological agents.

circumstances under which it may act without prompting, the data sources it may access, and the safeguards necessary to prevent errors and ethical violations.

Conclusion

The following Practical Guidance is based on professional responsibility obligations for lawyers and outlines the ethical duties required whenever a lawyer relies on computational tools that generate, select, summarize, or act upon legal information. References to “generative artificial intelligence” include any current *or* future computational system that generates outputs, recommendations, or actions.

Critically, any use of AI must not diminish or abdicate professional judgment. A lawyer remains fully responsible for any outputs and work product generated with the assistance of AI. Moreover, the greater the degree of autonomy afforded to an AI system (i.e., agentic systems), the more important it is for a lawyer to implement supervisory controls and verification mechanisms.

PRACTICAL GUIDANCE

Applicable Authorities	Practical Guidance
<p>Duties of Competence and Diligence</p> <p>Rule 1.1</p> <p>Rule 1.3</p>	<p>The duty of competence requires more than the mere detection and elimination of false or fabricated AI-generated results (i.e., hallucinations). Competence, at its core, requires that the lawyer (and not the technology) retains control over strategic decision-making and the ultimate exercise of professional judgment in the representation.</p> <p>When using AI, a lawyer has two related duties of competence. First, a lawyer must attain baseline technological competence by developing a reasonable understanding of the relevant AI system’s capabilities, data sources, limitations, and material risks and benefits before deploying it in connection with legal services. Second, a lawyer must exercise independent professional judgment by reviewing, verifying, and correcting AI-generated outputs consistent with the learning and skill reasonably necessary for the representation. Competent use of AI-enabled tools, therefore, requires not only familiarity with the technology itself but also the ongoing application of the lawyer’s independent professional judgment to all outputs used in providing legal services.</p> <p>As noted, AI systems are known to produce outputs that may be false, misleading, fabricated, incomplete, or biased, and reliance on such outputs without appropriate scrutiny and review is inconsistent with a lawyer’s professional obligations.</p> <p>A lawyer’s professional judgment cannot be delegated to AI and remains the lawyer’s responsibility at all times.</p> <p>These considerations are particularly important when utilizing agentic AI systems that incorporate increased autonomy into legal workflows. Agentic systems may initiate tasks, sequence legal steps, prioritize issues, or even interact with external tools without continuous lawyer prompting. Such functionality necessarily increases the risk that strategic decisions typically within the ambit of the lawyer will be shaped by automated processes rather than by the lawyer’s independent judgement. The risk associated with hallucinated outputs can also be compounded when using agentic AI.</p> <p>As a result, lawyers remain responsible for reviewing and approving all outputs, decisions, advice, and filings. Competence in this context includes understanding system-level characteristics, such as autonomous task execution, retrieval mechanisms, and common failure modes, consistent with rule 1.1, as well as implementing appropriate</p>

	<p>supervision and verification processes.</p> <p>The greater the level of system autonomy, the greater the lawyer’s obligation to implement oversight mechanisms sufficient to ensure that professional judgment remains with the lawyer. Moreover, because AI systems evolve over time through updates or model changes, a lawyer’s duty of competence includes periodic reassessment of the system’s capabilities and risks, including when deploying the system for new legal tasks.</p>
<p>Duty of Confidentiality</p> <p>Bus. & Prof. Code, § 6068, subd. (e)</p> <p>Rule 1.6</p> <p>Rule 1.8.2</p>	<p>Generative AI products often utilize the information that is input by the user, including prompts and uploaded documents or resources, to further train or refine the AI, and might also share such information with third parties or use it for other purposes. Even if the product does not utilize or share inputted information, it may lack reasonable or adequate security (e.g., it could store or transmit information in ways that raise confidentiality concerns).</p> <p>The confidentiality risks presented by generative AI vary depending on whether the system is used as an isolated prompt-based tool, integrated with other applications and data, or deployed as an autonomous agent (including with ongoing data access and autonomy of action). The degree of lawyer diligence and supervision must correspond to the level of system access and autonomy.</p> <p>As a general matter, a lawyer must not input any confidential information of the client into a generative AI solution that may present material risks to confidentiality or security, absent informed client consent (see rule 1.0.1(e)) as to the underlying risks.</p> <p>A lawyer has a duty to make reasonable efforts to protect client confidences when using generative AI. In fulfilling this duty, a lawyer should take reasonable steps to understand how a generative AI product collects, uses, stores, and discloses information provided by the user. Reasonable efforts require more than reliance on generalized marketing assurances. This may include reviewing terms of use, privacy policies, or vendor documentation, recognizing that such materials may vary in technical complexity. It is important to understand that, in some cases, a user may agree to the terms of use and privacy policy simply by using the product. Depending on the circumstances, reasonable efforts may include consulting with qualified IT professionals or cybersecurity experts to ensure that any AI system in which a lawyer would input confidential client information adheres to reasonable security, confidentiality, and data retention protocols. Breaches of confidentiality could occur via both prompt content and uploaded materials.</p>

	<p>Agentic AI systems may be configured to access internal firm systems, such as email, messaging platforms, document management systems, knowledge bases, client files, or calendaring systems. Unlike isolated prompt-based tools, such systems may have persistent or automated access to large volumes of confidential client information, raising acute confidentiality considerations under rule 1.6 and Business and Professions Code section 6068(e)(1).</p> <p>Lawyers must carefully evaluate and limit the scope of such access. Unrestricted or poorly configured agentic systems may unintentionally disclose confidential information (including across different matters) and even expose privileged material. As a result, a lawyer remains responsible for ensuring that any agentic AI system is configured in a manner consistent with their duty of confidentiality and that appropriate monitoring and access controls are implemented. A lawyer must not deploy an agentic AI system in a manner that permits autonomous external transmission of client information, including automated communications, filings, or data transfers, without appropriate safeguards and human review.</p>
<p>Duty to Comply with the Law</p> <p>Bus. & Prof. Code, § 6068(a)</p> <p>Rule 8.4</p> <p>Rule 1.2.1</p>	<p>A lawyer must comply with the law and cannot counsel a client to engage, or assist a client in conduct that the lawyer knows is a violation of any law, rule, or ruling of a tribunal when using AI tools, including laws and rulings that are applicable to the use of such technology.</p> <p>There are many applicable legal issues surrounding generative AI, including but not limited to compliance with AI-specific laws, privacy laws, cross-border data transfer laws, intellectual property laws, and cybersecurity concerns. Because the legal and regulatory landscape governing generative AI is rapidly evolving, a lawyer should continuously analyze and remain reasonably informed about the relevant laws and regulations applicable to the lawyer, the client, and the specific use of the technology at issue.</p>

<p>Duty to Supervise Lawyers and Nonlawyers, Responsibilities of Subordinate Lawyers</p> <p>Rule 5.1</p> <p>Rule 5.2</p> <p>Rule 5.3</p>	<p>Managerial and supervisory lawyers should establish clear policies regarding the uses of generative AI, including more autonomous or agentic AI tools, and make reasonable efforts to provide reasonable assurance that the conduct of the lawyers and nonlawyers under supervision complies with the lawyers’ professional obligations when using such technologies. This includes providing appropriate training on the ethical and practical aspects, and pitfalls, of any AI-enabled tools that perform tasks with limited or no real-time human direction.</p> <p>As AI technologies continue to evolve, managerial and supervisory lawyers should review and update applicable firm policies, controls, and associated training accordingly to address new risks and use cases, including those that may reasonably be anticipated as AI systems become more autonomous or are integrated into legal workflows. Moreover, lawyers should ensure that their lawyer and nonlawyer staff receive periodic training on the appropriate use of AI and safeguards surrounding the use of AI.</p> <p>A subordinate lawyer must not use AI-enabled tools, including agentic AI systems, at the direction of a supervisory lawyer in a manner that violates the subordinate lawyer’s professional responsibility and obligations and must exercise independent professional judgment in evaluating whether such use is consistent with those obligations.</p>
<p>Communication Regarding Generative AI Use</p> <p>Rule 1.2</p> <p>Rule 1.4</p>	<p>A lawyer must evaluate their communication obligations throughout the representation based on the facts and circumstances, including the technology, risks associated with AI use, scope of the representation, and sophistication of the client.</p> <p>A lawyer must consider disclosure to their client that they intend to use AI in the representation, including how the technology will be used, and the benefits and risks of such use. For example, disclosure may be appropriate where AI use materially affects decision-making processes. The use of AI in any form does not alter the lawyer’s obligation to ensure that the client understands who is providing legal advice and how substantive decisions affecting the representation are being made.</p> <p>A lawyer must also follow any applicable client instructions or guidelines that may restrict or limit the use of AI in connection with the representation of the client.</p>

<p>Charging for Work Produced by Generative AI and Generative AI Costs</p> <p>Rule 1.5</p> <p>Bus. & Prof. Code, §§ 6147–6148</p>	<p>A lawyer may use generative AI to create work product and may charge for actual time spent (e.g., crafting or refining generative AI inputs and prompts, or reviewing and editing generative AI outputs). Hourly billing must reflect the time actually spent on the client’s matter, recognizing that generative AI may reduce the time required for certain tasks. Lawyers may consider using alternative billing structures, such as flat fees, for more predictable assignments and building the time savings created by AI use into their competitive rates for such assignments. Regardless of the billing structure chosen, fee agreements should clearly communicate how fees and costs are calculated and the extent to which generative AI factors into the lawyer’s rates.</p> <p>The treatment of generative AI costs depends on whether such costs constitute general overhead or recoverable expenses incurred specifically for a client’s matter. Subscription fees for generative AI tools that provide general office functionality, such as drafting assistance, research capabilities, or document review, typically constitute overhead expenses similar to library maintenance or general computer systems and thus should be absorbed within the lawyer’s fee rather than charged separately to clients.</p> <p>However, lawyers may charge clients for generative AI costs that are specifically incurred for the client’s matter and go beyond general office operations, such as per use fees for specialized AI tools or costs for custom AI implementations developed for a particular matter. When charging such costs, the lawyer must ensure the fee agreement clearly discloses that generative AI costs may be billed separately, the charges must reasonably reflect the lawyer’s actual cost, and no markup or profit element may be added without the client’s informed written consent.</p>
--	--

<p>Candor to the Tribunal; and Meritorious Claims and Contentions</p> <p>Rule 3.1</p> <p>Rule 3.3</p>	<p>The lawyer’s duty of candor to the tribunal cannot be delegated to AI. The use (or ability) of AI to research, draft, summarize, or generate legal analysis does not diminish the lawyer’s personal responsibility for the truthfulness and legal sufficiency of any submission to the court.</p> <p>In that respect, a lawyer must review all outputs produced using AI tools for accuracy, including but not limited to analysis and citations to authority before submission to the court, and must independently verify and correct any errors or misleading statements made to the court, regardless of whether such outputs were generated with or without real-time human direction.</p> <p>A lawyer should also check for any rules, orders, or other requirements in the relevant jurisdiction that may necessitate the disclosure of the use of AI tools, including generative or autonomous AI systems, and should comply with any such disclosure obligations as applicable.</p> <p>Agentic AI tools may present additional risks. Where agentic AI is used in connection with court filings, lawyers must ensure that no document is transmitted to the court without lawyer review and approval. Lawyers must not permit AI systems to autonomously file documents, communicate with the court, or make representations on the lawyer’s behalf.</p>
<p>Prohibition on Discrimination, Harassment, and Retaliation</p> <p>Rule 8.4.1</p>	<p>Lawyers should be aware of the risk of biased outputs and take appropriate steps to ensure compliance with relevant laws, regulations, and rules of professional conduct. This can be of particular concern when using agentic AI as multiple agents can compound any underlying bias.</p> <p>For example, the use of AI in connection with employment actions or the screening of clients or candidates may present a risk of unlawful or unethical bias.</p> <p>Lawyers should engage in continuous learning about AI biases and their implications in legal practice. Firms should establish policies and mechanisms to identify and address potential AI biases.</p>
<p>Professional Responsibilities Owed to Other Jurisdictions</p> <p>Rule 8.5</p>	<p>A lawyer should analyze and understand the relevant laws and regulations of each jurisdiction in which a lawyer is licensed and make reasonable efforts to remain current with changes in those laws and regulations in order to comply with applicable professional obligations.</p>