

### Board of Trustees Directive to COPRAC:

I'm tasking the Committee on Professional Responsibility and Conduct, which is charged with studying and providing consultation and assistance to the Board on matters involving professional responsibility, with exploring the use of and recommendations for regulation of AI in the legal profession. By the Board's November 2023 Board meeting, COPRAC should issue its recommendations, which may include practical guidance, an advisory opinion or other resources, changes to the Rules of Professional Conduct or other rules or statutes, or other recommendations to ensure that AI is used competently and in compliance with the professional responsibility obligations of lawyers.

DG: Is there a benefit to considering new rule (or statute?) requiring a license or permit for those offering foundation models to practicing attorneys or pro se litigants to perform legal tasks/processes in the State of California?

DG: should there be a recommendation for a strong push for skill building and education of lawyers for generative AI in particular?

### GUIDELINES AND PRINCIPLES FOR USE OF GENERATIVE ARTIFICIAL INTELLIGENCE (AI or GAI) IN THE PRACTICE OF LAW

Notes on the scope and application of these guidelines and principles:

- These guidelines and principles largely identify and apply, but in some cases extend, existing professional responsibility obligations for lawyers, including the California Rules of Professional Conduct (RPC) and the State Bar Act set forth in the Business & Professions Code (B&P). They have been extended or applied to encompass the use of generative AI within legal practice.

DG: how about unauthorized practice of law (ie NOT "within the practice of law")?

- These principles are intended to address oncoming issues and imminent concerns with regard to the use of generative AI. However, they may apply to more traditional applications of AI and data modeling as well e.g. for automating processes or analytics.
- These principles are intended for use in governing best practices with the use of externally provided generative AI applications and services.

The below table lists the Guidelines and Principles (with corresponding existing applicable or related rules and statutes) that should be followed whenever using generative artificial intelligence in the legal profession. The second column presents an example of how an attorney or law firm might behave inconsistently with the principle, in which case ethical responsibilities may be violated. The last column provides an example of how to behave consistently with the principle, and applicable rule. Note that many of the applicable authorities create overlapping duties. As such, some of the conduct and compliance examples in columns two and three may be applicable to multiple sections.

RB to chart below: excellent! I read many examples to be repetitive/same across unique rules. suggest each rule example(s) be unique with the acknowledgment there is rule overlap by their nature

EKD to chart below: Clarify the differences between each rule so that the examples are more distinct for each rule.

<u>Guidelines and Principles (and corresponding rules)</u>	<u>Conduct that Breaches the Rule or Guideline</u>	<u>How to Comply with the Rule or Guideline</u>
----------------------------------------------------------------	--------------------------------------------------------	-----------------------------------------------------

<p><b>Duty of Confidentiality</b></p> <p>RPC 1.6 RPC 1.8.2 B&amp;P §6068(e)</p>	<p>You share confidential client information with a service provider by inputting data into the software or application you are using, when responding to prompts.</p> <p>If you are using a service that is not secure, the disclosure to the service provider in and of itself is a violation of the duty of confidentiality.</p> <p>NM: Shouldn't we also say that failing to diligence a particular platform's confidentiality rules is an issue? Not sure if this goes better in the duty of competence or diligence...</p> <p>EKD: Eleanor to draft revised language. Example from Dazza's Zoom Chat: Good: You share confidential information of your client with a service provider through prompts that violate your duty because, for example, the terms and conditions of the service provider permit them to share the information with third parties or use the prompts as part of training their models. Bad: Ensure you don't share confidential information in the first place, such as by adequately anonymizing the information in your prompts, or possibly ensure contractual and other safeguards are in place, including client consent.</p> <p>Also, if the terms and conditions of the service provider permit them to share the information with third parties, or use the prompts as part of their training models, this is another disclosure</p>	<p>Do not input any confidential information of the client into the AI software or applications which do not protect that information.</p> <p>NM: Again, I think we need to be clear that this should only apply to platforms that are not safe.</p> <p>Generally attempt to anonymize all client information you enter in response to the software's prompts.</p> <p>Even if not using the client's name or description, avoid entering details that can be used to identify the client.</p> <p>Obtain informed written consent before inputting any client information, and ensure that the contract with the AI provider does not permit the provider to share information you input into the software/application, to third parties or for their own use, in any manner.</p>

	<p>revealing your client's confidential information.</p> <p>The service provider may also automatically collect data that you input for their own use, which is also generated in their output to other customers.</p>	
<p><b>Duty of Reasonable Diligence</b></p> <p>RPC 1.3</p> <p>EKD: Erika will look discipline case history in CA in other jurisdictions as well as history of adoption of rule 1.3 in CA.</p> <p>BK: I would like to discuss incorporating the concept of "explainability" into Dilligence, Competence and the Duty to Supervise. I think that requiring "explainability" in the use of AI in law may go some distance to alleviating our concerns about over-reliance, and a potentially declining standard of practice resulting from AI use. At least it is worth discussing in the Working Group.</p>	<p>Not fact-checking or verifying citations from the outputs of generative AI.</p> <p>Using AI for legal analysis without regard for defensibility, <i>i.e.</i>, accepting analytical output from large-language models (LLMs) or similar generative AI technology without checking the accuracy (legal or factual) of the content.</p> <p>Accepting at face value the output of AI that contains recommendations or decisions contrary to the client's best interests, <i>e.g.</i>, potentially prioritizing the interests of a buyer when your client is the seller, or the employer when your client is the employee, etc.</p> <p>EKD: Nima will draft revised language to this and the competence section. From Working Group meeting: Add language such as "(<i>i.e.</i>, accepting the GAI output without reviewing for accuracy, applying the necessary analysis, particular client issues, etc.)"</p> <p>From Dazza's Zoom Chat:</p> <p>Good: Critically review, confirm, or correct the output of generative AI to ensure the content accurately reflects and supports the interests and</p>	<p>Ensure diligence and prudence with respect to facts and law. Maintain strong practices for fact-checking, and verifying the accuracy of legal citations, or contract references.</p> <p>Critically review, confirm, and correct the output of generative AI to ensure the content accurately reflects and supports the interests and priorities of the client in the matter at hand, including as part of advocacy for the client.</p> <p>Use AI-generated legal arguments, only as a springboard, and not as final product. For example, it can be used as a starting point, but you must review, analyze, supplement and improve upon it, before determining that the legal arguments you are presenting are the best arguments you can make for your client.</p> <p>Take steps to avoid lack of diligence, or reliance on AI to such a degree that it hinders critical attorney analysis fostered by traditional research and writing. For example, supplement any AI-generated research with research conducted by a human being; and supplement any AI-generated argument with critical</p>

	<p>priorities of the client in the matter at hand, including as part of advocacy for the client.</p> <p>Bad: Accepting at face value the output of GAI that contains recommendations or decisions contrary to the client's best interests, e.g. potentially prioritizing the interests of a buyer when your client is the seller, or the employer when your client is the employee, etc.</p>	<p>analysis and review of authorities by a human being.</p>
<p><b>Duty of Competence</b></p> <p>RPC 1.1</p>	<p>Making use of AI and accepting outputs as fact without understanding to a reasonable degree how the technology works and/or critically reviewing how outputs are generated.</p> <p>NM: I think we need to qualify this as I've suggested. We can't expect all practitioners to understand the tech in full.</p> <p>Failing to understand how and whether the AI platform in use exploits confidential client information, <i>i.e.</i>, in tandem with the duty of confidentiality is the obligation that the attorney understand the AI platform's terms of use, privacy policies, and other rules for collecting and using data.</p> <p>NM: This to me is really important. Part of the duty of competence must be that the lawyer know what they are using and how it impacts their client's confidences.</p>	<p>Understand and skillfully integrate generative AI with other relevant apps and tools in your workflow.</p> <p>Learn to compose prompts to generate high-quality outputs that augment and improve upon existing human expertise.</p> <p>For example, after a human has drafted a legal brief, you can be creative in composing prompts that use AI to serve as a devil's advocate to help you anticipate challenges to the legal arguments you plan to present.</p> <p>However, if you use AI to generate the brief initially, and you fail to understand the shortcomings of the AI, such as its limits or inaccuracies that can be produced by certain prompts, you might end up using faulty output, and incorporating those errors into a brief that is filed with the court.</p> <p>The duty of competence requires more than mere detection and elimination of false AI generated results. Competence requires the continuous application of legal reasoning and analysis regarding</p>

		<p>all of the potential options and impacts presented, including those that are included or omitted from or by AI tools. Competence requires vigilance against over-reliance on AI tools, and the active practice of law and application of trained judgment by the lawyer.</p> <p><b>BK: Does this capture our discussion in a useful way, or does it seem too vague/lofty?</b></p> <p>Attempting to understand at a reasonable degree the applicable terms of use and other policies governing the use and exploitation of client data by AI is necessary to satisfy the duty of competence.</p>
<p><b>Duty of Candor to the Tribunal; and Meritorious Claims and Contentions</b></p> <p>RPC 3.1 RPC 3.3</p> <p>Cal. Code Civ. Proc., §128.7 FRCP 11(b)(2), (3) and (4)</p>	<p>Filing a brief or memorandum of points and authorities that was generated by AI, without sufficiently reviewing the legal citations for accuracy, and without notifying the court that your brief was generated by AI software.</p> <p>Using AI-generated legal writing in a brief, without notifying the court which sections were AI-generated, resulting in possible misrepresentations to the court, by violating FRCP 11 or Cal. Code Civ. Proc., §128.7 representations</p> <p><b>NM: I think we should temper this entire section if others agree. I think it would be good to caveat by saying that this only violates the duty of candor if applicable local rules or other rules of procedure require such disclosures. Would seem to be difficult (and possibly</b></p>	<p>Be aware that AI software can, and has, fabricated legal authorities that do not exist.</p> <p>Review all legal arguments and citations to authority for accuracy, and correct any errors or misleading arguments generated by AI, before you submit the document to the court.</p> <p>Carefully review all documents and language produced by AI to ensure that no fabricated legal citations exist in the document, and to ensure that all citations are valid, existing legal authority which supports the proposition for which it has been cited.</p> <p>Carefully review all language generated by AI, for an accurate recitation of any facts in your case, or of the cases cited in the brief.</p>

	<p>unworkable) to mandate that ALL litigants include such disclosures without regard for local rules or the preferences of judges. That may be where this all ends up, but seems like we'd want the legislature to weigh in on that before going further.</p>	<p>If sections of your brief were generated by AI without substantive modification by a human being, include a footnote in each such section to indicate to the court which parts were generated by AI instead of a human being.</p>
<p><b>Duty of Candor, generally</b></p> <p>RPC 8.4(c) RPC 4.1(a)</p>	<p>Sending a demand letter to opposing counsel or an opposing party that was generated by AI software which contains fabricated legal citations due to a failure to check the legal contentions for accuracy.</p> <p>NM: See my comment above, which applies here too.</p>	<p>Carefully review all documents and language produced by AI to ensure that no fabricated legal citations exist in the document, and to ensure that all citations are valid, existing legal authority which supports the proposition for which it has been cited.</p> <p>NM: Also, I'm not sure this is additive to the last row, but happy to discuss. May consider jettisoning or combining this section with the last.</p>
<p><b>Duty to Inform the Client</b></p> <p>RPC 1.4(a)(2)</p>	<p>Failure to disclose use of AI where client would be surprised by and object to the way in which the attorney used generative AI for their case.</p> <p>If the terms of an agreement with a client require disclosure of this type and the attorney fails to disclose.</p>	<p>Clients pay for attorney analysis, not, AI automations. Disclosure to, and written consent from, the client should be obtained before use of generative AI in the client's case.</p> <p>RB: attorney methods being singled out seems odd</p> <p>NM: This seems slightly inconsistent with the conduct we're claiming violates the duty. I think the rule should be that failure to specify, ideally in an engagement letter, the risks and possibility that AI can be used on an engagement is the real issue. In other words, if an attorney is considering using generative AI for a case, that attorney should inform and advise the client of the risks, ideally receiving informed</p>

		<p>consent. However, I wouldn't REQUIRE written informed consent, but just note that the practice is encouraged. In all cases, however, I WOULD say it's required that the client be notified, preferably via an engagement letter that such use is possible and that there are risks associated with it. This should be, in my view, on a case-by-case basis, including that if a broader AI project is contemplated, the client is informed further of the risks and benefits commensurate with the size and scale of the project.</p> <p>The terms of the client engagement should describe the use of technology and specifically address the responsible use of generative AI. If needed, an amended engagement letter agreed upon with the client can address this issue.</p>
<p><b>Duty to Supervise Lawyers and Non-Lawyers (including non-human non-lawyers)</b></p> <p>RPC 5.1 RPC 5.3</p>	<p>Making use of AI applications without adequate best practices and human oversight, evaluation, and accountability mechanisms in place</p> <p>RB: Overall guiding principle for G/AI use</p>	<p>Check all language drafted by AI for accuracy, using authoritative legal sources by an accountable human being before submission to a Court.</p> <p>Designate responsible parties who decide on permitted uses of AI within your office or law firm. For example, determine, which cases AI can be used in and which cases lack client's consent to use AI. Also, determine which tasks AI can or cannot be used to perform.</p> <p>Implement policies that require a designated responsible party to</p>

		<p>sign off on AI use on a client by client, or matter by matter basis.</p> <p>Attorneys should first be trained on all the pitfalls of any AI brief-writing or legal argument producing technology before being allowed to use it. They should be shown how to properly use the AI technology of the firm, before being given access to the programs for use in producing legal arguments or drafting briefs. They should also be warned of the limits and errors that can be produced in certain uses.</p> <p>NM: This is important, and I think we should cross-add it to the duty of competence.</p>
<p><b>Professional Responsibilities Owed to Other Jurisdictions</b></p> <p>RPC 8.5</p> <p>RB: current examples?</p>	<p>Deploying AI to employees and agents of your firm who practice in jurisdictions that have, for example, banned that technology.</p>	<p>Analyze the relevant laws and regulations of each jurisdiction in which AI is deployed in/by your firm and ensure your compliance with such rules, e.g. by being able to turn off the tool for users in jurisdictions that limit or restrict the use of AI in the legal profession.</p>

#### Issues List for Scenarios:

- Tools for AI Use (deployment model)
  - Bespoke product in a closed system (e.g., Harvey) vs. general, open AI system (e.g., ChatGPT)
- Firm Size/Application/Use (usage type)
  - Use case
  - Privacy/ security / legal or regulatory issues (relate back to competence)
- Legal process (federal court, newsletter, etc.)

#### AI Scenario 1 – Small Firm



Attorney is the owner of a small law firm that primarily handles family law cases and offers free consultations. As part of each initial consultation, Attorney explains that the law firm uses generative artificial intelligence (GAI) to assist with its work

In addition to explaining how the firm uses GAI as part of the initial consultation, Law Firm's retention agreement contains a page that discloses the use of GAI and contains a space for the client's signature:

Law Firm uses generative AI (GAI), a product that generates outputs in response to instructions or requests that are input into the product. GAI generated outputs are based on an algorithmic model that is trained on vast amounts of data, including, but not limited to, all of the case law and other relevant authorities in any jurisdiction in the United States.

Law Firm uses GAI, for example, to (a) supplement human-generated research and writing, such as legal memos, motions, and briefs; and (b) assists Law Firm to defend the client's position by drafting a mock response brief on behalf of opposing counsel to help Law Firm anticipate and defend against attacks on the proposed arguments. Law Firm also uses GAI perform other work, including, but not limited to:

- Draft and review documents, settlement agreements, and proposed court orders.
- Draft communications for, and/or directly communicate with, clients, courts, and other parties through email, phone, and chat.
- Research relevant laws, precedents, and legal trends.
- Draft questions to use in depositions.

Law firm requires its attorneys to complete the Law Firm's formal training prior to an GAI use, and requires all Law Firm attorneys review any GAI output and independently verify the legal authorities cited in any GAI output through other sources, such as Westlaw, LexisNexis, and the firm's law library. Law Firm imposes **strict** measures to ensure that no confidential information that could be used to identify the client is disclosed when using GAI.

**RB: specify .. strict is subjective to interpretation**

Law Firm also uses other products that utilize artificial intelligence for calendar management, document management, document review, compliance review, and other uses. Law Firm inputs confidential information into these products. However, these products are walled-off to protect confidential information. No confidential information that is put into these programs can be viewed by anyone outside the law firm.

Client agrees that Law Firm may use generative artificial intelligence products and other products that utilize artificial intelligence in Law Firm's representation of client including, but not limited to, in the manners described herein and in other manners that comply with Law Firm's and its attorneys' ethical obligations.

---

*Client Signature and Date*

---

*Attorney Signature and Date*

Relevant Authorities:

Rules of Professional Conduct 1.1, 1.4, 1.6, 2.1, 5.1, 5.3

Business and Professions Code sections 6068, subdivision (e)

Other Rule to Consider that are Not Addressed by this Scenario (Question mark indicates request to discuss whether it's directly relevant to Generative AI):

Rules 1.5, 1.7, 1.8.2, 1.9, 3.3, 3.4?, 4.1?, 7.1-7.3?, 8.4(d), 8.4.1

Business and Professions Code sections 6068, subdivisions (c)?, (d)