



Judicial Council of California

455 Golden Gate Avenue · San Francisco, California 94102-3688

www.courts.ca.gov

REPORT TO THE JUDICIAL COUNCIL

Item No.: 25-109

For business meeting on July 18, 2025

Title

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work

Report Type

Action Required

Effective Date

September 1, 2025

Rules, Forms, Standards, or Statutes Affected

Adopt Cal. Rules of Court, rule 10.430;
adopt Cal. Standards of Judicial Administration, standard 10.80

Date of Report

June 16, 2025

Recommended by

Artificial Intelligence Task Force
Hon. Brad R. Hill, Chair

Contact

Jessica Devencenzi, 916-263-1374,
jessica.devencenzi@jud.ca.gov
Saskia Kim, 916-643-6951,
saskia.kim@jud.ca.gov
Jenny Grantz, 415-865-4394,
jenny.grantz@jud.ca.gov

Executive Summary

The Artificial Intelligence Task Force recommends adopting one rule of court and one standard of judicial administration to address the use of generative artificial intelligence for court-related work. The task force developed this proposal as part of its charge from the Chief Justice to oversee the development of policy recommendations on the use of artificial intelligence in the judicial branch. Adopting the proposed rule and standard will help promote responsible innovation in court operations while protecting confidential information, ensuring appropriate oversight, and maintaining public trust in the judicial branch.

Recommendation

The Artificial Intelligence Task Force recommends that the Judicial Council, effective September 1, 2025, adopt California Rules of Court, rule 10.430 and California Standards of Judicial Administration, standard 10.80, to address the use of generative artificial intelligence for court-related work. The proposed rule and standard are attached at pages 15–19.

Relevant Previous Council Action

The Chief Justice created the Artificial Intelligence Task Force in May 2024 in response to growing interest in generative artificial intelligence (generative AI) and public concern about the impact of this technology on the judicial branch. The task force is responsible for overseeing the development of policy recommendations on the use of AI in the judicial branch.

Analysis/Rationale

Generative AI is an emerging technology that can generate content in many forms and languages and on almost any subject at a user's request. Generative AI has many potential benefits and appears to have particular promise for courts' management and administrative functions. Generative AI also poses significant risks, though many of these risks can be mitigated with careful training and use, coupled with oversight.

The Artificial Intelligence Task Force is working to address the benefits and risks of generative AI throughout California's judicial branch. Use of generative AI for court-related work is one of the task force's current areas of focus. At the February 2025 Judicial Council meeting, the task force announced the *Model Policy for Use of Generative Artificial Intelligence* (model policy), which is offered as a resource for courts wishing to permit the use of generative AI for court-related work. The model policy addresses the confidentiality, privacy, bias, safety, and security risks posed by generative AI systems and addresses supervision, accountability, transparency, and compliance when using those systems. Courts can adopt the model policy as written or add, modify, or delete provisions as needed to address specific goals or operational requirements.

The task force recommends adopting a rule of court and a standard of judicial administration to address the confidentiality, privacy, bias, safety, and security risks posed by use of generative AI in court-related work. Generative AI is a tool that can be used to assist judicial officers and court staff to fairly administer justice, and this recommendation aims to promote responsible innovation in court operations while protecting confidential information, ensuring appropriate oversight, and maintaining public trust in the judicial branch.

Rule 10.430

Under rule 10.430, any court that does not prohibit the use of generative AI by court staff or judicial officers must adopt a policy that applies to the use of generative AI by court staff for any purpose and by judicial officers for any task outside their adjudicative role. The rule applies to the superior courts, the Courts of Appeal, and the Supreme Court. As discussed below, standard 10.80 covers the use of generative AI by judicial officers for tasks within their adjudicative role.¹

¹ Use of generative AI by Judicial Council staff will be covered by a separate policy, which is currently being developed by the Judicial Council Information Technology office.

Policies adopted under rule 10.430 must:

- Prohibit the entry of confidential, personal identifying, or other nonpublic information into a public generative AI system, meaning any system that is publicly available or that allows information submitted by users to be accessed by anyone other than judicial officers or court staff;
- Prohibit the use of generative AI to unlawfully discriminate against or disparately impact individuals or communities based on membership in certain groups, including any classification protected by federal or state law;
- Require court staff and judicial officers who create or use generative AI material to take reasonable steps to verify the accuracy of the material, and to take reasonable steps to correct erroneous or hallucinated output in any material used;
- Require court staff and judicial officers who create or use generative AI material to take reasonable steps to remove any biased, offensive, or harmful content in any material used;
- Require staff and judicial officers to disclose the use of or reliance on generative AI if the final version of a written, visual, or audio work provided to the public consists entirely of generative AI outputs; and
- Require compliance with all applicable laws, court policies, and ethical and professional conduct rules, codes, and policies when using generative AI.

Courts can comply with rule 10.430 by adopting the model policy or a policy that is substantially similar. The provisions marked “optional” in the model policy are not needed to comply with rule 10.430.

The task force considered several alternatives when drafting rule 10.430. First, the task force considered having the rule apply directly to court use of generative AI rather than requiring courts to implement policies meeting the rule’s requirements. Second, the task force considered requiring courts to adopt the model policy instead of giving courts the option to adopt their own policy. Third, the task force considered making the rule more expansive to include the model policy’s optional provisions.

The task force ultimately decided that the recommended version of rule 10.430 is preferable because it gives courts the flexibility to write a policy that will meet their specific goals and operational requirements while ensuring that all court policies address the major risks of generative AI. As discussed in the Advisory Committee Comment to subdivision (d), courts can comply with the rule by adopting a use policy that contains language substantially similar, but not identical, to subdivision (d). Courts can also adopt policies that are more restrictive than rule 10.430 or that have additional provisions not covered by the rule.

The task force also concluded that it will be beneficial to use the model policy to illustrate and expand on the rule's requirements rather than relying solely on a rule of court to set the parameters for court use of generative AI. The model policy can provide background, suggestions, examples, and other material that would not be suitable for a rule of court. The model policy can also be revised more quickly to respond to changes in generative AI technology and its uses.

Standard 10.80

Standard 10.80 covers the use of generative AI by judicial officers for tasks within their adjudicative role, and its provisions are similar to those in rule 10.430. The standard states that judicial officers:

- Should not enter confidential, personal identifying, or other nonpublic information into a public generative AI system;
- Should not use generative AI to unlawfully discriminate against or disparately impact individuals or communities based on membership in certain groups, including any classification protected by federal or state law;
- Should take reasonable steps to verify that generative AI material, including any material prepared on their behalf by others, is accurate, and should take reasonable steps to correct any erroneous or hallucinated output in any material used;
- Should take reasonable steps to remove any biased, offensive, or harmful content in any generative AI material used, including any material prepared on their behalf by others; and
- Should consider whether to disclose the use of generative AI if it is used to create content provided to the public.

Additionally, the Advisory Committee Comment to subdivision (b) reminds judicial officers to comply with applicable laws, court policies, and the California Code of Judicial Ethics when using generative AI.²

The task force considered having rule 10.430 cover the use of generative AI by judicial officers for any purpose but determined that a standard of judicial administration would be more appropriate for addressing the use of generative AI for tasks within a judicial officer's adjudicative role. The standard identifies the major risks of generative AI and allows judicial officers to determine the best way to address those risks in their adjudicative work.

² In particular, the task force anticipates likely future developments in ethical guidance relating to judicial officers' use of generative AI in their adjudicative work.

Policy implications

This proposal will create a rule of court and a standard of judicial administration to address the confidentiality, privacy, bias, safety, and security risks posed by use of generative AI in court-related work. Adopting the proposed rule and standard will help promote responsible innovation in court operations while protecting confidential information, ensuring appropriate oversight, and maintaining public trust in the judicial branch.

This proposal is, therefore, consistent with the *Strategic Plan for California's Judicial Branch*, specifically the goals of Independence and Accountability (Goal II) and Modernization of Management and Administration (Goal III).

Comments

This proposal was circulated for comment from March 13 to April 17, 2025, as part of a special invitation-to-comment cycle. The proposal received 19 comments: 3 from superior courts, 1 from the Joint Rules Subcommittee of the Trial Court Presiding Judges Advisory Committee and Court Executives Advisory Committee, 2 from judicial officers, 2 from law professors, 2 from attorney or bar associations, 1 from a legal aid organization, 1 jointly from a public interest association and a labor union, 1 from a public interest association, 2 from legal technology companies, 2 from attorneys, and 2 from non-attorneys. Two commenters agreed with the proposal, six agreed with the proposal if modified, five did not agree with the proposal, and six did not state a position. A chart with the full text of the comments received and the task force's responses is attached at pages 20–136. The principal comments and the task force's responses are summarized below.

Scope of the rule and standard

Many commenters suggested changing the scope of the rule and standard. These commenters primarily argued that the rule and standard should be more restrictive, but some commenters argued that the rule and standard should be made more permissive or should be abandoned entirely.

Six commenters suggested revising the rule or standard to completely prohibit the use of generative AI for adjudicative tasks such as writing opinions and orders. Similarly, one commenter suggested that provisions applying to use of generative AI by judicial officers for tasks within their adjudicative role should be mandatory rather than discretionary. One commenter suggested that courts should be prohibited from using generative AI for anything other than nonjudicial, public-facing applications such as streamlining access to public records.

Two comments from judicial officers suggested that the rule and standard should not be adopted at all. One argued that the rule and standard are not needed because existing rules, laws, and canons address the task force's core concerns regarding generative AI, while the other argued that the proposed rule should be replaced with a prohibition on the use of generative AI in adjudicative work and then be revised in the future to include additional provisions once technology, education, and guidance have developed further. This commenter also questioned the need to adopt a standard on the use of AI in adjudicative work now, before ethical guidance

is issued. Additionally, two commenters suggested that rather than allowing courts to develop their own generative AI use policies, the Judicial Council should adopt a uniform set of rules for generative AI use that would apply to all courts.

The task force is not recommending changes in response to these suggestions for several reasons. First, the task force recognizes that the possibility that generative AI might be used to draft orders, opinions, and other adjudicative materials raises significant concerns. However, the task force concluded that this issue is more appropriately addressed by the California Code of Judicial Ethics and related ethical guidance. For example, the canons prohibit judicial officers from abrogating their responsibility to personally decide the matters before them and considering evidence and facts that were not properly judicially noticed, and require judicial officers to “act at all times in a manner that promotes public confidence in the integrity and impartiality of the judiciary.”³ The task force concluded that the canons therefore likely prohibit judicial officers from having generative AI decide issues or write their opinions for them, and that the Supreme Court’s judicial ethics committees are the appropriate bodies to ask for guidance on this subject.

The task force also determined that regulating specific uses of generative AI is more difficult than it appears. For example, one commenter suggested prohibiting judicial officers from “us[ing] or rely[ing] on generative AI for any task that may affect the substance of an adjudicative decision,” but this language would prohibit judicial officers from using legal research tools developed by trusted legal research providers, or even using a grammar-checking tool that uses a generative AI model to make grammar suggestions. While it is likely uncontroversial to say that judicial officers should not prompt ChatGPT to decide issues or draft an opinion for them, it is less clear whether and to what extent it is acceptable for judicial officers to use generative AI tools for tasks like researching and outlining legal arguments. The task force therefore concluded that judicial officers should have the discretion to decide whether specific uses of generative AI are appropriate for adjudicative tasks, consistent with the California Code of Judicial Ethics and any applicable ethical guidance.

Second, the task force determined that courts and judicial officers are in the best position to identify acceptable uses of generative AI to meet their specific needs. The risks of generative AI depend heavily on the specific tool and how it is used. The acceptability of some uses (such as legal research) depends on the tool (such as a purpose-built legal research tool or an all-purpose chatbot), and the acceptability of some tools depends on how they are used (such as improving grammar in a single paragraph versus writing large sections of a document). It would be extremely difficult for the task force to create a list of acceptable tools and uses, and such a list would likely be both under- and overinclusive because the task force would have to speculate about how specific tools work or how courts might use them. Additionally, putting such a list in a rule of court would make it difficult to keep up with technological advancements. For these

³ Cal. Code Jud. Ethics, canons 2A, 3B(7).

reasons, the task force recommends that the rule and standard address specific risks of generative AI rather than specific generative AI tools or uses.

Third, the task force acknowledges that creating rules and standards to address emerging technology is challenging because those rules might become outdated, restrict innovation, or inadvertently exclude additional technologies posing the same risks the rules are meant to address. This concern was a significant factor in the task force's decision to recommend a rule and standard that focus on the overarching risks of generative AI, such as confidentiality, bias, accuracy, and transparency, rather than attempting to allow or prohibit specific generative AI tools and uses. The task force concluded that these risks must be addressed regardless of how the technology develops in the future and that it is therefore appropriate to recommend adoption of the rule and standard.

Finally, the current proposal strikes the best balance between uniformity and flexibility. Rule 10.430 will require all courts that do not prohibit the use of generative AI by court staff or judicial officers to impose specific requirements addressing its primary risks. However, the task force recognizes that use of generative AI will look very different depending on the court. For example, some courts might only permit use of things like purpose-built legal research tools and grammar checkers, while other courts might develop generative AI systems for internal uses such as answering questions about the court's human resources policies. Similarly, courts have differing levels of information technology staff and resources and will therefore have different answers to questions about how generative AI tools should be approved, deployed, and supervised. The task force determined that each court is in the best position to determine how it can meet rule 10.430's requirements and whether its generative AI use policy should be more restrictive or detailed than the rule.

Applicability of rule 10.430

Judge Karnow of the Superior Court of San Francisco County noted that as proposed in the invitation to comment, rule 10.430 would not apply to courts that are silent on the use of generative AI for court-related work because they neither permit nor prohibit its use. Instead, the rule requires courts to adopt a use policy only if they choose to permit generative AI use. The concern is that this leaves an unintended vacuum where the rule would not apply, yet use of generative AI for court-related work might still be occurring without the safeguards and protections contained in the rule.

Because this was not the intent of the proposal, the task force agrees with this concern and recommends that rule 10.430(b) read as follows: "Any court that does not prohibit the use of generative AI by court staff or judicial officers must adopt a generative AI use policy. This rule applies to the superior courts, the Courts of Appeal, and the Supreme Court."

The task force will provide a model policy courts can use if they wish to prohibit the use of generative artificial intelligence for court-related work.

Provisions requiring disclosure of the use of generative AI

Twelve commenters suggested revising rule 10.430(d)(5) and standard 10.80(b)(5), which concern disclosure of the use of generative AI in court-related work. As proposed in the invitation to comment, the rule would apply to court staff using generative AI for any purpose and to judicial officers using generative AI for tasks outside their adjudicative role and would require disclosure if generative AI outputs constitute a substantial portion of the content used in the final version of a written or visual work provided to the public. The standard states that judicial officers using generative AI for tasks within their adjudicative role should consider whether to disclose the use of generative AI if it is used to create content provided to the public.

Several commenters suggested that rule 10.430(d)(5) should require disclosure of any use of generative AI, including when generative AI outputs do not constitute a substantial portion of the work and when generative AI is used or relied upon only to develop drafts but not the final work. Similarly, several commenters suggested that standard 10.80(b)(5) should be mandatory rather than discretionary, and that it should require disclosure if judicial officers use generative AI to any extent in the creation of any document or statement shared with the public.

Conversely, several commenters suggested that rule 10.430(d)(5) should be made discretionary rather than mandatory, or that a disclosure requirement might be unnecessary in some or all circumstances. For example, a superior court suggested that, provided other mandatory safeguards were in place, requiring disclosure might sometimes “impose an undue administrative burden without meaningfully enhancing public trust or accountability and would likely discourage the use of generative AI in instances when it is appropriate.”

Many commenters also suggested revising rule 10.430(d)(5) to define the term “substantial portion.” The commenters were concerned that without a definition or examples of what it means to constitute a “substantial portion” of a work, it will be difficult for judicial officers and court staff to determine whether disclosure is required. For example, one commenter asked whether “substantial” is a percentage of the final text or a material part of the final text (even if only a small percentage of that final text). Several commenters also suggested revising the model policy to provide examples and further explanation of when disclosure is required.

The task force considered the approaches suggested by the commenters and ultimately determined that the disclosure requirement in rule 10.430 should be revised. The task force recommends that rule 10.430(d)(5) read as follows:

Require disclosure of the use of or reliance on generative AI if the final version of a written, visual, or audio work provided to the public consists entirely of generative AI outputs. Disclosure must be made through a clear and understandable label, watermark, or statement that describes how generative AI was used and identifies the system used.

This disclosure requirement will cover things like generative AI chatbots because it is important to inform the public when they are interacting with AI and not a person, for example, as well as

in other circumstances where the final content provided to the public has not been created or edited by a person.

The task force considered keeping the previously proposed language and defining “substantial portion” but concluded that doing so would be premature because courts are only beginning to identify potentially beneficial uses of generative AI. At this stage in the development and use of generative AI, it is important to require courts to disclose when works provided to the public consist entirely of generative AI outputs. This enhances transparency and public confidence. It is much less clear whether and to what extent disclosure is necessary or helpful when generative AI is used as an assistive tool similar to non-AI tools that are already in use, such as when it is used for legal research or editing documents written by humans.

The task force is concerned that if the disclosure requirement is too broad, it is likely to sweep in uses of generative AI where disclosure might not convey meaningful information about the quality of the resulting work and might cause unjustified mistrust. As some commenters noted, generative AI disclosures may be seen as a signal that the resulting material is inherently unreliable or that humans were not involved in creating the material, even if the disclosure explains otherwise. For example, if a staff attorney uses a generative AI legal research tool from a trusted legal research provider to perform legal research, the use of generative AI might not pose any more risk than using a non-generative AI tool from the same provider. The task force is therefore concerned that a broad mandatory disclosure requirement could discourage use of generative AI tools, even for acceptable purposes, and that mandatory disclosure would not be an effective way to address concerns about the reliability or trustworthiness of generative AI outputs in many circumstances.

For these reasons, the task force plans to continue gathering information about how courts are using, or plan to use, generative AI and will determine whether a different disclosure requirement is needed once it is clearer how such a requirement will impact courts and the public and further enhance public trust and confidence.⁴

The task force is not recommending revisions to standard 10.80(b)(5). As explained in the discussion of the scope of the rule and standard, above, the task force determined that the question of whether and to what extent judicial officers may use generative AI for adjudicative tasks should be addressed by the California Code of Judicial Ethics and related ethical guidance.

The task force acknowledges that having different disclosure requirements depending on whether generative AI is used by a judicial officer for a task within their adjudicative role could create difficulties for courts, for example by making it difficult to determine whether the use of generative AI to create adjudicative material must be disclosed if a staff attorney wrote some of

⁴ As noted above, the Advisory Committee Comment to subdivision (d) of rule 10.430 provides that courts adopting a generative AI use policy may make their policy more restrictive than the rule requires and may include provisions not covered by the rule. This means, for example, that a court may impose broader disclosure requirements than the requirements contained in the rule.

the material. However, because the task force does not believe the rule and standard should set mandatory requirements for judicial officers using generative AI within their adjudicative role, harmonizing the requirements in the rule and standard would require making the rule's requirement discretionary as well. The task force determined that it is preferable to require mandatory disclosure in some circumstances and that the more limited scope of the mandatory disclosure requirement in rule 10.430 will minimize the circumstances in which the rule and standard will come into conflict.

Finally, the task force anticipates that the other requirements in the rule, such as the requirements to address bias in generative AI materials and to take reasonable steps to verify accuracy, as well as existing ethical rules for judicial officers and attorneys, address the most significant risks of generative AI.

Definitions of “artificial intelligence,” “generative AI,” and “public AI system”

Several commenters suggested revising rule 10.430(a) and standard 10.80(a) to clarify the definitions of “artificial intelligence,” “generative AI,” and “public AI system.” The commenters were concerned that the proposed definitions might make it difficult for judicial officers and court staff to determine whether a particular tool or system is covered by standard 10.80 or a use policy adopted to comply with rule 10.430. Two commenters suggested that the rule and standard should use existing or proposed statutory definitions, such as those in Civil Code section 3110.

The task force agrees that all three definitions should be revised in both rule 10.430(a) and standard 10.80(a). The task force recommends using the following definition of “generative artificial intelligence,” which is based on one proposed by Loyola Law School professor Rebecca Delfino:

“Generative artificial intelligence” means a computer-based system that uses machine learning or similar techniques to produce new content—such as text, images, audio, video, code, or data visualizations—in response to user inputs. Generative AI systems create content that is not pre-programmed or explicitly retrieved but synthesized based on underlying models trained on large datasets and may include integration with other sources, such as real-time access to proprietary databases.⁵

The task force also recommends deleting the definition of “artificial intelligence” from the rule and standard because it is no longer needed; the definition of “generative artificial intelligence” no longer refers to “artificial intelligence,” and that term is not used elsewhere in the rule or standard.

⁵ The definition proposed by Professor Delfino referred to “integration with real-time or domain-specific sources.” The task force is concerned that these terms might be confusing to laypeople and has therefore revised the definition to refer to “other sources, such as real-time access to proprietary databases.”

The task force concluded this revised definition of generative AI will more accurately describe existing and potential future generative AI systems. Additionally, although the task force agrees that consistency with statutory definitions can be beneficial in some circumstances, the existing statutory definitions, such as those in Civil Code section 3110, are not a good fit for the rule and standard. Those definitions are part of statutory schemes for regulating AI providers and use terminology that laypeople might find confusing, such as the reference in section 3110(a) to “explicit and implicit objectives.”

The task force recommends changing the term “public AI system” to “public generative AI system” and revising the definition as follows, based on a suggestion from the Superior Court of San Francisco County:

“Public generative AI system” means a generative AI system that allows anyone other than court staff or judicial officers to access the data that courts input or upload to the system, or to use that data to train AI systems. “Public generative AI system” does not include any system that the court creates or manages, such as a generative AI system created for internal court use, or any court-operated system the court uses to provide those outside the court with access to court data, such as a court-operated chatbot that answers questions about court services.

The task force concluded that this definition will make it clearer which generative AI systems are covered by rule 10.430(d)(1) and standard 10.80(b)(1), which prohibit entering nonpublic information into public generative AI systems. In particular, this definition will make it easier to understand what types of security features and user data policies to look for when determining whether nonpublic information can be entered into a particular generative AI system.

Definition of “adjudicative role”

The Joint Rules Subcommittee of the Trial Court Presiding Judges Advisory Committee and the Court Executives Advisory Committee suggested revising the rule and standard to define the term “adjudicative role” so that it is clearer which tasks are covered by the rule and which are covered by the standard.

Although the task force agrees that the term “adjudicative role” may seem vague, the task force is not recommending revisions in response to this suggestion. The task force concluded it is appropriate to leave the term undefined to avoid potential conflicts with other rules or guidance that use similar terms.⁶ Additionally, judicial officers have discretion and are best situated to determine whether a particular task falls within their adjudicative role.

⁶ For example, the task force notes that canon 2A of the California Code of Judicial Ethics refers to “performance of the adjudicative duties of the judicial office” but does not define “adjudicative duties.”

Provisions addressing confidentiality

The Superior Court of San Francisco County suggested revising rule 10.430(d)(1), which prohibits entry of nonpublic information into public generative AI systems, to replace the word “nonpublic” with “nondisclosable” because nonpublic information is not necessarily confidential.

The task force recommends that rule 10.430(d)(1) continue to refer to “nonpublic” information. The task force expects that inputting nonpublic information into public generative AI systems may be problematic even if the nonpublic information is not confidential. For example, public generative AI systems can potentially be trained on any information users include in prompts or upload to the system, and information given to the system by one user can potentially appear in responses the system provides to other users.

Provisions addressing bias and discrimination

The Joint Rules Subcommittee suggested removing the word “unlawfully” from rule 10.430(d)(2) because it is “redundant and unnecessary.”

The task force recommends that rule 10.430(d)(2) refer to unlawful discrimination. The task force acknowledges that it is not strictly necessary to prohibit the use of generative AI to unlawfully discriminate because such discrimination is already prohibited. The task force included this provision in the rule primarily to ensure courts are aware of the risk that generative AI systems can produce biased or discriminatory outputs.

Provisions requiring review of generative AI material

Mark Griffin, the interim chair of the California Lawyers Association’s Law Practice Management and Technology Section, suggested requiring court staff and judicial officers to review their generative AI prompts (the user’s queries or inputs into the generative AI system) for bias.

The task force is not recommending revisions in response to this suggestion. Although the task force agrees that court staff and judicial officers should be aware that biased prompts can lead to biased outputs, the task force is concerned that requiring prompts to be unbiased could make it more difficult for judicial officers and court staff to perform certain tasks. For example, a staff attorney conducting legal research for a case alleging bias or discrimination might have to write prompts describing biased or discriminatory language. The task force anticipates that this issue can more appropriately be addressed through education and guidance materials.

Separately, two commenters suggested revising the rule and standard to clarify that those who use generative AI are responsible not only for reviewing generative AI material for accuracy, completeness, and bias, but also for verifying and correcting any material that contains inaccurate, incomplete, or biased content. As proposed in the invitation to comment, the rule and standard could be read to require people to *review* their generative AI material without requiring them to *correct* the material.

The task force agrees with this concern and recommends that rule 10.430(d)(3) read as follows: “Require court staff and judicial officers who create or use generative AI material to take reasonable steps to verify that the material is accurate, and to take reasonable steps to correct any erroneous or hallucinated output in any material used.”

The task force recommends that rule 10.430(d)(4) read as follows: “Require court staff and judicial officers who create or use generative AI material to take reasonable steps to remove any biased, offensive, or harmful content in any material used.”

The task force recommends similar revisions to standard 10.80(b)(3) and (4).

The task force recommends using the phrase “take reasonable steps” so that it will be possible to determine whether judicial officers and court staff have complied with the rule. The task force recommends removing the words “complete” and “incomplete” from rule 10.430(d)(3) for similar reasons.

Suggestions for additional provisions

Commenters suggested revising the rule and standard to cover additional topics, such as education and training requirements, benchmarking and documentation requirements, annual policy reviews, recordkeeping, and procurement. The task force appreciates these suggestions and is carefully considering them. The task force did not include these suggestions in its current recommendation because they cover topics that were not expressly covered by the invitation to comment.

Similarly, the task force appreciates the commenters’ suggestions for the model policy. The task force will be updating the model policy to conform with changes made to the rule and standard in response to the public comments (such as changes to defined terms) and will consider whether and how to make the commenters’ suggested revisions to the policy.

The task force will continue to consider how to address the risks posed by use of generative AI in court-related work. In addition to the proposed rule and standard and the model policy, the task force has developed a list of frequently asked questions and is considering whether additional guidance documents are needed. The task force will work with the Center for Judicial Education and Resources to ensure that judicial officers and court staff receive education and training regarding generative AI, including on emerging uses and risks. The task force has also been monitoring policy and other developments in jurisdictions outside California and will continue to do so.

Alternatives considered

The task force considered not recommending adoption of a rule or standard but ultimately determined that the proposal was warranted because it sets uniform requirements for courts that do not prohibit the use of generative AI for court-related work and because it helps courts, judicial officers, and court staff identify and address the primary risks of generative AI when used for court-related work. As discussed above, the task force considered several alternatives

when drafting the proposed rule and standard and in response to the public comments. The task force concluded that the current recommendation strikes the best balance between addressing the major risks of generative AI and giving courts the flexibility to address those risks in a way that will meet their specific goals and operational requirements.

Fiscal and Operational Impacts

Adopting rule 10.430 will require any court that does not prohibit the use of generative AI by court staff or judicial officers to adopt a generative AI use policy, which in turn might require training for judicial officers and court staff. Adopting standard 10.80 might also require training for judicial officers. The rule and standard in this proposal do not require courts to permit use of generative AI and therefore do not require courts to incur costs related to the purchase or use of generative AI tools.

The Superior Court of Los Angeles County commented that the proposed effective date of September 1, 2025, might not give courts enough time to implement tools to enforce their generative AI use policies. The court suggested a six-month implementation timeline. The court also noted that this proposal might be more difficult to implement for larger courts than for smaller courts due to their level of development and use of AI-related tools and applications. Similarly, the Superior Court of Placer County commented that a two-month implementation period is sufficient, but it might be helpful to provide a six-month grace period for compliance with adopted policies to allow courts time to coordinate with their vendors.

In response to these comments, the task force recommends revising rule 10.430(b) to state that courts that do not prohibit the use of generative AI must adopt a use policy by December 15, 2025. The rule will still become effective on September 1 if approved by the Judicial Council, but courts will have more time to create use policies and any tools necessary to implement those policies.

Attachments and Links

1. Cal. Rules of Court, rule 10.430, at pages 15–17
2. Cal. Standards of Judicial Administration, standard 10.80, at pages 18–19
3. Chart of comments, at pages 20–136

Rule 10.430 of the California Rules of Court is adopted, effective September 1, 2025, to read:

Title 10. Judicial Administration Rules

Division 2. Administration of the Judicial Branch

Chapter 6. Court Technology, Information, and Automation

Rule 10.430. Generative artificial intelligence use policies

(a) Definitions

As used in this rule, the following definitions apply:

- (1) “Court staff” means all employees, contractors, volunteers, and any other persons working for or on behalf of the court.
- (2) “Generative artificial intelligence” or “generative AI” means a computer-based system that uses machine learning or similar techniques to produce new content—such as text, images, audio, video, code, or data visualizations—in response to user inputs. Generative AI systems create content that is not pre-programmed or explicitly retrieved but synthesized based on underlying models trained on large datasets and may include integration with other sources, such as real-time access to proprietary databases.
- (3) “Judicial officer” means all judges of the superior courts, all justices of the Courts of Appeal and the Supreme Court, all temporary and assigned judges, and all subordinate judicial officers.
- (4) “Public generative AI system” means a generative AI system that allows anyone other than court staff or judicial officers to access the data that courts input or upload to the system or to use that data to train AI systems. “Public generative AI system” does not include any system that the court creates or manages, such as a generative AI system created for internal court use, or any court-operated system the court uses to provide those outside the court with access to court data, such as a court-operated chatbot that answers questions about court services.

(b) Generative AI use policies

Any court that does not prohibit the use of generative AI by court staff or judicial officers must adopt a generative AI use policy by December 15, 2025. This rule applies to the superior courts, the Courts of Appeal, and the Supreme Court.

1 **(c) Policy scope**

2
3 A use policy created to comply with this rule must cover the use of generative AI
4 by court staff for any purpose and by judicial officers for any task outside their
5 adjudicative role.

6
7 **(d) Policy requirements**

8
9 Each court's generative AI use policy must:

- 10
11 (1) Prohibit the entry of confidential, personal identifying, or other nonpublic
12 information into a public generative AI system. Personal identifying
13 information includes driver's license numbers; dates of birth; Social Security
14 numbers; National Crime Information and Criminal Identification and
15 Information numbers; addresses and phone numbers of parties, victims,
16 witnesses, and court personnel; medical or psychiatric information; financial
17 information; account numbers; and any other content sealed by court order or
18 deemed confidential by court rule or statute.
- 19
20 (2) Prohibit the use of generative AI to unlawfully discriminate against or
21 disparately impact individuals or communities based on age, ancestry, color,
22 ethnicity, gender, gender expression, gender identity, genetic information,
23 marital status, medical condition, military or veteran status, national origin,
24 physical or mental disability, political affiliation, race, religion, sex, sexual
25 orientation, socioeconomic status, and any other classification protected by
26 federal or state law.
- 27
28 (3) Require court staff and judicial officers who create or use generative AI
29 material to take reasonable steps to verify that the material is accurate, and to
30 take reasonable steps to correct any erroneous or hallucinated output in any
31 material used.
- 32
33 (4) Require court staff and judicial officers who create or use generative AI
34 material to take reasonable steps to remove any biased, offensive, or harmful
35 content in any material used.
- 36
37 (5) Require disclosure of the use of or reliance on generative AI if the final
38 version of a written, visual, or audio work provided to the public consists
39 entirely of generative AI outputs. Disclosure must be made through a clear
40 and understandable label, watermark, or statement that describes how
41 generative AI was used and identifies the system used.
42

- (6) Require compliance with all applicable laws, court policies, and ethical and professional conduct rules, codes, and policies when using generative AI.

Advisory Committee Comment

Subdivision (a). The definition of “court staff” in this subdivision is intended for use in this rule only.

Subdivision (c). California Standards of Judicial Administration, standard 10.80 covers the use of generative AI by judicial officers for any task within their adjudicative role.

Subdivision (d). This subdivision does not require any court to permit the use of generative AI by court staff or judicial officers. Courts may entirely prohibit the use of generative AI and may also set restrictions on how generative AI may be used for court-related work, such as allowing or prohibiting the use of specific generative AI tools, allowing use of generative AI only for particular tasks, or requiring approval for the use of generative AI. Courts that are required by subdivision (b) to adopt a use policy because they are not prohibiting the use of generative AI for court-related work can comply with subdivision (d) by adopting verbatim the nonoptional sections of the *Model Policy for Use of Generative Artificial Intelligence*, or by adopting a policy that uses substantially similar language. Courts adopting a generative AI use policy under this rule may make their policy more restrictive than the rule requires and may include provisions not covered by rule 10.430.

Standard 10.80 of the California Standards of Judicial Administration is adopted, effective September 1, 2025, to read:

Title 10. Standards for Judicial Administration

Standard 10.80. Use of generative artificial intelligence by judicial officers

(a) Definitions

As used in this standard, the following definitions apply:

- (1) “Court staff” means all employees, contractors, volunteers, and any other persons working for or on behalf of the court.
- (2) “Generative artificial intelligence” or “generative AI” means a computer-based system that uses machine learning or similar techniques to produce new content—such as text, images, audio, video, code, or data visualizations—in response to user inputs. Generative AI systems create content that is not pre-programmed or explicitly retrieved but synthesized based on underlying models trained on large datasets and may include integration with other sources, such as real-time access to proprietary databases.
- (3) “Judicial officer” means all judges of the superior courts, all justices of the Courts of Appeal and the Supreme Court, all temporary and assigned judges, and all subordinate judicial officers.
- (4) “Public generative AI system” means a generative AI system that allows anyone other than court staff or judicial officers to access the data that courts input or upload to the system or to use that data to train AI systems. “Public generative AI system” does not include any system that the court creates or manages, such as a generative AI system created for internal court use, or any court-operated system the court uses to provide those outside the court with access to court data, such as a court-operated chatbot that answers questions about court services.

(b) Use of generative artificial intelligence

A judicial officer using generative AI for any task within their adjudicative role:

- (1) Should not enter confidential, personal identifying, or other nonpublic information into a public generative AI system. Personal identifying information includes driver’s license numbers; dates of birth; Social Security numbers; National Crime Information and Criminal Identification and Information numbers; addresses and phone numbers of parties, victims, witnesses, and court personnel; medical or psychiatric information; financial

1 information; account numbers; and any other content sealed by court order or
2 deemed confidential by court rule or statute.

- 3
- 4 (2) Should not use generative AI to unlawfully discriminate against or
5 disparately impact individuals or communities based on age, ancestry, color,
6 ethnicity, gender, gender expression, gender identity, genetic information,
7 marital status, medical condition, military or veteran status, national origin,
8 physical or mental disability, political affiliation, race, religion, sex, sexual
9 orientation, socioeconomic status, and any other classification protected by
10 federal or state law.
- 11
- 12 (3) Should take reasonable steps to verify that generative AI material, including
13 any material prepared on their behalf by others, is accurate, and should take
14 reasonable steps to correct any erroneous or hallucinated output in any
15 material used.
- 16
- 17 (4) Should take reasonable steps to remove any biased, offensive, or harmful
18 content in any generative AI material used, including any material prepared
19 on their behalf by others.
- 20
- 21 (5) Should consider whether to disclose the use of generative AI if it is used to
22 create content provided to the public.
- 23
- 24

25 **Advisory Committee Comment**

26

27 **Subdivision (a).** The definition of “court staff” in this subdivision is intended for use in this
28 standard only.

29

30 **Subdivision (b).** This subdivision provides guidelines to judicial officers for the use of generative
31 AI for tasks within their adjudicative role. California Rules of Court, rule 10.430 covers the use
32 of generative AI by judicial officers for tasks outside their adjudicative role. In addition to the
33 guidelines provided in this subdivision, judicial officers should be mindful of complying with all
34 applicable laws, court policies, and the California Code of Judicial Ethics when using generative
35 AI.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
1.	Hon. Lamar Baker Associate Justice Court of Appeal, Second Appellate District, Division Five	N	The Artificial Intelligence Task Force is due great thanks for even taking on this difficult and complex issue. The proposed standard, however, does not reflect the humility and caution that is required under the circumstances. I am concerned the task force believed it was obligated to develop a standard specifying the conditions under which judges can use AI in adjudicating cases rather than considering a more fundamental question: whether courts, in the immediate future, should make any use of AI at all when deciding cases (other than, perhaps, as incorporated by the legal research functions of Lexis and Westlaw) and the degree to which such use will seriously undermine public confidence in the judiciary. According to the proposal memo, standard 10.80 “covers the use of generative AI by judicial officers for tasks within their adjudicative role.” As I read it, the standard would permit an appellate judge to upload the appellate briefs and record in an appeal to a generative AI system or program (so long as the program is not public or the briefs and record do not include confidential or nonpublic information), ask the AI program to draft an opinion resolving the appeal, and file that AI-drafted opinion as the opinion of the court without informing the parties (and, it appears, even the other judges on the appellate panel) of the use of AI--so long as the authoring judge reads the opinion before filing it. In my view, sanctioning such a scenario is a mistake and will undermine public confidence in the judiciary by the standard’s mere promulgation. Insofar as the task force believes the risks are mitigated by its anticipation, reflected in the proposal memo, of “likely future developments in ethical guidance relating to judicial officers’ use of generative AI in their adjudicative work,” I do not understand the need to promulgate a standard authorizing use of AI in adjudicative work now, before such ethical guidance issues. The two should, at a minimum, go hand in hand. AI is a very fast-moving field, but court policy need not, and should not, try to match that speed--and certainly not in a manner that might authorize what is later determined to be ethically questionable conduct.	The task force recommends adopting standard 10.80 because it has determined that the question of whether and to what extent judicial officers may use generative AI to carry out their adjudicative duties is more appropriately addressed by the California Code of Judicial Ethics and related ethical guidance. The canons likely prohibit judicial officers from having generative AI write their opinions for them, and the Supreme Court’s judicial ethics committees are the appropriate bodies to ask for guidance on this subject. The task force therefore concluded that it should not recommend that the Judicial Council either permit or prohibit the use of generative AI by judicial officers. Instead, the task force recognizes that some judicial officers may choose to use generative AI tools for tasks within their adjudicative role, and it therefore recommends adopting standard 10.80 to provide guidance regarding the risks of those tools.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			I accordingly recommend moving more slowly before authorizing changes that could rather dramatically change how courts decide cases. In my view, the only standard that needs to be promulgated now with respect to use of AI in a court's adjudicative work is a standard that says do not use it (except, perhaps, as incorporated by the legal research functions of Lexis and Westlaw). There is no rush. We have adequate time to continue with our traditional method of resolving cases while observing how AI develops and taking a more incremental approach to questions about the extent to which AI programs should be a part of the adjudicative process.	Additionally, the task force determined that it is necessary to recommend adoption of a rule and standard to address the risks of generative AI because generative AI can be used for tasks outside the adjudicative role and can likely be safely used for some adjudicative tasks, such as legal research using purpose-built legal research tools from trusted providers. Additionally, generative AI is increasingly being incorporated into existing software products and may already be difficult to avoid in some circumstances. The task force will continue to consider how to address the risks posed by use of generative AI in court-related work. In addition to the proposed rule and standard, the task force has developed FAQs and is considering whether additional guidance documents are needed. The task force will work with the Center for Judicial Education & Resources (CJER) to ensure that judicial officers and court staff receive education and training regarding generative AI,

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
				including on emerging uses and risks.
			If the task force is unpersuaded and believes it is for some reason imperative to now allow judges to use AI when resolving disputes, I would at a minimum recommend doing so on a far more limited basis, akin to a small pilot program, with much greater public transparency about what courts involved in that program are doing (and not doing).	The task force appreciates the commenter's concern but concluded that individual courts are in the best position to determine which uses are appropriate for their specific needs and circumstances. The task force is concerned that placing branchwide limitations on specific uses of generative AI will unnecessarily limit innovation and will prevent courts from identifying safe, effective uses of generative AI that do not pose ethical risks.
2.	Susan J. Bassi Publisher, Investigative Journalist Public Records & Local News Advocate Los Gatos	N	As a Silicon Valley resident and member of the media covering California's courts and technology, I respectfully submit this public comment with respect to SP25-01, with strong opposition to the proposed adoption of laws, rules, or policies permitting the use of artificial intelligence (AI) by court staff or judicial officers in California's judicial system. For over a decade, our team of investigative reporters has closely monitored and reported on California's courts, particularly the family court system, where there is no jury oversight, and where the press is largely absent. In 2024, the California Commission on Judicial Performance reported that family court judges now account for the highest number of complaints filed against judicial officers, marking a troubling milestone since the agency began keeping records. These complaints followed the implementation of technology in the courts that has	Please see responses to Susan Bassi's specific suggestions below.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			seen troubling patterns that would only worsen if judges and court staff were permitted to use AI, even with so- called proposed training. Our work has included the <i>Tainted Trials, Tarnished Headlines, Stolen Justice</i> investigative series, published in the Davis Vanguard, which exposed secret and undocumented meetings involving judges, prosecutors, family law attorneys, custody evaluators, and journalists from hedge fund–owned media outlets. These meetings were unrecorded, unregulated, and lacked transparency, raising serious concerns about bias and backchannel influence in legal proceedings. Our reporting has also highlighted: • Failures in the public disclosure of judicial conflicts, supported by requests made under California Rules of Court, Rule 10.500. • Mismanagement of courtroom technology, particularly in the Silicon Valley region where courts lag behind the private sector in tech competency. • Judicial misuse of social media during elections and politically sensitive cases, further indicating a lack of preparedness to ethically integrate emerging technologies like AI into court operations.	
			Why AI Has No Place in California Courtrooms—Yet Despite claims that AI could improve efficiency, multiple industry-wide studies and news investigations reveal that AI introduces serious risks, especially in contexts where human rights and liberties are at stake. These include: • Bias and Discrimination: The Stanford HAI Center and MIT Technology Review have documented how AI models reflect racial, gender, and socioeconomic biases in legal and hiring decisions. This has resulted in	The task force determined that it is necessary to recommend adoption of a rule of court and a standard of judicial administration to address the confidentiality, privacy, bias, safety, and security risks posed by use of generative AI in court-related work. The task force concluded that adopting the

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			discriminatory outcomes, something California’s judiciary must avoid at all costs. • Lack of Transparency: AI systems, especially those driven by large language models, operate as “black boxes.” The Harvard Berkman Klein Center warns that their logic and decision-making cannot be meaningfully audited or explained. This is incompatible with the requirement for transparency and judicial reasoning in constitutional law. • Inaccuracies and Hallucinations: Courts rely on facts, evidence, and precedent. However, AI models have a well-documented tendency to fabricate legal citations and misstate facts, a phenomenon known as “hallucination.” Several attorneys have already been sanctioned for submitting AI-generated briefs containing fictitious case law (New York Times, 2023).	proposed rule and standard will help promote responsible innovation in court operations while protecting confidential information, ensuring appropriate oversight, and maintaining public trust in the judicial branch. Additionally, the task force will continue to consider how to address the risks posed by use of generative AI in court-related work. In addition to the proposed rule and standard, the task force has developed FAQs and is considering whether additional guidance documents are needed. The task force will work with the Center for Judicial Education & Resources (CJER) to ensure that judicial officers and court staff receive education and training regarding generative AI, including on emerging uses and risks.
			Constitutional and Ethical Implications The role of a judge is not merely clerical. It is interpretive, ethical, and deeply human. Judicial discretion involves empathy, context, and constitutional analysis, elements that AI cannot currently replicate. Further, California Government Code § 69957 mandates that a verbatim transcript be provided in certain cases (e.g., family	The task force determined that the question of whether and to what extent judicial officers may use generative AI to carry out their adjudicative duties is more appropriately addressed by the

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			law, custody, and domestic violence). Yet, over one million hearings per year go undocumented, and many judges refuse to allow litigants to record proceedings, leaving no historical record for AI. AI systems are trained on large datasets. So what data will be used to train these models in California courts? With so many hearings off-record, we risk developing AI that reflects gaps in transparency, unofficial influences, and judicial behavior that evades oversight. This makes AI deployment in California’s courts particularly dangerous, not only ethically but also legally.	California Code of Judicial Ethics and related ethical guidance. The task force concluded that the canons likely prohibit judicial officers from having generative AI write their opinions for them, and that the Supreme Court’s judicial ethics committees are the appropriate bodies to ask for guidance on this subject.
			Judicial Council’s AI Task Force Concerns The composition of the Judicial Council’s AI Task Force itself raises red flags. The Task Force is made up almost entirely of judges, CEOs of court systems, and insiders with limited technological expertise. At a recent presentation, the lack of baseline understanding about AI among task force members was evident, underscoring the need for education before regulation.	The task force believes that it can make fair and impartial recommendations and that it is sufficiently informed to make the recommendations in this proposal.
			Rather than drafting policy in isolation, the Judicial Council should: • Invite public interest technologists, civil liberties groups (e.g., ACLU), and academic AI ethicists into the conversation. • Prioritize training for judges and court staff in digital literacy and ethics. • Focus on modernizing public records access and ensuring all court data is equitably and transparently available before any AI model is used or trained.	The purpose of the invitation to comment is to invite all interested stakeholders into the conversation. To the extent this suggestion is beyond the scope of the current proposal, the task force may consider it as time and resources permit.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			Efficiency vs. Accountability If AI is to be used at all, it should be limited to non-judicial, public-facing applications such as: • Streamlining access to public records • Automating clerical tasks like form completion • Improving notice systems for hearings and filings • <u>Scrapping Judge and Court Staff 700 Forms to compare with court assignments and outcomes.</u> However, any such use must be accompanied by: • Full public oversight • Independent audits (outside the courts and legal profession) • Strict data governance protocols Let's be clear: If attorneys are at risk of being replaced by AI, as some recent reports predict (Business Insider, 2024), then so are judges and court staff. Any public investment in AI must account for this potential displacement, not exacerbate it through premature or unregulated implementation.	The task force appreciates the commenter's concern but concluded that individual courts are in the best position to determine which uses are appropriate for their specific needs and circumstances. The task force is concerned that placing branchwide limitations on specific uses of generative AI will unnecessarily limit innovation and will prevent courts from identifying safe, effective uses of generative AI that do not pose ethical risks. Additionally, this suggestion is beyond the scope of the current proposal, but the task force may consider it as time and resources permit.
			Conclusion The analogy is simple: Allowing California courts to use AI today is like handing a modern teenager a payphone and expecting them to call their friends without any	Please see previous responses to Susan Bassi's specific suggestions.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			coins. Individuals in the judiciary and employed in the courts simply lack the knowledge, infrastructure, and oversight to implement AI responsibly. For now, no new rule, law, or policy should be passed that allows judicial officers or court staff to use AI in any official capacity, especially in ways that affect legal rulings or public information. Instead, California must focus on: • Ensuring complete and accurate human-generated records of all proceedings • Improving public transparency • Educating both the legal community and the public on what AI is, what it is not, and how it might one day be used in the courts.	
3.	California Employment Lawyers Association by Barbara Figari Cowan, Chair	NI	On behalf of the California Employment Lawyers Association (CELA), a statewide organization of more than 1,300 attorneys who represent workers in employment and civil rights litigation, we respectfully submit this comment in response to SP25-01 regarding proposed standards on the use of generative artificial intelligence (AI) in the judicial branch. CELA supports the Judicial Council’s recognition of the growing impact of generative AI on the practice of law and the importance of proactive safeguards. However, we strongly urge the Council to adopt a single, uniform statewide standard that applies across all courts in California. Fragmented, local approaches would create unnecessary complexity and inequity for court users—particularly for self-represented litigants and those practicing in multiple jurisdictions. Consistency promotes fairness, transparency, and efficiency. Procedural fairness requires that all court users be able to rely on a predictable framework. Varying local rules around the use of AI—whether in filings, court-generated content, or internal judicial processes—could result in litigants receiving	Please see responses to CELA’s specific suggestions below. The task force concluded that the current proposal strikes the best balance between uniformity and flexibility. Use of generative AI will look very different depending on the court, and each court is in the best position to determine how it can meet rule 10.430’s requirements and whether its generative AI use policy should

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			different treatment depending solely on the venue. This is particularly concerning in employment and civil rights cases, where many of our clients already face significant barriers in accessing justice.	be more restrictive or detailed than the rule.
			We also express concern about the potential displacement of human jobs—particularly those held by court clerks, research attorneys, and legal support staff—if generative AI is integrated into court operations without strict limitations and transparency. These workers form the backbone of the judicial system and possess irreplaceable institutional knowledge, cultural competence, and human judgment. Replacing skilled staff with automated tools not only threatens livelihoods but risks eroding the quality and empathy of court services. As advocates for workers’ rights, CELA urges the Judicial Council to explicitly consider the labor impact of AI adoption and to incorporate safeguards against unnecessary job loss or deskilling of essential roles.	While this suggestion is beyond the scope of the current proposal, the task force appreciates this information.
			Finally, we are also concerned about the potential for generative AI tools to introduce or replicate bias, misstate legal authority, or fabricate information (“hallucinations”). Courts must maintain strict standards to ensure that decisions are grounded in verified fact and law—not in machine-generated content that lacks human oversight. A uniform rule should clearly prohibit reliance on generative AI for legal reasoning or fact-finding in judicial decision-making, and should require clear disclosure whenever such tools are used in drafting any court-generated materials.	The task force appreciates with the commenter’s concerns regarding the risks of generative AI tools and notes that it is recommending adoption of a rule of court and a standard of judicial administration to address the confidentiality, privacy, bias, safety, and security risks posed by use of generative AI in court-related work.
			Further, the integrity of the judicial system depends on public trust. That trust could be undermined if courts vary in their use of AI tools without clear guidance or explanation. A uniform rule reinforces the branch’s commitment to transparency and high standards, especially as technology evolves faster than the law can respond.	Please see previous response on this issue.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			In addition, a statewide rule would reduce administrative burden. Rather than asking every court to individually interpret, draft, and implement policy, one consistent standard allows for centralized oversight, training, and review. It also levels the playing field by ensuring that all litigants—regardless of location or resources—have the same expectations and protections when it comes to AI use in the courtroom.	
			We recommend that any rule adopted should (1) prohibit reliance on AI for legal conclusions or fact-finding by courts; (2) require that any use of AI tools in court-generated documents be disclosed; (3) confirm that parties are responsible for verifying any AI-generated content they submit; and (4) be subject to ongoing review as technology and use cases evolve.	The task force concluded that the question of whether and to what extent judicial officers may use generative AI to carry out their adjudicative duties is more appropriately addressed by the California Code of Judicial Ethics and related ethical guidance. The task force concluded that the canons likely prohibit judicial officers from having generative AI write their opinions for them, and that the Supreme Court's judicial ethics committees are the appropriate bodies to ask for guidance on this subject. In light of all the comments regarding requirements to disclose use of generative AI, the task force has recommended a revised version of rule 10.430(d)(5). However, the task force concluded that mandatory disclosure of use of generative AI

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
				in all circumstances could unnecessarily prohibit courts from using generative AI in circumstances where the technology can be used safely and ethically, and that mandatory disclosure would not be an effective way to address concerns about the reliability or trustworthiness of generative AI outputs in many circumstances.
			CELA thanks the Judicial Council for addressing this important issue and for the opportunity to comment. We would welcome future participation in any discussions or working groups as the rule is developed and implemented.	No response required.
4.	Court Watch Silicon Valley [No further commenter information provided]	N	Court Watch Silicon Valley is an informal association of voters, parents, and grandparents working in or around Silicon Valley’s technology sector, who have been directly impacted by practices in the local court system. We are writing to express our strong opposition to the Judicial Council’s proposed rule changes under SP25-01, which would authorize and expand the use of artificial intelligence (AI) in both adjudicative and administrative functions within California’s courts. Our concerns are particularly acute given the documented conduct of the Santa Clara County Superior Court, a jurisdiction at the heart of Silicon Valley and one deeply entangled in longstanding transparency and accountability issues.	The task force notes that the proposed rule and standard do not require courts to permit the use of generative AI by court staff or judicial officers. The proposed rule and standard also do not determine whether use of generative AI is appropriate for any particular task. Rather, the purpose of the rule and standard is to identify and address the risks of using generative AI in court-related work. The task force determined that it is necessary to recommend adoption of the rule and standard

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			***¹ Conclusion AI use in all industries depends on accurate information for input, which the courts simply cannot provide. Therefore, we urge the Judicial Council to limit the use of AI strictly to internal administrative functions aimed at improving public access, court efficiency, and technological modernization before incurring expense and risk to allow its use by judges and court staff. AI should not be used in adjudicative roles or for the analysis of confidential data until: • Accurate and complete digital records are maintained and publicly accessible. • Oversight mechanisms are in place to prevent misuse or selective access. • Ethical and privacy concerns are fully addressed. • Judicial officers and staff are properly trained, and TESTED for the responsible use of AI. California's courts must earn public trust through transparency and accountability before they can responsibly integrate artificial intelligence into the judicial process.	because there are circumstances where the technology can be used safely and ethically, such as legal research using purpose-built legal research tools from trusted providers. Additionally, generative AI is increasingly being incorporated into existing software products and may already be difficult to avoid in some circumstances. The task force is concerned that placing branchwide limitations on specific uses of generative AI will unnecessarily limit innovation and will prevent courts from identifying safe, effective uses of generative AI that do not pose ethical risks. The task force will continue to consider how to address the risks posed by use of generative AI in court-related work. In addition to the proposed rule and standard, the task force has developed FAQs and is considering whether additional guidance documents are needed. The task force will

¹ The commenter provided extensive comments here that did not address the proposal. These comments are, therefore, not included in this chart.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
				work with the Center for Judicial Education & Resources (CJER) to ensure that judicial officers and court staff receive education and training regarding generative AI, including on emerging uses and risks.
5.	Rebecca A. Delfino Associate Dean for Clinical Programs and Experiential Learning Faculty Director Moot Court Programs Law Professor Loyola Law School, Loyola Marymount University, Los Angeles	NI	Introduction The Judicial Council’s Artificial Intelligence Task Force has taken an important and commendable first step in addressing the profound implications of the use of generative artificial intelligence (AI) by California courts. The proposed California Rule of Court, Rule 10.430 (Rule 10.430), California Standards of Judicial Administration Standard 10.80 (Standard 10.80), and the accompanying Model Policy for Use of Generative Artificial Intelligence in the California Courts (Model Policy) reflect the awareness that AI technologies—especially those capable of generating text, images, or legal analysis—will increasingly shape how courts operate, communicate, and render decisions. The proposals seek to guide the appropriate, ethical, and effective use of generative AI in court operations and ensure that generative AI tools are used to respect legal and constitutional obligations, safeguard public trust, and preserve the integrity of court functions. At the same time, as with any emergent regulatory framework, these initial proposals would benefit from further refinement. The rapid pace of generative AI development, the diversity of use cases within the courts, and the need to maintain public trust in judicial integrity all underscore the importance of ensuring that this framework is as clear, consistent, and comprehensive as possible. This comment builds upon the strong foundation in Rule 10.430, Standard 10.80, and Model Policy by offering specific, constructive recommendations to strengthen the proposed framework. Drawing on principles of good governance, institutional	Please see responses to Rebecca Delfino’s specific suggestions below.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			integrity, and practical implementation, the suggestions offered here aim to enhance the transparency of AI use, ensure accountability at both the institutional and individual levels, and reinforce the judiciary’s leadership in the ethical deployment of advanced technologies. My recommendations are organized in two parts. Part I addresses general concerns that cross the entire framework of Rule 10.430, Standard 10.80, and the Model Policy. Part II of the comment provides specific and individual suggestions for the Rule and Model Policy, aimed at strengthening the proposals by adding accountability, clarifying obligations, and reinforcing public transparency.	
			I. Framework-Level Observations A. Definition of “Generative AI” The current definition of “generative AI” in Rule 10.430, Standard 10.80, and the Model Policy reads: “Generative AI means artificial intelligence trained on an existing set of data (which can include text, images, audio or video) with the intent to generate new data objects when prompted by a user. Generative AI creates new data objects contextually in response to user prompts, based only on the data it has already been trained on.” This definition captures the general idea behind generative AI, but could be more precise, broader, and more adaptable.	No response required.
			First, the definition should be refined to enhance its clarity. It uses technical language such as “data objects,” a term that may be unfamiliar or ambiguous to many users, including judicial officers and court staff. More intuitive language—such as “content,” “text,” or “images”—would make the definition more accessible. Moreover, the phrase “with the intent to generate” is problematic. The intent of the system’s developers is difficult to assess and not always relevant for policy purposes. What matters more is the system’s functionality: whether it can generate content in response to user input. Finally, the clause stating that the system responds “based only on the data it has already been trained on” may be misleading. Many modern generative AI tools incorporate additional capabilities such as retrieval-	The task force appreciates these suggested definitions and recommends revising the definition of “generative AI” in the rule and standard based on the commenter’s suggestion. Additionally, in light of all the comments received on this issue, the task force is recommending a

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			augmented generation (RAG), which allows them to access external or real-time data in addition to their training set. Second, the definition does not reflect the full range of technologies it aims to cover. It focuses primarily on traditional large language models trained on static datasets. However, it does not account for a growing range of generative systems that produce not only text but also code, legal summaries, charts, and audio-visual outputs. Additionally, it omits multi-modal systems, interactive chat-based models, and domain-specific tools—such as legal research platforms—that combine generative AI with real-time access to proprietary databases. Third, the definition does not account for future developments. As generative AI continues to evolve—especially with systems that integrate static training data with real-time querying, embedded legal databases, and user-contextual interactions—the current language may prove too limited. A more forward-looking framework would help ensure lasting relevance. To address these concerns, the following revised definition is proposed: “Generative artificial intelligence” means a computer-based system that uses machine learning or similar techniques to produce new content—such as text, images, audio, video, code, or data visualizations—in response to user inputs. Generative AI systems generate content that is not pre-programmed or explicitly retrieved but synthesized based on underlying models trained on large datasets and may include integration with real-time or domain-specific sources. This revised definition avoids technical jargon, reflects the full range of generative capabilities, and anticipates the future evolution of AI systems used in the judicial context. It would help ensure that the rule and model policy remain effective and adaptable as technology advances.	revised definition of “public generative AI system” in the rule and standard, and recommends deleting the definition of “artificial intelligence” in the rule and standard because that definition is no longer needed due to the proposed revisions to the definition of “generative AI.”

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			B. Transparency and Disclosure Requirements The requirements for court staff and judicial officers to disclose the use of generative AI in Rule 10.430,[1] Standard 10.80,[2] and Model Policy [3] reflect a shared goal: to promote public trust by ensuring that judicial use of generative AI is transparent when it meaningfully contributes to publicly available work. However, these provisions are framed with differing levels of obligation and clarity, and when read together, they reveal internal inconsistencies in language, legal effect, and implementation standards. Aligning these elements through more consistent language could help better advance the framework’s intended goals. [1] Rule 10.430: “Each courts generative AI use policy must: (5) Require disclosure of the use or reliance on generative AI if generative AI outputs constitute a substantial portion of the content used in the final version of the written or visual work provided to the public.” [2] Standard 10.80: “A judicial officer using generative AI for any task within their adjudicative role: (5) should consider whether to disclose the use of generative AI if it is used to create content provided to the public.” [3] Model Policy. “VI Transparency: (a) If generative AI outputs constitute a substantial portion of the content used in the final version of a written work or visual work that is provided to the public, the. Work must contain a disclaimer or watermark. (b) Labels are watermarks used to disclose the use of generative AI should be easily visible and understandable, accurately informing the audience that generative AI has been used in the creation of the content and identifying the system used to generate it.” Rule 10.430 helpfully establishes a clear, mandatory baseline: each court’s policy on generative AI use must include a requirement to disclose when AI outputs constitute a substantial portion of a final written or visual work shared with the public. This language provides helpful clarity and a firm obligation. However, that clarity is somewhat diluted by the wording in Standard 10.80, which applies to individual judicial officers. Rather than requiring disclosure, it states that a judicial	In light of all the comments received on this issue, the task force has recommended a revised version of rule 10.430(d)(5). However, the task force concluded that mandatory disclosure of use of generative AI in all circumstances could unnecessarily prohibit courts from using generative AI in circumstances where the technology can be used safely and ethically, and that mandatory disclosure would not be an effective way to address concerns about the reliability or trustworthiness of generative AI outputs in many circumstances. Additionally, as time and resources permit, the task force will consider whether to revise the model policy to more specifically address the use of generative AI by research attorneys.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			officer “should consider whether to disclose” the use of generative AI in public-facing content in their adjudicative work. This shift from a mandatory requirement in the Rule to a more discretionary guideline in the accompanying Standard may lead to variations in interpretation and practice across courts and individual judges. In addition, Rule 10.430, Standard 10.80, and the Invitation to Comment (SP25-01) fail to articulate a rationale for imposing a mandatory disclosure requirement when judicial officers employ generative AI in non-adjudicative contexts, while permitting discretion when the same technology is used in adjudicative decision-making. This distinction raises concerns. The adjudicative function—where legal reasoning is developed, decisions are rendered, and public confidence in impartiality is most essential—arguably demands the highest level of transparency. If the purpose of disclosure is to maintain public trust and to ensure accountability in the use of emerging technologies, it is unclear why that obligation should be stronger when a judicial officer uses AI to draft a standing order or informational notice but weaker when the same tool helps shape the resolution of a legal dispute. Without a clear policy rationale for this differentiation, the framework risks sending inconsistent signals about the values it seeks to uphold and the contexts in which transparency matters most. A uniform disclosure obligation—one that applies to both adjudicative and non-adjudicative use when the AI’s contribution is substantial or material—would better reflect the core principles of judicial integrity and public accountability that underlie the framework. Furthermore, the Model Policy, intended to guide implementation, adds another layer of complexity. Like Rule 10.430, it states that a disclosure is required—specifically, that a disclaimer or watermark must be included if generative AI constitutes a substantial portion of a final written or visual work provided to the public. It also requires that this label be “easily visible and understandable” and that it identify the system used to generate the content. These detailed requirements are helpful, but they are found only in the Model Policy and not echoed in the Rule or the Standard. As a result, it is unclear whether the visibility and specificity requirements are binding, optional, or merely best practices. Divergent standards—	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			mandatory, discretionary, and directive—send conflicting signals about whether disclosure is a firm requirement or optional practice and to whom it applies. These internal inconsistencies in the disclosure and transparency requirements across Rule 10.430, Standard 10.80, and the Model Policy are not merely theoretical—they create practical dilemmas for judicial officers and court staff. A clear example arises in the context of judicial drafting practices involving research attorneys. Imagine a scenario in which a research attorney prepares a draft opinion or memorandum in which a substantial portion of the content is generated with the assistance of a generative AI tool. Under Rule 10.430, which mandates that each court’s AI policy “require disclosure of the use or reliance on generative AI if generative AI outputs constitute a substantial portion of the content used in the final version of the written or visual work provided to the public,” the use of that AI-generated content triggers a disclosure obligation. Similarly, under the Model Policy, a substantial portion of AI-generated content in the final public-facing work would require a visible disclaimer or watermark. However, if the judge receives that draft from the research attorney and adopts it without independently knowing about or considering the AI-assisted drafting process, it is unclear whether the judge must disclose the AI use—especially in light of Standard 10.80, which merely states that a judicial officer “should consider” disclosure and does not contain a firm requirement. Because the standard relies on the judge’s awareness and discretion, it creates a potential gap in disclosure when the generative AI use originates with staff rather than the judicial officer directly. This example highlights a compliance dilemma: the court’s policy, under Rule 10.430, may require disclosure, and the Model Policy contemplates mandatory labeling, but the judge may not feel personally obligated to disclose under the permissive language of the Standard—particularly if the generative AI use occurred earlier in the drafting chain. This disconnect blurs the line of responsibility and opens the door to inconsistent outcomes, where disclosure depends not on the extent	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			of AI use but on who used it and whether they viewed themselves as bound by a mandatory or discretionary standard. These inconsistencies could lead to under-disclosure, jeopardizing the transparency the framework seeks to ensure, or to overcautious practices where judges avoid useful AI tools altogether for fear of accidental non-compliance. Clarifying and harmonizing the obligations across all levels of court actors—judicial officers, staff attorneys, and court administrators—is essential to creating a workable, fair, and transparent system. Another concern involves striking the right balance between transparency and judicial independence, particularly regarding the use of generative AI by court staff and judicial officers for internal purposes. If a generative AI system is used solely for internal purposes—for example, to help summarize a case, draft a bench memo, or brainstorm issues—and the outputs are neither shared with the parties nor appear in any written or visual work provided to the public, then under the current language of Rule 10.430 and the Model Policy, no disclosure would be required. That result is consistent with a long-standing norm in judicial practice: internal deliberative tools and communications—including memos from law clerks or research attorneys—are not disclosed. The line drawn by the Rule and Model Policy around “written or visual work provided to the public” appears intentional and grounded in respect for this boundary. However, this limitation also raises a key policy tension: some internal uses of generative AI may materially shape judicial reasoning or decisions, even if those outputs never appear on the public record. The public may reasonably expect to be informed not only when AI drafts the text of a published ruling but also when it substantively influences the decision-making process. In an era where AI can go beyond rote summarization and actively generate legal arguments or identify perceived weaknesses in a claim, even internal use may carry normative weight.	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			In this respect, excluding non-public uses may be technically consistent with the rule’s drafting but normatively incomplete. By drawing the line strictly at “what is seen,” the framework could enable significant, even outcome-determinative use of generative AI without disclosure—not out of bad faith, but because the current rule does not reach that far. That gap risks undermining transparency, especially if it becomes widely known that AI is playing an influential (albeit invisible) role in judicial decision-making. Beyond these inconsistencies, a misalignment in scope and triggering criteria also exists. The Rule and Model Policy reference the concept of a “substantial portion” of AI-generated content, but so not define what that means. The phrase “substantial portion,” as used in Rule 10.430 and the accompanying Model Policy, is vague and potentially problematic for courts and judges seeking to comply with disclosure or certification requirements regarding the use of generative artificial intelligence. Its ambiguity stems primarily from the absence of a clear, objective threshold. The term “substantial” is inherently relative—it may mean a majority in some contexts or merely something of importance in others. Without a defined metric or standard, judges and court personnel are left to speculate whether a given use of AI qualifies as “substantial,” which could lead to inconsistent interpretations and application across different judicial officers and courts. Moreover, the “substantial portion” is context-dependent. What might constitute a substantial use of AI in drafting a routine procedural order could differ significantly from its use in composing the reasoning of a complex opinion. This contextual variability further undermines uniformity and predictability in compliance. The phrase also lacks the support of an established body of judicial interpretations, unlike similar terms in copyright law or employment law, where courts have had decades to define and refine their meaning. In this context—where policies governing judicial reliance on AI are still emerging—such ambiguity risks either chilling the appropriate use of technology or, conversely, enabling under-disclosure of its influence on judicial reasoning.	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			To resolve these problems, a single, unified standard should be adopted across the Rule, Standard, and Model Policy. First, disclosing the use of generative AI should be mandatory whenever AI contributes materially or substantially to any court work. This requirement should apply to both court staff and individual judicial officers in their adjudicative and non-adjudicative roles, closing the gap between institutional requirements and individual responsibility. Second, the threshold for disclosure should be clearly defined—for example, by stating that a “substantial portion” includes any AI-generated content that materially influences the reasoning, substance, or language of the final work or that comprises more than 20% of its content. Third, the requirement for a clear, visible, and understandable disclaimer or watermark that names the AI system used should be incorporated into the Rule and Standard and the Model Policy. Finally, rather than mandating disclosure of all internal uses—which would intrude into the judicial deliberative process and likely face resistance—a middle path could be adopted: Proposed Harmonized Disclosure Provision (Model Language) Disclosure Requirement: “Judicial officers and court staff must disclose the use of generative artificial intelligence when (1) its outputs constitute a substantial portion of any written or visual content provided to the public, or (2) its use materially informs the reasoning, analysis, or resolution of a case, even if the AI-generated content does not appear in a written or visual work provided to the public.” • “Materially informs” means that the use of generative AI contributes in a non-trivial way to the reasoning, analysis, or outcome of a judicial decision, including shaping conclusions, influencing legal interpretation, or framing the resolution of factual or legal issues, regardless of whether the AI-generated content appears in the final written work. • “Substantial portion” means any AI-generated text, analysis, or recommendation that materially influences the reasoning, outcome, or language of a final judicial decision, order, or final work or comprises more than 20% of the output. Disclosure must be made through a visible	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			and understandable label or watermark that identifies the use of generative AI and names the system used. This definition offers both a qualitative and quantitative benchmark, improving clarity, encouraging consistent application, and fostering transparency in the judicial use of generative AI. Furthermore, this approach would preserve judicial discretion over internal deliberation, recognize that some internal uses are minor or administrative and don't merit disclosure, and encourage greater transparency where AI plays a substantive, if hidden, role in shaping the outcome. Although the current rule's focus on public-facing work is justifiable, it may not go far enough to account for the evolving ways in which generative AI can influence judicial decision-making. A more nuanced, optional disclosure pathway for significant internal uses could better align with the spirit of transparency the framework seeks to uphold. By speaking in a single, consistent voice, the Rule, Standard, and Model Policy can more effectively achieve their shared goal: maintaining public confidence in the integrity and transparency of the courts in an era of rapidly evolving technology. Harmonizing the standards and adopting a clear, uniform approach fosters transparency while promoting responsible and consistent use of generative AI throughout the judiciary.	
			II. Specific Suggestions for Rule 10.430 and the Model Policy A. Rule 10.430 While Rule 10.430 provides a useful starting point by requiring each California court to adopt a policy on generative AI use, the current version is too narrow in scope and lacks the structural safeguards necessary to ensure effective, ethical, and equitable implementation across the judicial system. Given the rapid evolution of generative AI tools and their increasingly sophisticated applications in legal and	Please see responses to individual suggestions below.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			judicial contexts, the Rule should do more than merely mandate the adoption of a policy. It must establish a framework for ongoing oversight, transparency, staff education, and equitable deployment. The following additions to Rule 10.430 are designed to promote the principled use of Generative AI and preserve public trust in judicial integrity.	
			1. Regular Policy Review Generative AI systems are developing at a pace that far outstrips the ordinary cadence of rulemaking and administrative reform. Absent regular policy updates, courts risk operating under outdated assumptions about generative AI's capabilities, vulnerabilities, or ethical implications. A routine, mandated review cycle ensures that courts can respond in real-time to emerging threats—such as new forms of deepfake manipulation or data leakage vulnerabilities—and can integrate best practices as they are developed across jurisdictions. An annual review is not only good governance; it is essential to risk management in a dynamic technological landscape. To address these concerns, the additional language is proposed: Each court shall review its generative AI policy at least annually, updating it to reflect technological advances, emergent risks, and evolving best practices.	Revising the rule and standard to implement this suggestion would require further public comment because it is beyond the scope of issues presented in this invitation to comment. The task force may consider it as time and resources permit. Additionally, the task force will consider whether to implement this suggestion, and others that are beyond the scope of this invitation to comment, via other means including the model policy or other guidance documents.
			2. Encourage Transparency and Public Disclosure Public confidence in the judiciary depends not only on fair outcomes but also on institutional transparency. As courts increasingly rely on generative AI tools—whether for drafting notices, processing filings, or analyzing legal materials—litigants and the public have a legitimate interest in understanding the nature and limits of those tools. Posting generative AI policies online enables court users to understand when and how generative AI may be used in judicial or administrative communications. This form of transparency is a low-cost but high-impact	As discussed above, this suggestion is beyond the scope of the current proposal, but the task force may consider it as time and resources permit.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			mechanism for reinforcing public trust, particularly when misinformation about AI use is widespread. To address these concerns, the additional language is proposed: Courts are encouraged to publish their generative AI policies and summaries online for public access.	
			3. Require Training for Judicial Officers and Court Staff Judicial decisions rest on informed and independent judgment. That judgment is compromised when those responsible for exercising it are unaware of the tools they are using. Given the black-box nature of many generative AI systems, meaningful use requires a baseline understanding of their technical architecture, strengths, and limitations. Regular training ensures that judicial officers and court personnel can recognize when generative AI outputs are unreliable, when bias may be introduced, and when human review is especially critical. Without this foundational knowledge, courts risk overreliance on tools that may generate plausible—but substantively incorrect—outputs. To address these concerns, the additional language is proposed: Judicial officers and staff shall receive regular training on generative AI capabilities, limitations, risks, and approved court uses.	As discussed above, this suggestion is beyond the scope of the current proposal, but the task force may consider it as time and resources permit.
			4. Add a Misuse Reporting Mechanism No technology is error-proof, and generative AI tools—especially those powered by probabilistic models—are uniquely prone to unintentional misuse, hallucinated outputs, and systemic bias. A clear mechanism for reporting incidents or suspected misuse is critical for early detection and course correction. Moreover, a uniform process for elevating serious or systemic issues to the Judicial Council would enable centralized tracking of trends, inform future policy, and ensure consistency across	As discussed above, this suggestion is beyond the scope of the current proposal, but the task force may consider it as time and resources permit.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			jurisdictions. Just as courts have protocols for reporting security breaches or ethical misconduct, generative AI failures warrant formal oversight. To address these concerns, the additional language is proposed: Each court shall establish an internal reporting process for Generative AI misuse or failures, with escalation to the Judicial Council where appropriate.	
			5. Require Evaluation Before Adoption of New Tools Generative AI vendors often market tools without sufficient empirical support for their reliability or alignment with legal standards. Before integration into judicial workflows, generative AI products must undergo a vetting process to assess technical accuracy, adherence to data privacy standards (particularly when sensitive or confidential filings are involved), and bias mitigation protocols. This is especially important for tools that summarize legal arguments, predict outcomes or draft documents with legal effect. Procurement without prior evaluation exposes courts to reputational, legal, and operational risk. To address these concerns, the additional language is proposed: Before procuring or deploying a generative AI tool, courts must evaluate its reliability, security, and alignment with privacy laws.	As discussed above, this suggestion is beyond the scope of the current proposal, but the task force may consider it as time and resources permit.
			6. Address Access Disparities and Resource Inequity and Encourage Cross-Court Collaboration and Sharing of Best Practices Without centralized support, Rule 10.430 could exacerbate inequalities between well-resourced urban courts and smaller or rural jurisdictions. Wealthier courts may benefit from the cost and efficiency gains of vetted generative AI tools, while others may lack the personnel or technical infrastructure to deploy or evaluate such tools responsibly. Uniform access to vetted tools, shared training modules, and centralized policy templates would help prevent a two-tiered system in which only some courts can take advantage of innovation—or meet compliance burdens—	As discussed above, this suggestion is beyond the scope of the current proposal, but the task force may consider it as time and resources permit.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			effectively. Equal access to reliable tools is essential to maintaining a fair and unified judicial system. The Judicial Council is encouraged to provide guidance and support to ensure smaller or under-resourced courts are not disadvantaged in implementing generative AI tools or policy compliance. Relatedly, while many courts will likely encounter similar challenges in deploying generative AI tools—ranging from staff training to tool evaluation—Rule 10.430 should promote shared learning and collective problem-solving. The challenges courts face in implementing generative AI are not unique. Issues such as vetting tools, managing training gaps, assessing risks, and communicating with the public will arise in every jurisdiction. Yet without a mechanism to promote information-sharing, courts may duplicate efforts, develop inconsistent approaches, or miss opportunities to learn from one another. By encouraging collaboration through the Judicial Council, the rule can help ensure that successes in one court inform practices in others—particularly beneficial for smaller or under-resourced courts. To address these concerns, the additional language is proposed: Courts are encouraged to share lessons learned, policy templates, and successful use cases with the Judicial Council to support cross-jurisdictional innovation and consistency. This provision recognizes that policymaking in the generative AI space is still evolving and that courts can be partners in that process. A culture of collaboration would reduce redundancy, promote higher-quality policies, and allow the Judicial Council to aggregate experiences and identify system-wide trends or challenges.	
			7. Establish Oversight Structures Rule 10.430 appropriately requires each court to adopt a policy governing generative AI use, but it is silent on who within the court is responsible for	As discussed above, this suggestion is beyond the scope of the current proposal, but the task

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			overseeing implementation or responding to questions about policy interpretation, compliance, or updates. For a generative AI policy to function effectively, courts must have clear lines of responsibility. Without designated oversight, questions about how the policy applies to specific tools or use cases may go unanswered. Moreover, regular policy review—essential to ensure alignment with evolving technologies—requires internal leadership. Whether oversight is assigned to an individual officer (e.g., a court executive or technology lead) or a committee, a formal point of accountability is critical to ensure compliance and support implementation. To address these concerns, the additional language is proposed: Each court shall designate a responsible officer or committee to oversee generative AI policy implementation and updates. This addition would promote internal accountability and ensure that generative AI policies are not merely aspirational documents but actively maintained, interpreted, and enforced. It also provides a designated point of contact for court staff with questions or concerns, enhancing operational clarity and promoting responsible adoption. By requiring internal oversight and encouraging external collaboration, the Judicial Council can move from mandating AI policy adoption to fostering a coherent, well-supported, and ethically grounded framework for judicial innovation.	force may consider it as time and resources permit.
			B. Model Policy for Use of Generative Artificial Intelligence As generative AI tools become more integrated into court operations, the absence of clearer limits, documentation practices, and vendor standards risks undermining public confidence in the judiciary and exposing courts to unintended harm. The following proposed revisions are intended to fortify the Model Policy by setting clearer expectations and promoting sound governance in the judicial use of generative AI.	The task force will consider these suggested revisions to the model policy as time and resources permit.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			1. Clarify Permitted and Prohibited Uses The Model Policy does not currently distinguish between acceptable administrative uses of generative AI and prohibited adjudicative uses. This omission leaves court personnel uncertain about the boundaries of permissible AI use and creates the risk that AI-generated content could be relied upon, even inadvertently, to resolve litigants' legal claims. While generative AI may provide efficiency in low-risk contexts, its use in drafting judicial decisions or analyzing legal arguments threatens to displace independent judicial reasoning. Courts need practical guidance to distinguish between operational assistance and adjudicative overreach. To address these concerns, the following additional language is proposed to provide examples of appropriate and inappropriate uses of the technology: <i>Permitted Uses:</i> Preparing administrative memoranda, summarizing policy documents, generating FAQs or procedural information for court users. <i>Prohibited Uses:</i> Drafting judicial rulings or decisions, analyzing arguments from litigants, and generating materials representing authoritative judicial reasoning.	The task force will consider these suggested revisions to the model policy as time and resources permit.
			2. Require Documentation and Accountability The absence of a documentation requirement in the Model Policy makes it difficult for courts to audit or evaluate how generative AI is used. As reliance on these tools grows, courts need internal records that trace which tools were used, for what purposes, and how the outputs were reviewed. This is especially true for medium- and high-risk applications, such as summarizing briefs, drafting communications to the public, or assisting with complex filings. Documentation not only ensures accountability but also provides a mechanism for institutional learning, quality control, and responsible innovation. To address these concerns, the following additional language is proposed:	The task force will consider these suggested revisions to the model policy as time and resources permit.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			Courts must maintain detailed records of medium- and high-risk generative AI usage, specifying the generative AI tools employed, the purposes, reviewers involved, and final verification of outputs.	
			3. Annual Review and External Audits Generative AI is an evolving technology that regularly introduces new features, risks, and legal implications. Without periodic reassessment, courts may continue using tools that have become outdated, less secure, or noncompliant with new legal or ethical standards. While internal review is essential, high-risk applications—such as those that influence case processing or are visible to the public—also warrant external, independent evaluation. External audits provide objectivity, reveal blind spots, and reinforce public trust in court governance. To address these concerns, the following additional language is proposed: Courts must review their generative AI policy annually and consider third-party audits for high-risk applications.	The task force will consider these suggested revisions to the model policy as time and resources permit.
			4. Minimum Standards for Vendor Vetting Courts have a legal and ethical obligation to ensure that the technology they use aligns with public sector requirements for data privacy, non-discrimination, and transparency. Not all generative AI vendors meet these standards, and courts should not assume that commercial products are suitable for judicial use without independent review. A procurement framework with minimum vetting criteria will help safeguard against tools that embed bias, mishandle data, or fail to explain how their outputs are generated. To address these concerns, the following additional language is proposed: Courts shall assess vendor compliance with privacy, nondiscrimination, and transparency standards before procuring generative AI tools.	The task force will consider these suggested revisions to the model policy as time and resources permit.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			5. Clarify Data Retention and Consent for Public Interactions Many generative AI platforms retain user input data and may use it to improve their models. If a court user submits information to an AI-powered system (for example, to receive procedural guidance), they may not realize that their input is being stored or that their interaction is with an AI system. Without clear policies on data retention and user consent, courts risk violating individual privacy rights and undermining transparency. Court users should be informed when interacting with an AI system and whether their data is stored or shared. To address these concerns, the following additional language is proposed: Policies must clarify whether and how generative AI tools store input or output data and whether court users will be informed or required to consent when interacting with generative AI-generated materials.	The task force will consider these suggested revisions to the model policy as time and resources permit.
			These additions to the Model Policy would transform it from a general statement of caution into a practical governance tool that supports the responsible, transparent, and ethical use of generative AI in judicial proceedings. As these technologies become more deeply embedded in court operations, clear standards will be essential to preserving judicial independence and public confidence in the administration of justice.	The task force will consider these suggested revisions to the model policy as time and resources permit.
			Conclusion The Judicial Council’s Artificial Intelligence Task Force’s proposed Rule 10.430, Standard 10.80, and Model Policy provide a thoughtful and forward-looking foundation for addressing the use of generative artificial intelligence in California’s courts. The Judicial Council is to be commended for taking this important step. The initiative reflects a deep awareness of the promise and risks accompanying this emerging technology and a commitment to ensuring its use aligns with the judiciary’s core values of integrity, impartiality, and public accountability.	No response required.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			Even so, to fully achieve the stated goals of transparency, consistency, and ethical use, additional refinements would strengthen the overall framework. As this comment identifies, definitional ambiguities and internal inconsistencies, particularly in the disclosure standards and their scope, create uncertainty for courts and judicial officers. Likewise, the absence of certain structural safeguards, such as oversight roles, training requirements, and guidance on internal use, may limit the effectiveness and uniformity of implementation across jurisdictions. The suggested revisions support the creation of a clear, practical, and enduring framework. They are grounded in the shared goal of maintaining public trust in the courts while fostering thoughtful and appropriate innovation. By harmonizing language across the rule, standard, and policy, clarifying key terms, and incorporating governance best practices, the Judicial Council can further position California's judiciary as a national leader in the responsible integration of generative AI.	
6.	David Freeman Engstrom LSVF Professor of Law Co-Director, Deborah L. Rhode Center on the Legal Profession Stanford Law School	NI	I was honored to present to you last fall, and I commend you on the tremendous work that the Task Force has achieved since then. I am writing to offer public comment on the <i>Model Policy for Use of Generative Artificial Intelligence</i>. Below I detail several ways that the <i>Model Policy</i> as currently written may have implications for the ability of California courts to innovate. I focus in particular, but not exclusively, on the <i>Model Policy's</i> potential impacts court-university partnerships, which I see as a vitally important way that California's courts can continue their leadership at the frontier of justice innovation. As you may recall from our conversation, Stanford Law School's Deborah L. Rhode Center on the Legal Profession and Legal Design Lab are partnering with the Superior Court of Los Angeles County (LASC) to develop and implement a blueprint for more innovative, modern, and accessible courts. In a major report released earlier this month, Stanford and LASC detailed a pioneering plan for justice innovation that includes several projects that seek to leverage emerging generative AI technology.[1] As I interpret the <i>Model Policy</i>, when members of a university	Please see responses to David Freeman Engstrom's specific suggestions below.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			research team serve as research contractors to a California court, they would be considered “court staff” under its terms. [1] DAVID FREEMAN ENGSTROM ET AL., DEBORAH L. RHODE CTR. ON THE LEGAL PROF., A BLUEPRINT FOR EXPANDING ACCESS TO JUSTICE IN LOS ANGELES SUPERIOR COURT’S EVICTION DOCKET (2025), https://drive.google.com/file/d/11OGqy5_U1NfoZAod-9v1PacQTiEHxBsq/view .	
			I. The Definition of “Public Generative AI System” May Generate Confusion The <i>Model Policy</i> defines “public generative AI system” as “a system that is publicly available or that allows information submitted by users to be accessed by anyone other than judicial officers or court staff, including access for the purpose of training or improving the system.”[2] I understand the need for the judiciary to create policies for high-risk categories of AI systems. The breadth of this definition, however, may be problematic. In practice, it may be difficult to delineate what is included in this high-risk category of tools and what is not. The examples of generative AI systems listed in the <i>Model Policy</i> include diverse tools with different risk profiles. The free versions of the foundation models (such as ChatGPT, Claude, Copilot) raise different privacy and security concerns than do the subscription-based versions of these same tools or proprietary, enterprise tools (such as Westlaw Precision and Lexis+AI). Additionally, as the <i>Model Policy</i> recognizes, public LLLMs have been integrated into many commonly used software systems. Many systems currently in use by California courts, and also a growing number going forward, may fall under this broad definition, limiting the tools available to court staff. To avoid implementation headaches, this definition could benefit from additional specificity and/or guidance—for example, by refocusing on acceptable security features and user data policies.	In light of all the comments received on this issue, the task force is recommending revised definitions of “generative AI” and “public generative AI system” in the rule and standard. The task force recommends deleting the definition of “artificial intelligence” in the rule and standard because it is not needed due to the proposed revisions to the definition of “generative AI.”

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			Relatedly, while the <i>Model Policy</i> governs generative AI use by court staff and judicial officers, this definition is also potentially problematic when considering court-deployed, public-facing AI tools, which some courts are developing (including LASC and the Stanford team). The current definition seems to apply the same treatment to tools wherein users interact with AI outputs but cannot access underlying data and tools where underlying data is accessible. Yet these tools present fundamentally different risk profiles. [2] 10.430(5). Emphasis added.	
			II. The Disclosure Requirements for Generative AI Content May be Unnecessary Second, Rule 10.430(d)(5) would require court staff to disclose the use of generative AI in written or visual outputs if a substantial portion of the content was created using these tools. One of the many promising uses of generative AI for courts in widening access to justice is this technology’s ability to create written self-help materials. Using generative AI tools, court staff can expedite the creation of self-help information, as well as streamline and simplify existing legal information resources. Under Rule 10.430(d)(3), court staff must review and ensure that content is accurate and complete. Given this mandated review, Rule 10.430(d)(5)’s disclosure requirement does not seem necessary in all instances. Further, such a disclosure on written legal information may have unintended consequences: readers may be less likely to trust content marked with a generative AI disclosure, perhaps anticipating hallucinations or inaccuracies or perhaps displaying reflexive (and, given review requirements, undue) aversion to machine-generated material. Instead of a brightline rule, the <i>Model Policy</i> could provide a series of considerations for making decisions about which pieces of generative-AI-facilitated content should be marked. Alternatively, the rule for court staff could be amended to reflect the discretion that is built into the analogous provision for judicial officers in Standard 10.80(b)(5): “<i>Should consider</i> whether to disclose the use of generative AI if it is used to create content provided to the public.”[3]	In light of all the comments received on this issue, the task force has recommended a revised version of rule 10.430(d)(5). The task force understands the concern that requiring disclosure might limit courts’ flexibility but concluded that mandatory disclosure is necessary in some circumstances to maintain public trust in the judicial branch.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			[3] Emphasis added.	
			III. Additional Guidance Would Be Helpful on Compliance Measures Research contractors to California courts, and therefore “court staff” under the <i>Model Policy</i> , would benefit from additional guidance in two related areas.	No response required.
			First, the <i>Model Policy</i> understandably prohibits the use of generative AI to “unlawfully discriminate against or disparately impact individuals or communities....” What is unclear is what kind of testing or benchmarking must be undertaken to make the showing that a tool is not having this prohibited effect. The LASC-Stanford team is building an AI-powered, user-directed informational triage tool that allows litigants to self-sort into appropriate legal help pathways from a comprehensive database of legal assistance resources, based on certain case features and litigant needs and preferences. The <i>Model Policy</i> is unclear as to what documentation a court or research team must develop on an AI tool’s operation and impact. More specific guidance on this essential issue can ensure trustworthy court use of AI while avoiding innovation-stymieing uncertainty about evaluation requirements.	The task force will consider these suggested revisions to the model policy as time and resources permit.
			Second, the <i>Model Policy</i> directs court staff and judicial officers to review generative AI material for accuracy and completeness. As detailed in our recent report, the LASC-Stanford team is also prototyping an AI-powered “default assistant” to aid Court staff in ensuring that default judgments are legally warranted. (While not reliant on just generative AI technology, this tool may be considered a public AI system under the broad definition in the <i>Model Policy</i> .) This tool will be available to select court clerks and research attorneys who will have the discretion to either accept or reject the tool’s recommendations before entering a clerk judgment or sending a recommendation to the judicial officer. The tool is designed to streamline the otherwise time-consuming manual review process. Here, too, the <i>Model Policy</i> is unclear as to what documentation a research team must develop to	The task force will consider these suggested revisions to the model policy as time and resources permit.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			demonstrate accuracy and reliability. More specific guidance can both ensure trustworthy court use of AI and avoid stymieing useful innovation.	
			* * * I appreciate your consideration, and thank you, again, for your leadership on these critical issues.	No response required.
7.	Fortuna Arbitration by Kimo Gandall, CEO Cambridge, Massachusetts	NI	Statement of Interest Dear members of the Judicial Council, My name is Kimo Gandall. I am a third-year student at Harvard Law School, a Professional Registered Parliamentarian, and CEO of Fortuna-Insights. I am a lifelong Californian, born and raised in Orange County, and completed my undergraduate studies at UC Irvine (Zot, zot, zot!). My father moved to California from the Hawaiian Islands, and my mother relocated from Waco, Texas. Both of my parents continue to reside in California, as do I. My co-founder, Kenny McLaren, is also a native Californian. Together, Kenny and I founded several companies in California, including WildSafari Studios, during our high school years. Our current company, Fortuna-Insights, Inc., is the parent of Fortuna Arbitration (“Arbitrus.ai”), a legal artificial intelligence firm with significant business relationships in California, involving both investors and clients. We write this letter out of deep concern regarding SP25-01, which represents an impractical attempt to regulate artificial intelligence. Such regulation will substantially increase the cost of engaging private AI vendors. Fortuna supports what Ezra Klein describes in his recent book, Abundance, as “outcome-driven governance.” On Governor Newsom’s podcast, Klein and the Governor advocated shifting from a scarcity mindset toward policies that facilitate the construction of more housing, energy infrastructure, and other essential public	Please see responses to Fortuna Arbitration’s specific suggestions below.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			assets. That is, a mindset away from solely minimizing harm through regulation, and rather employing the government as a vehicle of public service to improve the lives of citizens. The judiciary plays a crucial role in this transformation, as the rule of law underpins a prosperous and healthy society. However, we remain deeply concerned about increasing costs and inefficiencies within California’s public sector, particularly the judiciary. Current case backlogs and prolonged litigation processes not only impose financial burdens but also delay timely justice for Californians. As the Judicial Council considers regulations on generative AI, we urge a balanced approach that acknowledges AI’s transformative potential while maintaining rigorous ethical standards and accountability. Embracing AI will help create a judicial system that is more efficient, cost-effective, and accessible for all Californians. Thus, it is out of love for the state of California that we write you to reconsider the implementation of SP25-01. Respectfully submitted, Kimo Gandall CEO, Fortuna Arbitration	
			Executive Summary Fortuna Arbitration appreciates the opportunity to provide comments on the proposed regulation SP25-01 concerning generative artificial intelligence (AI) in California’s judicial system. Our primary concern is that the proposed regulations, particularly Rule 10.430 and Standard 10.80, while well-intentioned, may inadvertently increase costs, complicate judicial processes, and hinder the adoption of innovative technologies that can significantly improve judicial efficiency and accessibility.	Please see responses to Fortuna Arbitration’s specific suggestions below.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			California’s judiciary currently faces budgetary constraints and significant case backlogs. Rather than imposing restrictive regulations that risk adding further inefficiencies, Fortuna Arbitration strongly recommends a balanced approach leveraging California’s rich university resources and private-sector innovation for decentralized oversight and ongoing evaluation of AI tools. This strategy would enable rigorous ethical standards and technical accountability without sacrificing the transformative potential of AI technologies. Moreover, we express concern regarding the proposed standard’s language addressing “disparate impact.” While we unequivocally support anti-discrimination principles, we caution against a blanket prohibition that misunderstands how AI systems function; indeed, AI systems simply mirror historical biases present in legal precedents rather than generating discriminatory intent independently. We advocate instead for transparency in AI training processes, ongoing bias monitoring, and the mitigation of harmful outcomes, ensuring AI’s responsible integration within existing legal frameworks. Fortuna Arbitration believes deeply in AI’s potential to enhance the judicial process significantly. By adopting thoughtful, informed, and flexible regulatory frameworks, California can position itself as a leader in judicial innovation while upholding fairness, accountability, and accessibility for all. Fortuna Arbitration also recognizes that the psychology of humans revolves heavily around social accountability, consequences, and punishment as mechanisms to regulate behavior. People inherently understand that negative actions typically lead to repercussions, influencing their ethical and moral decisions. AI becomes unsettling precisely because it lacks this fundamental psychological restraint; it cannot experience guilt, fear consequences, or be genuinely punished. But the Judicial Council should also understand that these are real businesses, with real engineers and executives who stake their name to their product—entire enterprises defined by their opportunity to provide a quality service. Those people <i>do</i> care, are held accountable by the market, and can be independently regulated.	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			We respectfully submit that the Judicial Council reconsider its implementation of SP25-01.	
			The Fortuna Team [Biographical and contact information omitted by the AI Task Force.] ²	No response required.
			Comment on the Implementation of Rule 10.430 and Model Policy Generally Fortuna’s first comment is general opposition to the implementation of Rule 10.430 on the grounds of feasibility. The California judiciary is facing a \$97 million dollar cut—and during a time when funds are scarce, the judiciary should take extraordinary action to innovate, not regulate. Blain Corren, <i>Judicial Council Allocates Funding to Trial Courts with \$97 Million Required Cut</i> , https://newsroom.courts.ca.gov/news/judicial-council-allocates-funding-trial-courts-97-million-required-cut . Fortuna recommends that the Judicial Council instead focus on a decentralized regulatory regime, comprised of watchdogs from California’s rich university system, and training for judges to review technical recommendations. To do otherwise, would impose regulatory costs on the judiciary and the judiciary’s vendors that are not currently reasonable.	Please see responses to Fortuna Arbitration’s specific suggestions below.
			<i>Fortuna Recognizes the Potential Harm of Artificial Intelligence</i> Fortuna recognizes the potential harms posed by artificial intelligence that accompany increased efficiency. Recent incidents underscore the dangers of unvetted AI outputs – for example, a New York attorney was sanctioned after citing	The task force determined it is necessary to recommend adoption of a rule of court and a standard of judicial administration to address the confidentiality,

² The task force has omitted the commenters’ biographical information and redacted email addresses and phone numbers from this comment for privacy reasons.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			fictitious cases produced by ChatGPT. Sara Merken, New York lawyers sanctioned for using fake ChatGPT cases in legal brief, REUTERS (2023), https://www.reuters.com/legal/new-york-lawyers-sanctioned-using-fake-chatgpt-cases-legal-brief-2023-06-22/#:~:text=New%20York%20lawyers%20sanctioned%20for%20using%20fake%20ChatGPT%20cases%20in%20legal%20brief,-By%20Sara%20Merken&text=NEW%20YORK%2C%20June%2022%20(Reuters,an%20artificial%20intelligence%20chatbot%2C%20ChatGPT. Moreover, studies reveal that general-purpose AI chatbots “hallucinate” (produce false information) in 58%–82% of legal queries; and specialized legal AI chatbots, such as Lexis and Westlaw, can hallucinate 17% and 33%, respectively. Magesh et al., Hallucination-Free? <i>Assessing the Reliability of Leading AI Legal Research Tools</i> (Preprint, 2024), https://dho.stanford.edu/wp-content/uploads/Legal_RAG_Hallucinations.pdf. But attorney misconduct is not unique to AI—indeed, lawyers are frequently sanctioned in California for various misconduct. See The State Bar of California, https://apps.calbar.ca.gov/discipline. In Los Angeles alone, there were 8 disbarments in 2024. Across the state, attorneys are frequently disciplined for comingling client funds, and improper conduct. While bots do currently hallucinate, there is no reason that the Judicial Council should treat them any differently than a non-lawyer legal assistant. That relationship is already governed by Rule 5.3 of California’s Rules of Professional Conduct. <i>See Id.</i> (“Lawyers often utilize nonlawyer personnel, including secretaries, investigators, law student interns, and paraprofessionals. Such assistants, whether employees or independent contractors, act for the lawyer in rendition of the lawyer’s professional services. A lawyer must give such assistants appropriate instruction and supervision concerning all ethical aspects of their employment.”). There is no reason AI should be treated differently; just as if a paralegal invents or misrepresents a case still subjects his supervising attorney to liability, so does the	privacy, bias, safety, and security risks posed by use of generative AI in court-related work. The task force concluded that adopting the proposed rule and standard will help promote responsible innovation in court operations while protecting confidential information, ensuring appropriate oversight, and maintaining public trust in the judicial branch.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			AI system. The managing of risks in each individual case—as Rule 5.3 already does—should be appropriated by private parties.	
			<i>A Technical Review of AI</i> We borrow the following passage from our foundational whitepaper, <i>We Built Judge.ai. And You Should Buy It</i> . We intend the following to be purely educational about the structure and architecture of AI systems, which will support our comment on managing the risks. Artificial intelligence, or AI, is a term introduced in 1955 by John McCarthy, an emeritus professor at Stanford. It broadly describes “the science and engineering of making intelligent machines.” For many, AI is most familiar through applications like ChatGPT, which belong to a category of technologies known as “large-language models” (LLMs). These LLMs are a specific type of Deep Neural Network (DNN), which in turn are a subset of algorithms. For the sake of the lawyers who are still with us, let’s start with the basics. An algorithm is simply a set of well-defined instructions or procedures designed to accomplish a specific task. Mathematical algorithms, like most quantitative methods, function by correlating independent variables, or “features,” with dependent variables, or “labels.” Among the types of mathematical algorithms AI users are most likely to encounter are Deep Neural Networks (DNN). ... LLMs, likewise, are a form of DNN that employs a network architecture called a “transformer.” Transformers excel in handling text data and have become the foundation for many advanced models in natural language processing due to their ability to efficiently capture long-range dependencies via selfattention mechanisms—that is, using a transformer allows a model to pinpoint and	No response required.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			connect related pieces of information across long distances in text by simultaneously considering all words and assigning proper weight to their relationships, even those humans would not otherwise recognize. In an LLM, tokens are both the features and labels. Here’s the basic process: LLMs will take any given text, tokenize it into discrete units (tokens), transform these units into embeddings containing a numerical vector value, and then feed into the prediction model. The tokens from the source are the feature; the token in the next sequence (think ‘sentence’) of the model is the label. These tokens serve as the features for the LLM, providing the necessary input to understand and generate text. Contra to the very many LinkedIn hype posts, the LLM is not ‘learning’ the truth value of any of these instances as mediated by experience—instead, the text itself is its own source of truth. Said again, humans experience language in relation to their experience and purpose; the LLM learns language only in proximity to external language. The LLM is like teaching a parrot to say “come in” when you knock on the door, learning by mere association. LLMs thus only understand language as a function of the next probable token. Type in “Knock, knock,” and ChatGPT invariably replies, “who’s there?” Whether or not this relationship indicates truth is a subject of debate. The Judicial Council should thus recognize that LLMs like, ChatGPT, generate text by predicting the next most probable word based on patterns in language, not through experiential knowledge or an understanding of ‘truth.’ Suppose the following example, which one might make based on the prior understanding of LLMs: ... suppose you are a judge trying to make a decision (the output) based on a series of briefs; you have a group of clerks, each with a different subset of knowledge. Each clerk is a node, focusing on input information (the features)—such as case precedents or damage calculations. These clerks work collaboratively, sharing information that seeks a certain pattern of relationship. Each clerk then filters out information, repeating the process iteratively with	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			other clerks (the layers), until a consensus is reached on the most probable answer. For a decision-making DNN, the label might be a discrete outcome based on what the court should do—affirm, reverse, or dismiss, for example. Alternatively, for a LLM it might be predicting the next sequence of tokens, which once finished, will constitute a whole case, or in the case of arbitration, the whole dispute. But as we also note, this example can be misinterpreted by public officials: In fact, it is debatable whether LLMs are capable of “understanding” or reason as we know or conceptualize those terms at all. While clerks can interpret ambiguous information and adapt their reasoning to the nuances of a case, nodes process input data blindly, relying entirely on numerical patterns. This fundamental limitation necessitates the use of vast amounts of labeled data to train DNNs. Through repeated exposure to examples, nodes can approximate understanding by identifying patterns—but this process is far removed from the intuitive reasoning employed by human clerks. When managing risks, the Judicial Council would do well to consider AI tools nothing more than a series of complex mathematical functions. Those functions do not harbor—and cannot harbor— racial, gender, or lifestyle animus, outside of those biases built into the legal system itself. We should not hold AI tools to a standard higher than we would any human actor. Finally, it is also critical to recognize that “AI” is not a monolithic concept. Different types of AI systems function in fundamentally different ways, and a regulatory approach that lumps them all together can misfire. In particular, traditional predictive or classifier models differ greatly from transformer-based generative models (such as large language models, <i>LLMs</i>). Each comes with distinct technical features, use cases, and risk profiles. SP25-01’s focus on generative AI should be calibrated so as not to inadvertently envelop or stifle other AI-driven	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			tools that courts already use or may soon adopt for legitimate purposes—especially when classifier models are used to self- regulate the outputs of LLM models.	
			<i>Decentralized Oversight—Universities</i> The Judicial Council of California should establish a decentralized oversight program for generative AI tools by partnering with multiple California universities (for example, Stanford, UC Berkeley, and others in the UC system). Rather than creating a new centralized AI oversight unit within the court system, the Council would invest in academic expertise to perform ongoing monitoring and evaluation of AI technologies used by the courts. These academic partners – drawing on their faculty, research staff, and advanced students in law and computer science – would operate as independent centers of excellence that collectively ensure AI tools meet the judiciary’s standards for fairness and accuracy. And indeed, California universities are already leading the race in these programs—as cited above, universities such as Stanford are already the best of AI legal research in the country. Under this proposal, the Judicial Council would coordinate a consortium of university partners, each tasked with a specific scope of work. The Judicial Council should employ these preexisting resources, instead of creating new obligations for state courts. Key responsibilities of the academic partners would include: • Bias Audits of AI Tools: Regularly examining and identifying potential legal biases in generative AI systems that might impact court users. Universities would design tests (similar to audit studies) to detect racial, gender, or other biases in AI outputs. Findings would be documented and reported to the Council, enabling preemptive mitigation of any biased behaviors before they affect judicial decisions. Independent validation by different institutions will ensure that bias detection is thorough and credible. Many are already doing so.	This suggestion is beyond the scope of the current proposal, but the task force may consider it as time and resources permit.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			• Performance Benchmarking & Reliability Testing: Testing and documenting the performance of generative AI tools on tasks relevant to court operations. This includes evaluating accuracy, tendency to hallucinate or err, and the tools’ effectiveness in tasks like legal research, drafting orders, summarizing documents, and language translation. For example, a university partner might benchmark an AI legal research assistant against known case law queries to measure its error rate (as Stanford researchers recently did, listed above). • Risk Assessment & Best-Practice Guidance: Serving as an advisory panel to the Judicial Council on the risks, safeguards, and best practices for safe AI adoption. The academic partners would stay abreast of the latest developments in generative AI (a fast-moving field) and issue guidance on issues like data security, confidentiality, and appropriate use- cases for AI in courts. They could flag emerging concerns (e.g. a new form of AI-generated deepfake evidence) and recommend policy responses. They would also help develop educational programs for judges and court staff—demystifying AI and training personnel on how to use these tools cautiously and effectively. Crucially, this partnership model is decentralized: oversight duties are distributed across multiple institutions rather than vested in a single new government office. Each university partner would work independently within its expertise area (for instance, one might focus on bias and ethics, another on technical performance and security, etc.), and the Judicial Council would aggregate their insights. The Council could formalize this through memoranda of understanding or research grants that outline deliverables (e.g. annual bias audit reports, quarterly AI performance briefs, on- call advisory services for emergent issues, etc.). By empowering external experts, the Judicial Council can ensure that oversight keeps pace with innovation. This approach mirrors strategies already being embraced in California’s broader AI governance: the state is encouraging collaborations between government, academia, and industry to respond to AI’s rapid evolution. It recognizes that California’s top	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			universities are an invaluable resource for public sector innovation, especially in a field as complex as AI. It is notable that universities are already working with the court system to effectuate this process: in Los Angeles, a Law School team is already implementing several AI systems, including an automated ‘default prove-up’ system to “review default judgements,” and a “referral tool” to help connect pro se litigants to legal tools. Stanford University, Human-Centered Artificial Intelligence, https://hai.stanford.edu/news/harnessing-ai-to-improve-access-to-justice-in-civil-courts.	
			<i>Decentralized Oversight—Private Institutions</i> California has one of the largest startup ecosystems—and for comparably minimal costs, private enterprises can work with the Court system to deploy automated systems across every internal vertical. This is already occurring. But more importantly, companies (like Fortuna) can assist the Judicial Council in providing comprehensive, peer-reviewed, and open-sourced analysis of their systems. If the court system must incorporate AI regulations, it is more efficient to price those regulations into each bid, with individual procurement officers placing those conditions into state contracts. Not all AI systems—including generative—need extensive (expensive) oversight. Some systems—like automated shepardization—work on an expedited workflow that only checks certain conditions. Others, like sentencing algorithms, clearly do, as they pose a high likelihood of imposing social costs if improperly deployed. The Judicial Council would do well in establishing a contractual framework for procurement officers to analyze bids, as opposed to imposing broad regulatory costs across the court system.	This suggestion is beyond the scope of the current proposal, but the task force may consider it as time and resources permit.
			Comment on Discrimination	Please see responses to Fortuna Arbitration’s comments in the

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			Fortuna’s second comment is on the policy of regulating AI “unlawful discrimination.” Fortuna argues that new binding rules regarding discrimination that target AI are unnecessary and potentially redundant. California’s courts are already bound by robust constitutional, statutory, and ethical constraints that prohibit bias, discrimination, and misuse of private information. These existing legal frameworks inherently extend to any tools courts use—including generative AI—ensuring accountability and fairness without the need for technology-specific mandates. Overly specific AI regulations could create confusion, stifle beneficial innovation, and imply a false gap in oversight where none truly exists. Instead of new binding rules, the Council should issue flexible best- practice guidelines on AI use, allowing courts to adopt new technology responsibly under the umbrella of existing law. This approach safeguards core principles (impartiality, privacy, and integrity) while promoting innovation and adaptability. The analysis below outlines the current legal and ethical constraints on court conduct and explains why they sufficiently govern AI usage in the judiciary.	“Policy Recommendation” section below.
			<i>Existing Regulation for Discrimination</i> From a broad state constitutional angle, guarantees of due process and equal protection apply fully to the judicial branch. <i>Article I, Section 7(a)</i> provides that no person may be denied equal protection under the laws by the state. This broad mandate means that courts, as state actors, must not treat individuals differently based on protected characteristics, whether decisions are made by humans or by relying on an AI tool. In addition, <i>Article I, Section 1</i> enshrines privacy as an inalienable right, added specifically to guard against modern threats to personal privacy aclunc.org. These constitutional provisions create a baseline: any court practice—including use of generative AI in proceedings or administration—that results in unlawful discrimination or undue intrusion on privacy would violate fundamental law.	Please see responses to Fortuna Arbitration’s comments in the “Policy Recommendation” section below.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			In terms of statutory requirements, California state law already forbids discrimination in any state- operated program on the basis of characteristics like race, religion, national origin, sex, disability, or other protected categories. Cal. Gov. Code, sec. 11135 (2024). This law applies to state entities and programs, which encompass the courts. Finally, the judiciary already has sufficient regulations to cover concerns over generative AI. Judicial officers are bound by the California Code of Judicial Ethics, which imposes explicit duties to prevent bias and ensure fairness. Canon 3 “Performing the Duties of Judicial Office Impartially and Diligently” states that “A judge shall perform judicial duties without bias or prejudice.” This canon also requires judges to demand similar impartial conduct from others under their authority: “A judge shall require order and decorum in proceedings” and must ensure all staff and court personnel under the judge’s direction uphold the same standards. This canon also requires judges to demand similar impartial conduct from others under their authority: “<i>A judge shall require order and decorum in proceedings</i>” and must ensure all staff and court personnel under the judge’s direction uphold the same standards.	
			Comment on Disparate Impact SP25-01 is the Judicial Council’s proposal to adopt Rule 10.430 and Standard 10.80, establishing policies for the use of generative AI in court-related work. Under the proposed Standard 10.80 (applicable to judges in their adjudicative role), judicial officers: “Should not use generative AI to unlawfully discriminate against or disparately impact individuals or communities” based on membership in any protected class (i.e. any classification protected by federal or state law). The intent behind the <i>disparate impact</i> language is clear and laudable: to prevent AI from injecting bias or causing unfair outcomes in the administration of justice. The	Please see responses to Fortuna Arbitration’s comments in the “Policy Recommendation” section below.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			concern arises from well- documented cases of algorithms exhibiting bias against racial minorities, women, or other protected groups. The Council’s proactive stance recognizes that if generative AI is to be used in courts, it must not undermine principles of equal justice. However, the phrasing of the clause raises practical questions and demonstrates a misunderstanding of the technical capabilities of AI: generative AI itself has no consciousness or intent—it produces outputs based on preexisting patterns in its training data. If that training data (e.g. decades of case law or statutes) contains historical biases or systemic disparities, the AI’s outputs may reflect those patterns. In such a scenario, is the AI “unlawfully discriminating,” or is it simply mirroring biases inherent in the law? The distinction is crucial for policy. While we agree that unlawful discrimination should be prohibited, the Judicial Council would be well-advised to consider the potential harm of a blanket <i>disparate impact</i> prohibition. Instead, the Judicial Council should aim to introspectively examine preexisting harmful policies in California; correct those policies; and expressly change the system, instead of attempting to reduce mere mathematical functions representing the current system. Most importantly, this regulation creates a perverse incentive: because AI has no intent and is largely a reflection of previous training data, the only way to accommodate statutes that already cause disparate impact would be to manually weigh the model in favor of certain discriminated against parties (at least, for more traditional algorithms, like classifiers. For transformer models, such as LLMs, reinforcement training or simple prompting would likely suffice). To solve the social harm of disparate impact, the legislature should change the relevant statutes, and the judiciary should overturn the relevant caselaw. Likewise, because the judiciary should judge each case on its individual merits, the council should not adopt this rule.	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			<i>Impartiality in the Legal System</i> It is an uncomfortable reality that the legal system—despite ideals of impartiality—has produced disparate outcomes for different communities over time. Numerous legal precedents and studies demonstrate that <i>protected groups (by race, ethnicity, gender, etc.) have faced unequal treatment or outcomes in practice</i>. These disparities form part of the corpus of “ground truth” on which any legal AI would be trained: Racial Bias in Sentencing and Death Penalty: The U.S. Supreme Court in <i>McCleskey v. Kemp</i> (1987) was presented with a rigorous statistical study of Georgia’s death penalty. The data showed defendants accused of killing white victims were 4.3 times more likely to be sentenced to death than those accused of killing Black victims. 481 U.S. 279, 287. Unadjusted figures were even starker: capital sentencing rates in cases with white victims were almost 11 times higher than in cases with Black victims. <i>Id.</i> at 326 (J. Brennan, Dissenting). The Court acknowledged a “discrepancy that appears to correlate with race” but ultimately declined relief, essentially reasoning that some level of racial bias in sentencing was “inevitable.” <i>Id.</i> at 312. This precedent chillingly illustrates that systemic racial disparities in outcomes have been long recognized yet tolerated in law. An AI trained on the body of criminal case law, especially older cases, will inevitably read countless opinions that reflect or even accept such disparities. Disparities in Pretrial Decisions (Bail and Indictment): Racial bias is not confined to sentencing – it can appear at the very start of a case. A 2023 empirical study of over 43,000 felony cases in New York City found that Black defendants faced higher bail amounts and were more likely to be indicted than similarly situated White defendants. Connor Concannon and Chongmin Na, <i>Examining Racial and Ethnic Disparity in Prosecutor’s Bail Requests and Downstream Decision-making</i>, 16 Race and Social Problems	Please see responses to Fortuna Arbitration’s comments in the “Policy Recommendation” section below.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			1, 1 (2022), pmc.ncbi.nlm.nih.gov pmc.ncbi.nlm.nih.gov. These early decisions had “<i>significant indirect effects</i>” that cumulatively contributed to unwarranted racial disparities in pretrial detention and case outcomes. <i>Id.</i> Even though some later stages showed mixed effects, the study confirms that at critical decision points, race was a factor in practice. The judicial decisions and prosecutorial requests recorded in such data carry forward a legacy of disparate impact that an AI might learn as “normal” patterns of how bail is set or who gets indicted. The Judicial Council is likely aware of further studies proving the general point that the law is already embedded with elements of disparate impact. <i>Bias and disparate impacts are already woven into the fabric of our case law and legal data.</i> Race, gender, and other protected characteristics have, in various ways, influenced outcomes in courts. These influences might be unjust, yet they exist as part of the “ground truth” of the law as written and applied. Any AI ingesting tens of thousands of judicial opinions will inevitably learn patterns reflecting these disparities. This is not to say the law is <i>only</i> bias—of course, the law also contains neutral principles and protections against discrimination – but the historical record is mixed, and an AI does not inherently know which patterns are the “bad” biases to avoid and which are valid legal rules. It simply learns what it is given. Empirically, if that heritage includes systemic disparities, the model will learn them. A 2024 study by Bozdag et al. found that Legal-BERT (a transformer model tuned for legal language) “inherits” gender bias from its training data, which included U.S. and EU case law blog.genlaw.org. Similarly, Sevim et al. (2023) concluded that legal text corpora contain significant gender bias across countries, and NLP models trained on those corpora reflect that bias blog.genlaw.org. These findings confirm that biases present in source data are picked up by AI. The model does not independently concoct stereotypes – it statistically <i>learns</i> them from the patterns in the text. For example, if many judicial opinions subtly associate women with certain roles (or minorities with crime, etc.),	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			a text-generating AI will likely reproduce such associations in its outputs unless corrective measures are taken. Given the above, we can draw a critical distinction: Generative AI tools replicate patterns; they do not originate policy. Any <i>disparate impact</i> observed in an AI’s outputs is traceable to some pattern in its training data—which, in the case of legal AI, is our body of law and legal practice. If an AI tool used by a court produced an outcome that disproportionately affects a protected group, it is almost certainly because that outcome aligns with a precedent or rule in the training data that had the same effect. The AI did not on its own decide to treat one group worse than another; it statistically inferred that outcome from how similar cases have been treated or discussed in the legal record. This dynamic is why the “unlawfully discriminate or disparately impact” clause, as an oversight mechanism, could be problematic. It implicitly treats the AI as a potential <i>actor of discrimination</i>, akin to a human who might choose to apply a prejudiced rule. But the AI’s “choices” are just regurgitations of human choices embedded in data. In other words, if we detect that an AI tool’s use is causing a disparate impact, that is likely symptomatic of an underlying bias in our laws or precedents. The AI is highlighting it, not independently creating it. From a policy perspective, this suggests a need to focus on <i>data bias and outcome monitoring</i> rather than simply forbidding the tool from producing any biased result. We must ask: Is it fair or useful to hold the AI to a higher standard than the source material it learned from? If even our current human judges and juries—bound by existing law—produce disparate outcomes, expecting an AI trained on their outputs to somehow not produce disparate outcomes might be unrealistic without further intervention. In fact, a rigid application of the <i>disparate impact</i> prohibition might perversely result in banning AI tools that are merely truthfully reflecting the state of the law.	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			<i>Legal Considerations</i> The language “<i>unlawfully discriminate or disparately impact</i>” lumps together two related but distinct concepts from discrimination law. “<i>Unlawful discrimination</i>” is straightforward – it refers to intentional disparate treatment or policies that explicitly violate anti-discrimination laws. No one would argue AI (or any tool) should be allowed to do that. However, “<i>disparate impact</i>” in legal doctrine refers to when a facially neutral practice has a disproportionate adverse effect on a protected group, even without discriminatory intent. <i>See Rosenfeld v. Abraham Joshua Heschel Day Sch., Inc.</i>, 226 Cal. App. 4th 886, 893 (Cal. App. 2014) (“Disparate impact exists where, “<i>regardless of motive</i>, a facially neutral employer practice or policy, bearing no manifest relationship to job requirements, in fact had a disproportionate adverse effect on members of the protected class.”). Under civil rights statutes (like Title VI for recipients of federal funds, or Title VII in employment), a disparate impact is not automatically illegal; it triggers an analysis of justification and a possible less-discriminatory alternative. In other words, not every disparity is “unlawful” – the law tolerates some disparities if they flow from legitimate practices and no less-biased alternative is feasible. The Supreme Court has even noted that requiring the elimination of all racial disparities in criminal justice would “throw into question the principles that underlie the entire system.” By stating generative AI “<i>should not... disparately impact individuals or communities</i>”, the proposed policy risks holding AI to a near-zero-tolerance standard for any disparate outcome. That goes beyond how we treat most policies implemented by human actors. For instance, if a new judicial procedure for setting bail was found to inadvertently result in more detentions of indigent defendants (who might disproportionately be from certain racial groups), courts would carefully study if the procedure is justified or could be improved—but they might not automatically discard it unless it violates the law. With AI, the current wording suggests <i>any</i> disparate impact is unacceptable, even if the AI is faithfully following	Please see responses to Fortuna Arbitration’s comments in the “Policy Recommendation” section below.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			existing law. This could lead to a paradox: an AI tool could be deemed violative of policy for echoing the very disparities our legal system has produced (which themselves might not have been deemed unlawful enough to change by courts). Additionally, enforcing this clause raises challenging questions: How will courts determine if an AI tool has a disparate impact? What metrics or evidence will be used? Unlike a hiring algorithm where one can compare selection rates by race or gender, a generative AI might be used for varied tasks – summarizing a case, suggesting a sentence, drafting an order. The <i>impact</i> of those usages on communities is indirect and may be hard to isolate. Would we examine the AI’s outputs over time for statistical bias? Or scrutinize a particular case outcome influenced by AI to see if it harmed a protected group? This is a nebulous area. There is a risk of over-deterrence: courts might avoid using AI at all for fear that any mistake or any appearance of bias could put them in violation of Rule 10.430/Standard 10.80. Such fear could rob the judiciary of efficiency gains and consistency that AI could offer in appropriate tasks.	
			<i>Policy Recommendation</i> 1. Emphasize “Compliance with Anti-Discrimination Law” and Intent: The rule should make clear that AI must not be used in ways that <i>violate existing anti-discrimination laws</i>. For example, “<i>Generative AI must not be used to engage in unlawful discrimination (such as basing decisions on protected characteristics in a manner prohibited by law)</i>.” This covers the intentional or direct misuse of AI to target a protected class – which is clearly unacceptable. It also aligns with how courts understand discrimination (e.g., an AI should not be programmed with different rules for different races, etc.). This framing is stronger on <i>unlawful conduct</i>, and avoids the ambiguity of “disparate impact” by itself.	The task force is not recommending revisions in response to this suggestion. The task force concluded that the requirement not to “use generative AI to unlawfully discriminate” is the clearest way to inform users that they must keep their legal obligations in mind when using generative AI. The task force concluded that identifying specific potential means or forms of discrimination in the rule and standard could be

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
				read to exclude certain unlawful behavior not specifically listed.
			2. Replace or Qualify “Disparately Impact” with “Identify and Mitigate Bias”: Instead of a flat prohibition on any disparate impact, the policy could require that courts assess and mitigate potential biases in AI tools. For instance: <i>“Courts should evaluate generative AI outputs for potential biased or disparate patterns affecting any protected group, and take appropriate remedial action if such patterns are detected.”</i> This shifts the role to one of vigilance and correction. It acknowledges that some bias may emerge (since the AI is trained on imperfect data), but insists that courts be proactive in catching it – much as the proposed standard already says judicial officers “should review generative AI material... for biased, offensive, or harmful output” courts.ca.gov . The difference is we treat disparate impact as a risk to manage and minimize , not an on/off switch that disqualifies the AI altogether. This approach is akin to how agencies handle disparate impact under civil rights laws: identify if it’s happening, then adjust the practice or provide justification and seek less discriminatory alternatives justice.gov justice.gov . A court could similarly adjust how it uses AI or require tweaks in the AI system if a bias is found.	In light of all the comments received on this issue, the task force is recommending revisions to rule 10.430(d)(3) and (4) to require users to take reasonable steps to verify, correct, or remove inaccurate or biased material. The task force is recommending similar revisions to standard 10.80(b)(3) and (4).
			3. Require Transparency from AI Tools Regarding Training Data: To better align understanding, the policy could mandate that any AI tool used by courts document its training sources and known limitations. For example, a provision might state: <i>“Any generative AI system adopted should come with documentation of its training data (e.g., corpus of case law, statutes) and any bias testing performed. Courts should favor tools that have undergone bias audits and that allow for human interpretability of their outputs.”</i> This doesn’t appear in the current draft but would greatly help in oversight. If we know the AI was trained, say, on all California appellate cases from 1850–2024, we can anticipate that older cases in that set may carry historical prejudices, and we can guide users accordingly (or even filter out certain eras or terms). This addresses the issue upstream, by acknowledging bias in training data and demanding	The task force is not recommending revisions in response to this suggestion. The task force determined that imposing this requirement would prevent use of most generative AI systems, including those developed by trusted legal research providers. Additionally, this suggestion imposes a higher bar for generative AI research tools than other tools. For

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			clarity. Such transparency can inform judges and staff to be skeptical of certain outputs and cross-check them.	example, standard legal research systems do not warn about or filter case law from the 1800s.
			4. Continuous Monitoring and Feedback Loop: Finally, the Council might consider adding a requirement that the use of AI in courts be continually monitored for impacts on different communities , with periodic reports or audits. For instance: <i>“The Judicial Council (or a designated committee) should periodically review the effects of generative AI use in court operations, including any evidence of disparate impacts on litigants or communities, and update policies or training as needed.”</i> This creates an ongoing oversight mechanism. If a pattern emerges where AI-assisted decisions appear skewed, the Council can take targeted action – maybe adjusting the tool or limiting its use in that context. This dynamic approach is more adaptive than a static prohibition and acknowledges that our understanding of AI bias will evolve. It also signals to communities that the judiciary is not complacent – it is actively watching for and addressing any inequitable outcomes.	This suggestion is beyond the scope of the current proposal, but the task force may consider it as time and resources permit.
			Fortuna writes today to recommend a change of the language “prohibit the use of generative AI to unlawfully discriminate against or disparately impact individuals...” to strike the “disparately impact.” Proposed Rewritten Clause: In light of the above, a potential rewrite of the contentious clause in SP25-01 could be: <i>“Generative AI tools must be used in a manner consistent with all anti-discrimination laws. Courts and judicial officers should not rely on AI outputs to make decisions that would violate these laws or unjustly discriminate against individuals or groups. If a generative AI system produces recommendations or content that reflect historical biases or result in disparate impacts on protected classes, users should correct those biases, where expressly stated. Users should also make a conscious effort to promote Generative AI accountability and report any express discriminatory behavior. Courts shall take steps to mitigate</i>	Please see previous responses.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			<i>any such bias, including reviewing AI outputs for fairness and adjusting use or policy as necessary to prevent unjust outcomes.”</i> Such language maintains a strong stance against discrimination (no one wants AI that produces unjust results) but shifts the focus to <i>mitigation</i> and <i>responsibility</i> rather than an outright bar. It implicitly accepts that some bias may surface (since it references “if... reflect historical biases”) but demands action when it does. This is more realistic and still aligns with the Judicial Council’s ethical mandate. It treats AI similar to how we treat a junior clerk or an advisory guideline: a helpful input that must be checked and that must not be followed if it would lead to illegal or inequitable results.	
			Conclusion We wanted to end our comment with a sort of ‘legal inspiration.’ Right now, when you think of a “law”, you think of words. It’s a paragraph; it has subparts; it’s got one of those § things. And you’re right—but not for long. In short order, the law is going to be an AI. And there’s a simple reason for that; to repeat our founding words: <i>Predictable, efficient, and cheap—very cheap.</i> Let us repeat that: The law is going to <i>be</i> an AI. We mean that literally. The AI will not merely read the statute and estimate what it means. It will <i>be</i> the <i>statute</i>—a statute that you can talk to, a statute you can ask for legal advice. And the Judiciary—celebrating California’s innovation and output-driven government—should embrace this. Importantly, AI reflects the current system—if that system is discriminatory or has incidental disparate impacts, the legislature should either (1) change that system or (2) create social policies to minimize the incidents of disparate impacts. If sentencing has a racially disparate impact, then the solution is to solve the	No response required.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			underlying issues of crime and poverty—not to engage in an increasingly popular luddite sentiment. AI is not to blame for bad law. <i>Californians have always been innovators—we must remain that way.</i> This is inevitable — and Fortuna-Insights is seeing it firsthand. We built Arbitrus, an AI arbitrator that adjudicates contract disputes by the parties’ express stipulation. When Arbitrus performs this function, it’s really not acting as a judge at all—it <i>is</i> the contract. Its statements are not opinions; they are truths about the contractual arrangement, and the parties have stipulated that the AI is always right. When they don’t like the answer, they don’t yell at Arbitrus; they change the contract—thus reprogramming the bot. And this dynamic doesn’t just play out on the back end. With Arbitrus’ predictive function, parties can ask it for legal advice <i>before</i> they’ve even acted. It’s a revolutionary system — and we see it work for California companies every day. Make no mistake: That basic system is coming to state and federal judiciaries (and already is). As such legislators need to approach legal AI not as a “tool” for judges to use; it will start life as such a tool, but almost certainly will not end that way. When the machine perfects to the point that judges trust it implicitly and are right to do so, legal AI will represent nothing short of a new form of de facto government that collapses the judiciary and legislature into one hybrid entity. That’s future is what California needs to be planning for now. Because it’s not fifty years away. It might not even be <i>two</i>. Blink once? It’s out from under you.	
8.	Mark G. Griffin, Esq. Attorney and Interim Chair of California Lawyers Association, Law Practice	AM	As drafted, the Judicial Council of California’s Artificial Intelligence Task Force’s (“Task Force”) proposed rules, Cal. Rules of Court, rule 10.430, and Cal. Standards of Judicial Administration, standard 10.80, are well-written proposals. However, as drafted, the proposed rules overlook the biggest bias in artificial intelligence: human bias. Artificial intelligence is only as good as the human controlling it. To address	The task force appreciates the commenter’s concerns but is not recommending revisions in response to this suggestion. Although the task force agrees

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
	Management and Technology Section Hanson Bridgett LLP San Francisco		human bias, the Task Force should consider inserting the following language after Cal. Rules of Court, rule 10.430(d)(2): “Require court staff and judicial officers who generate or use generative AI material to review generative AI prompts for biased, offensive, or harmful input.” Additionally, the Task Force should consider inserting the following language after Cal. Standards of Judicial Administration, standard 10.80(b)(2): “Should review generative AI prompts, including any generative AI prompts prepared on their behalf by others, for biased, offensive, or harmful input.”	that court staff and judicial officers should be aware that biased prompts can lead to biased outputs, the task force is concerned that requiring prompts to be unbiased could make it more difficult for judicial officers and court staff to perform certain tasks. The task force concluded that this issue can more appropriately be addressed through education and guidance materials.
9.	Justin Xavier Howe Information Security/Security Operations Judicial Council of California San Francisco	NI	Suggested revision to JCC GenAI Policy The use of non-creative AI models should always require disclosure of the false-positive-rate in Judicial Applications Every statistical model, machine learning model, and Generative AI model has a false positive rate (or equivalently, a hallucination rate). This rate should be explicitly disclosed and documented in all judicial applications of this model, so that the sufficiency of such evidence can be evaluated. A model that exhibits a 10% false positive rate must be handled differently than a model exhibiting a 0.01% false positive rate within the Judiciary. The primary aim of this suggestion, is to remind Judicial Officers that every statistical model generates false positives. Warning: The JCC does not disclose these false-positive-rates in current publications, nor related ‘goodness-of-fit’ measures in the statistical analysis that it conducts.	The task force appreciates the commenter’s concerns but is not recommending revisions in response to this suggestion. The task force determined that imposing this requirement would prevent use of most generative AI systems, including those developed by trusted legal research providers.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			Pertinent Citations • “<i>We know AI [writing] detectors don’t work (DO NOT USE THEM), so I hear about instructors using their gut to figure out who is cheating.</i>” – Ethan Mollick ○ https://x.com/emollick/status/1699517598772125842 • “<i>There are at least seven confirmed cases of misidentification due to facial recognition technology, six of which involve Black people who have been wrongfully accused.</i>” ○ https://innocenceproject.org/news/artificial-intelligence-is-putting-innocent-people-at-risk-of-being-incarcerated/	
10.	Hon. Curtis Karnow Judge, Superior Court of California, County of San Francisco	N	I write on my own behalf and not that of my court. I discuss these issues: (1) whether any rule or standard is needed at this time; (2) the definitions and terms used; and (3) recommendations on training and alerts which may address the Committee’s underlying concerns. I recommend against issuing the rule and standard. 1. It may be awkward for a committee at least implicitly charged with developing new rules to decide that new rules are not worth the candle. But I suggest there is no current need for the new rule and standard; there are plenty of rules already, including those that address the underlying concerns of the current proposal. Courts and individual judges around the country have thought it necessary to come up with a plethora of rules, almost all of them useless.[1] The rules are a mixture of warnings and alerts, inconsistent definitions of AI, arbitrary barriers to using AI, requirements that the use of AI be noted without an apparent rationale; and, mostly, restatements of extant duties (such as that lawyers shouldn’t fill briefs with made-up cases).	The task force appreciates the commenter’s concerns but determined that it is necessary to recommend adoption of a rule of court and a standard of judicial administration to address the confidentiality, privacy, bias, safety, and security risks posed by use of generative AI in court-related work. The task force determined that adopting the proposed rule and standard will help promote responsible innovation in court operations while protecting confidential information, ensuring appropriate oversight, and maintaining public trust in the judicial branch.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			[1] See generally, <i>The Duke Project</i>: https://rails.legal/resource-ai-orders/. These orders are generally targeted at lawyers, not court personnel. I suggest that rules should be reserved only for behaviors which may lead to sanctions. Alerts and cautions do not fit that bill, and they erode the signal we try to send when enacting a rule. Standards of judicial administration certainly need not meet that high bar; but should not issue if they duplicate existing duties and otherwise are unlikely to change people’s behavior. The rush to have rules in this area parallels the rush among businesses to have “AI-enabled” services and products. Some of those efforts are just old products in new packages. The desire to be fashionable risks being unfashionably dated in the near future.[2] Courts, and especially the Judicial Council, should be wary of contributing to this. [2] For a blunt discussion of the hype, see e.g., Ed Zitron (Feb. 17, 2025) at https://www.wheresyoured.at/longcon/. Perhaps an animating concern here is to ensure that users of AI, such as court staff and judges, are aware of the risks. This sort of concern does not need a rule or standard. This sort of concern can be addressed with statements, alerts, and training for judicial officers and staff. Indeed, training and alerts are likely to be far more useful to users. There are real dangers in using AI products,[3] which caution against reliance on the systems, at least the current (early 2025) crop of them. And as this fast-moving area evolves, new alerts and modifications to training would be able to react on a time-scale far, far shorter than the process of enacting and revising new rules and standards. [3] I briefly describe some of them at https://works.bepress.com/curtis_karnow/70/. [This document, last edited earlier this year, is already out of date in some respects.] One of the central issues for developing AI is the alignment problem, and in that	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			connection it has become increasingly clear that the risks of AI include, for example, cheating and deception. <i>Alignment Faking In Large Language Models</i>, https://time.com/7259395/ai-chess-cheating-palisade-research/; [Chrome extension redacted]³ https://arxiv.org/pdf/2412.14093 discussed in https://www.forbes.com/sites/craigsmith/2025/03/16/when-ai-learns-to-lie/. We don't know enough yet to formulate a response to actual, pervasive, problems in California state courts: which is the usual backdrop for new policies and rules. I understand wanting to get out in front of technology, but that goal is best addressed by having this Committee act as the eyes and ears of the state judiciary, collecting both legal and technology developments, identifying the specific risks and solutions, and publicizing those to people who work in our courts.	
			2. The comments below apply to both the proposed rule and standards.	No response required.
			Definitions <u>“Public AI system”</u> -The definition can be read to include as a ‘Public AI system’ access by contractors for training purposes when they work on a private (internal) system. They are likely bound to keep data confidential. So is this what the “drafters mean? -How will users know if the system is ‘publicly available’? There are currently products which are available both as public and private systems. Furthermore, it very likely that in the near future AI will be embedded in larger applications and systems, in effect hidden from the view of the user, who will have no idea if the system is public or private. For example, one may imagine a seemingly ordinary auto-complete feature in a word processor which is guided by AI.	In light of all the comments received on this issue, the task force is recommending revised definitions of “generative AI” and “public generative AI system” in the rule and standard. The task force recommends deleting the definition of “artificial intelligence” in the rule and standard because it is not needed due to the proposed revisions to the definition of “generative AI.”

³ The task force has redacted a link to a Google Chrome extension.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			<u>“Artificial Intelligence”</u> I do not envy the committee the task of defining this term. The definition here is not useful because the notion of “typically requiring human intelligence” is vague. We don’t know if this applies to old-fashioned Westlaw research (reading a lot of cases and knowing which ones have certain words could be said to typically require human-like intelligence), or the relatively newer tools, marketed as “AI”-enabled by both Westlaw and Lexis, which are frequently used as ‘super’ search tools.[4] Human intelligence is also used to spell-check, but of course we don’t mean to include that sort of tool in this definition. [4] E.g., https://legal.thomsonreuters.com/en/products/westlaw-edge. See https://nbi-sems.com/blogs/news/lexisnexis-and-westlaw-will-launch-ai-legal-research-tools. [The deeper reason why definition is difficult here: As technology advances and takes over tasks typically done by humans, the tasks no longer become implicitly defined as what humans can do. For example, both the games of chess and Go were (at different times) considered to be the exclusive domain of human thinking as <i>contrasted with</i> the capabilities of computers. Then as computers took over the top rankings in these games, excellence at the games no longer became part of the human definition. So too with AI and the law: as computerized systems get better than humans in tasks (from spell checking to reviewing documents for relevance and privilege)—excellence at those tasks is no longer thought to be an especially human ability or part of the ‘human’ definition.] The Committee’s definition of ‘artificial intelligence’ in other words might refer to all currently available “AI” products, or none of them, of some of them. The definition provided is also confusingly close to the definition of artificial general intelligence [AGI], a system which can perform all intellectual tasks	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			humans can do, and in many cases exceed human performance. The Committee I am sure does not mean to invoke AGI: AGI does not currently exist, and there are different views when if ever it will arrive. An informed opinion predicts its arrival in 2026 or 2027.[5] The Committee may well have very different views about the risks of AGI than it does about the current crop of AI products. [5] Kevin Roose, “Powerful A.I. Is Coming. We’re Not Ready,” <i>The New York Times</i> (March 14, 2025), https://www.nytimes.com/2025/03/14/technology/why-im-feeling-the-agi.html. <u>Generative AI</u> This definition exhibits the difficulty of having a rule or standard that would be useful for more than a few months, because the technology changes rapidly. The “only” clause here excludes most current AI products, because these products not only use data on which the product was trained but also, on a prompt-by-prompt basis, they reach into the internet {and other sources, especially in private systems}.[6] This current crop of products is excluded by your definition. [6] E.g., https://textcortex.com/post/ai-chatbots-with-web-browsing; https://community.openai.com/t/chatgpt-can-now-access-the-live-internet-can-the-api/401928 Another example: LLMs are no longer the sole inhabitant of the generative AI space:- we have small language model as well which are highly useful. Nor should the Committee think that LLMs are the last word in AI; within a year or so the structure of AI programs might well be wholly different [7] and not properly described as “generative AI” at all.[8] [7] G. Scali, “Exploring the Future Beyond Large Language Models” (12 July 2023), https://thechoice.escp.eu/tomorrow-choices/exploring-the-future-beyond-	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			large-language-models/#:~:text=Beyond%20Large%20Language%20Models:%20The,learning%2C%20and%20meta%2Dlearning; Vivek Wadhwa, “The next wave of AI won’t be driven by LLMs. Here’s what investors should focus on instead,” <i>Fortune</i> (Oct. 18, 2024), https://fortune.com/2024/10/18/next-wave-ai-llms-investor-focus-tech/ [8] https://www.linkedin.com/pulse/symbolic-ai-generative-whats-difference-darren-culbreath-mkvp/#:~:text=Generative%20AI%20%2D%20Creating%20Novel%20Content,music%2C%20or%20even%20writing%20stories. 	
			<u>Disclosure of substantial portion</u> It’s unclear how this applies. Generally the AI generated output is used as a first draft; sources are checked (to the extent possible) and the draft is edited; perhaps every sentence ends up with at least a light edit, or some bits are deleted and others added. Does the end product contain a “substantial” portion? “Substantial” might not refer to a percentage of the final text, but rather to the <i>material</i> parts of a text, such as when a certain input (e.g. from an AI) is the source of the most important chunk of the final product, although reflected only in a small percentage of the verbiage. Is that what the Committee means?	In light of all the comments received on this issue, the task force has recommended a revised version of rule 10.430(d)(5).
			<u>Rule application</u> The rule only applies if a court permits the use of generative AI. Most courts neither permit nor not permit it, and that can be expected to continue. So the rule then doesn’t apply; is that correct?	The task force agrees that rule 10.430, as originally proposed, inadvertently excluded courts that do not take a position on use of generative AI. The task force recommends that rule 10.430(b) read as follows: “Any court that does not prohibit the use of generative AI by court staff or judicial officers must adopt a generative AI use policy. This

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
				rule applies to the superior courts, the Courts of Appeal, and the Supreme Court.”
			3. The core underlying concerns of the Committee are already handled by the current regime of rules, laws and Canons. Surely court staff and judges know that they should not publicize personal information like social security numbers; they know that they cannot be biased or discriminate against people in various groups; they know that their products should be accurate. They certainly know they are to comply “with all applicable law,” and so on. None of that is really at issue, and prohibitions along these lines are gratuitous. Yet they constitute most of the proposed rule and standard. What is at issue, I suggest, is understanding how AI can be used, <i>unwittingly</i>, to produce harmful results. This requires an identification of the <i>underlying risks of AI</i>, and making users aware of those. Those underlying risks are e.g., (i) the alignment problem and AI’s ability to cheat [see n.2], which may be more insidious than the sometimes obvious hallucinations that the Committee identifies; (ii) the bias which may be inherent in the AI’s training including the biases of the humans involved. There are likely others; and the real risks may change over time, all of which can be handled by ensuring alerts and training materials are kept updated. These problems, however, are not addressed by the Committee’s proposal.	The task force will continue to consider how to address the risks posed by use of generative AI in court-related work. In addition to the proposed rule and standard, the task force has developed FAQs and is considering whether additional guidance documents are needed. The task force will work with the Center for Judicial Education & Resources (CJER) to ensure that judicial officers and court staff receive education and training regarding generative AI, including on emerging uses and risks.
11.	LexisNexis by Aron Holewinski Field Client Manager	A	LexisNexis appreciates the opportunity to comment on SP25-01. As a trusted partner to the legal community and provider of secure generative AI tools through Lexis+ AI with Protégé, we support the Judicial Council’s framework that encourages responsible adoption while protecting confidentiality, transparency, and public trust. We offer the following comments:	Please see responses to LexisNexis’s specific suggestions below.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			1. Disclosure Standard: Rule 10.430(d)(5) requires disclosure when AI outputs comprise a “substantial portion” of public-facing content. We recommend providing clarification or examples—particularly where courts use secure, court-approved systems that support administrative or non-adjudicative drafting.	In light of all the comments received on this issue, the task force has recommended a revised version of rule 10.430(d)(5).
			2. Use of Private AI Systems: We commend the proposal’s clear prohibition on entering confidential information into public AI systems. Tools like Lexis+ AI with Protégé are designed with privacy-by-design principles, use AES-256 encryption, and do not share user input with any third party. These systems help courts responsibly integrate AI without risk of data exposure.	No response required.
			3. Secure Personalization: Protégé allows opt-in personalization (e.g., role, practice area, jurisdiction) while preserving user control and deletion rights. We suggest highlighting such personalization models as best practice for improving productivity without compromising neutrality or ethics.	In light of all the comments received on this issue, the task force is recommending a revised definition of “public generative AI system” in both the rule and standard. The revised definition outlines specific data privacy issues to be considered when using generative AI for court-related work.
			4. Benchmarking Against Peers: We support the Council’s rulemaking approach and encourage drawing on peer frameworks such as New Jersey’s Judiciary AI Principles, which emphasize independence, integrity, fairness, and service.	The task force has been monitoring policy and other developments in jurisdictions outside California and will continue to do so.
			5. Implementation Tools: We support the planned release of FAQs and sample use cases. We recommend including examples that distinguish between high-risk and low-risk uses of AI,	The task force will consider these suggested revisions to the FAQs as time and resources permit.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			outline safe disclosure practices, and clarify options for internal vs. public deployment.	
			LexisNexis remains committed to supporting California courts with secure, ethical AI solutions tailored to the unique needs of public institutions.	No response required.
12.	Orange County Bar Association by Mei Tsang, President	AM	I. INTRODUCTION Thank you for the opportunity to review and comment on the Artificial Intelligence Task Force’s proposal (SP25-01) [See www.courts.ca.gov/policyadmin-invitationstocomment.htm]. The emergence of generative AI raises complex issues concerning confidentiality, bias, privacy, security, and, especially, due process in the courts. We appreciate the task force’s proactive efforts to provide guidance and safeguards for court-related use of this evolving technology. There exists an urgent need to address inconsistencies and omissions in the current proposal. Generative AI is too often treated as though it were a simple, standardized tool like a calculator. In reality, these models may incorporate extrarecord or third-party training data, inadvertently expose confidential information, or introduce bias. “Public AI systems” might silently train on user-provided data, thereby creating serious confidentiality and ex parte concerns. Furthermore, any reliance on AI for actual judicial decision-making could imperil a litigant’s due process and constitutional rights.	Please see responses to Orange County Bar Association’s specific suggestions below.
			II. OVERALL IMPRESSIONS OF THE PROPOSAL A. Positive Aspects The efforts of the AI Task Force to draft policy that regulates and standardizes the use of Generative AI in the judiciary is paramount to the ever increasing backlog of cases, which will undoubtedly be further exacerbated by the use of Generative AI by legal practitioners and pro se litigants to facilitate litigation and other adversarial or legal proceedings. The draft policy’s proposed requirements, such as prohibitions	Please see responses to Orange County Bar Association’s specific suggestions below.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			on inputting confidential information and mandates to review AI-generated content for errors, demonstrate the need to address algorithmic bias and misinformation. We also support the separation of guidance for court staff (rule 10.430) and judicial officers in their adjudicative role (standard 10.80), which acknowledges the distinct responsibilities and ethical considerations each group faces. B. Key Concerns It is respectfully observed that adopting one rule and one standard to govern all uses of Generative AI in the judicial branch is ill-advised. Generative AI systems, the courts implementing them, and the roles of those who utilize them are too varied and complex to be effectively regulated under a single, uniform framework. Although uniformity can foster consistency, it risks oversimplifying the critical distinctions between relatively low-stakes administrative tasks and high-risk adjudicative responsibilities. This one-size-fits-all approach may inadvertently undermine fairness, transparency, and public trust by failing to address key nuances—ranging from confidentiality to due process—that arise in different contexts. The proposal identifies certain risks but does not fully explain how reliance on Generative AI for adjudicative tasks might impact or even compromise core judicial responsibilities. Reliance on AI-generated analyses or recommendations may, for example, compromise the transparency of a judge’s legal reasoning or introduce extrarecord data. The reference to a “substantial portion of the content” as a threshold for disclosure is also ambiguous; even minor AI-assisted additions may significantly affect outcomes and should be disclosed for clarity and public confidence. Although courts can adopt stricter rules, the baseline standard might encourage minimal disclosures, potentially failing to ensure due process. Additionally, the proposal relies on permissive phrases such as “should” or “should consider”—particularly in the context of adjudicative roles—which may afford judicial officers broad latitude to rely on Generative AI with minimal oversight or	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			transparency. When fundamental rights are at stake, due process demands more definitive language—such as “must” or “shall”—to underscore the necessity of compliance with strict rules concerning the use of Generative AI, understanding the inherent limits of technology in resolving substantive legal questions, and rigorous judicial oversight. III. DOES THE PROPOSAL ADDRESS ITS STATED PURPOSE? The proposal expressly aims to address the confidentiality, privacy, bias, safety, and security concerns posed by Generative AI in court-related work, while promoting responsible innovation in court operations and preserving public trust. A close reading of the proposal indicates that its ultimate goal is to establish uniform guidelines under which courts, judicial officers, and court staff may use emerging AI tools without compromising the integrity of judicial proceedings. The proposal attempts to meet these objectives by requiring disclosure, mandating oversight in certain scenarios, and prohibiting or limiting various AI-driven practices that could harm litigants or undermine fairness. The question remains, however, whether the proposed language and requirements meaningfully achieve these ends. The following sections examine specific areas to evaluate how effectively the proposal meets its stated purpose.	
			A. Confidentiality, Privacy, and Security 1. Confidentiality and Privacy The proposal takes an important step by prohibiting the entry of confidential or nonpublic information and personal identifying information into public AI systems. This restriction is designed to avoid unauthorized disclosure of sensitive data. However, additional measures could strengthen this protective framework:	Please see responses to Orange County Bar Association’s specific suggestions below.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			▪ Oversight and Auditing: The rule might outline how courts should periodically review or audit compliance with these prohibitions, taking into account the rapidly growing variety of generative AI platforms and plugins. ▪ Alignment with Local Ordinances: Courts in urban or highly populated regions may face stricter local mandates around privacy or data protection. Stating clearly that “federal or state law” includes any applicable municipal or county ordinance would help ensure consistent compliance.	
			2. Security Risks The proposal implicitly acknowledges cybersecurity concerns by cautioning against uploading nonpublic data. Yet Generative AI can also present broader vulnerabilities, such as breaches or hacking attempts that target external AI tools and APIs: ▪ Clear Cybersecurity Guidance: An explicit requirement for courts, judicial officers, and court staff to follow established cybersecurity protocols (e.g., secure credential management, penetration testing) would help protect against data exfiltration or model manipulation. ▪ Mandatory, Not Optional: Treating security as a core obligation—rather than an optional add-on—reinforces the high stakes for litigants whose data could be exposed.	Please see responses to Orange County Bar Association’s specific suggestions below.
			B. Bias, Safety, and the Broader Legal Framework 1. Bias and Safety The proposal sensibly requires users to check AI-generated content for biased, offensive, or harmful outputs. Generative AI can inadvertently replicate or	Please see responses to Orange County Bar Association’s specific suggestions below.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			exacerbate societal biases, posing a real threat to equity and fairness. However, the text offers little detail on how or how often these checks should be performed: ▪ Structured Reviews: Courts could employ standardized, periodic assessments for high-risk uses of Generative AI, ensuring that repeated or systemic biases are identified and remedied rather than dismissed as one-off anomalies. ▪ Remedial Measures: Guidance on what to do if bias is discovered—such as immediate removal, correction, or mandatory re-review—would underscore a commitment to preventing discriminatory harm.	
			2. Local Ordinances and Additional Obligations Some jurisdictions may impose stricter or more targeted rules around anti-discrimination, consumer protection, or privacy. The proposal should affirm that compliance with “federal or state law” necessarily includes abiding by any relevant local requirements. This clarification avoids confusion in courts that operate under a patchwork of municipal and county rules.	The task force is not recommending revisions in response to this suggestion because the applicability of municipal, county, and other local ordinances to the courts can be more complicated than the applicability of state and federal law. Courts should advise judicial officers and court staff regarding any applicable local ordinances.
			C. Maintaining Public Trust Generative AI can enhance efficiency, but it also risks creating the perception that judges rely on automated decision-making rather than personal legal analysis. The proposal’s disclosure requirement is one avenue to mitigate this concern: ▪ Strong Disclosure Protocols: A clear, user-friendly notification—whether appended to a public document or included in an official statement—assures litigants and the public that AI was used only for permissible purposes and that the judge or court staff verified its output.	Please see responses to Orange County Bar Association’s specific suggestions below.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			▪ Avoiding “Substantial Portion” Ambiguity: Even seemingly minor AI contributions can influence outcomes. Replacing the “substantial portion” threshold with any “material role” or “material influence” approach would promote transparency across the board.	
			D. Due Process and Judicial Integrity The proposal acknowledges that judges should not depend on AI to carry out adjudicative duties. However, the language around these limits is sometimes permissive, suggesting practices judges “should” or “should consider” rather than strictly prohibit or require: ▪ Prohibiting AI in Substantive Decision-Making: Due process demands that every judge personally weigh the facts and law before rendering a decision. Any meaningful framework must forbid AI-generated reasoning or factual determinations in final orders, rulings, or judgments. ▪ Clear Enforcement: Reinforcing that judicial officers are solely responsible for all substantive legal conclusions—without recourse to “the AI made me do it”—helps safeguard fundamental rights and anchors the public’s faith in an impartial judiciary. Any model rule must specify clear repercussions for any judicial officer that improperly violates due process by relying on AI-generated reasoning or factual determinations in final orders, rulings, or judgments.	Please see responses to Orange County Bar Association’s specific suggestions below.
			E. Extrarecord and Ex Parte Concerns Generative AI often relies on vast, behind-the-scenes datasets that may include information never presented or challenged in court: ▪ Extrarecord Data: A judge who inadvertently pulls in outside facts from an AI model, especially if those facts are inaccurate or incomplete, risks basing a ruling on material the parties had no opportunity to contest.	Please see responses to Orange County Bar Association’s specific suggestions below.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			▪ Manipulation Risks: There is also a possibility that third parties could systematically “train” or influence public AI models to shape outputs in certain ways. This raises the specter of ex parte communications, where a judge essentially receives input from an unseen party. Any rule or standard must therefore alert judicial officers to the hidden path by which extrarecord data can creep into legal decisions, threatening the adversarial process and the reliability of the record. By explicitly addressing confidentiality, security, anti-bias safeguards, and due process protections, the proposal would more closely fulfill its stated purpose of guiding responsible and equitable implementation of Generative AI in the courts. Strengthened disclosure requirements, categorical prohibitions on AI in adjudicative tasks, and clear avenues for oversight all serve to preserve transparency and bolster public trust in judicial outcomes.	
			IV. SUGGESTED REVISIONS A. Prohibition on Substantive Adjudicative Use An express statement should clarify that Generative AI may not draft, decide, or substantively shape judicial rulings or orders. Allowing the use of Generative AI for administrative or preliminary research tasks is plausible, but the content of judicial decisions must remain a product of the judge’s independent analysis. B. New Detail on Administrative vs. Adjudicative Functions Certain tasks—such as case scheduling, purely clerical tasks, or preliminary citation checks—may benefit from the use of AI while posing minimal due process risk. However, any aspect that informs the final merits of a case or affects legal conclusions should be off-limits, absent robust guardrails and mandatory disclosures. We urge the Judicial Council to define “adjudicative tasks” broadly	The task force is not recommending revisions in response to this suggestion. The task force determined that the question of whether and to what extent judicial officers may use generative AI to carry out their adjudicative duties is more appropriately addressed by the California Code of Judicial Ethics and related ethical guidance. The task force concluded that the canons likely prohibit judicial officers from having generative AI write their opinions for them, and that the Supreme Court’s judicial ethics committees are the

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			enough to capture motions, dispositive orders, and factual or legal determinations at any stage of litigation.	appropriate bodies to ask for guidance on this subject. The task force determined that courts and judicial officers are in the best position to identify acceptable uses of generative AI to meet their specific needs. The risks of generative AI depend heavily on the specific tool and how it is being used. It would be extremely difficult for the task force to create a list of acceptable tools and uses, and such a list would likely be both under- and overinclusive because the task force would have to speculate about how specific tools work or how courts might use them. Putting such a list in a rule of court would make it difficult to keep up with technological advancements. For these reasons, the task force recommends that the rule and standard address specific risks of generative AI rather than specific generative AI tools or uses. Additionally, generative AI is increasingly being incorporated into existing software products

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
				and may already be difficult to avoid in some circumstances. The task force is concerned that placing branchwide limitations on specific uses of generative AI will unnecessarily limit innovation and will prevent courts from identifying safe, effective uses of generative AI that do not pose ethical risks.
			C. Mandatory Disclosure Disclosure should be mandatory whenever AI contributes to an official document or statement that is shared with the public. The threshold of “substantial portion” could be replaced by a simple rule requiring disclosure of any material influence on the text or decision.	In light of all the comments received on this issue, the task force has recommended a revised version of rule 10.430(d)(5). However, the task force concluded that mandatory disclosure of use of generative AI in all circumstances could unnecessarily prohibit courts from using generative AI in circumstances where the technology can be used safely and ethically, and that mandatory disclosure would not be an effective way to address concerns about the reliability or trustworthiness of generative AI outputs in many circumstances.
			D. Certification of Compliance	In light of all the comments received on this issue, the task force is recommending revisions

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			In addition to disclosure, the court or judicial officer should certify in a verifiable manner (e.g., a short statement in the published ruling or accompanying document) that they reviewed and verified AI-generated content, took responsibility for its substance, and complied with all applicable rules. Such certifications are critical for ensuring accountability, especially if a litigant later challenges the AI-influenced ruling on appeal.	to rule 10.430(d)(3) and (4) to require users to take reasonable steps to verify, correct, or remove inaccurate or biased material. The task force is recommending similar revisions to standard 10.80(b)(3) and (4). The task force concluded this terminology will make clearer that courts cannot simply identify inaccurate or biased information in generative AI material; they must also verify the accuracy of the material and correct or remove any inaccurate or biased information.
			E. Removing Ambiguity Around “Substantial Portion” If the term “substantial portion” remains, it should be clearly defined. However, a better approach is to disclose all generative AI usage that informs the final document in any way.	Please see the committee’s previous response to the Orange County Bar Association regarding the task force’s recommended revisions to rule 10.430(d)(5).
			F. Emphasis on Due Process A clearer statement that due process concerns prohibit judges from delegating their legal analysis to AI would confirm the judiciary’s commitment to fairness and transparency. This emphasis would assure litigants that final determinations remain firmly under human judicial control.	The task force is not recommending revisions in response to this suggestion. The task force concluded that the Supreme Court’s judicial ethics committees are the appropriate bodies to ask for guidance on this subject.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			G. Potential Appellate Remedies We call out the need for explicit remedies if a judicial officer violates these standards. If reliance on AI introduces overt errors, biased language, or fundamental defects in a ruling, it is unclear how that misconduct would be addressed on appeal. We encourage the Judicial Council to consider clarifying that demonstrable violations of these AI policies may form part of the record for appellate review or judicial disciplinary proceedings.	This suggestion is beyond the scope of the task force’s charge.
			H. Training and Ongoing Education Additional or enhanced training would help judges and staff understand the limitations of Generative AI. Regular updates on best practices for validation, bias detection, and data protection would further mitigate the risks associated with rapidly evolving AI tools.	Revising the rule and standard to implement this suggestion would require further public comment because it is beyond the scope of issues presented in this invitation to comment. The task force may consider it as time and resources permit. Additionally, the task force will consider whether to implement this suggestion, and others that are beyond the scope of this invitation to comment, via other means including the model policy or other guidance documents.
			I. Alignment With Existing California Statutes The task force should align any new rule or standard with existing or pending California legislation defining Artificial Intelligence and Generative AI. Several bills (e.g., AB 2013, AB 2885, and SB 942) contain more precise definitions of “Artificial Intelligence,” “Generative AI,” and “Training Data.” In particular, adopting these legislative definitions would create a consistent legal framework and	In light of all the comments received on this issue, the task force is recommending revised definitions of “generative AI” and “public generative AI system” in the rule and standard. The task force recommends deleting the

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			reduce confusion for practitioners who must reconcile the Rules of Court with other AI-specific statutes.	definition of “artificial intelligence” in the rule and standard because it is not needed due to the proposed revisions to the definition of “generative AI.” However, although the task force agrees that consistency with statutory definitions can be beneficial in some circumstances, the existing statutory definitions, such as those in Civil Code section 3110, are not a good fit for the rule and standard. Those definitions are part of statutory schemes for regulating AI providers and use terminology that laypeople might find confusing, such as the reference in section 3110(a) to “explicit and implicit objectives.”
			SUGGESTED REVISIONS V. RULE 10.430. GENERATIVE ARTIFICIAL INTELLIGENCE USE POLICIES A. Section “(a) Definitions” - Original Text Title 10. Judicial Administration Rules Division 2. Administration of the Judicial Branch Chapter 6. Court Technology, Information, and Automation	No response required. (This portion of the comment appears to copy the proposed rule verbatim.)

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			Rule 10.430. Generative artificial intelligence use policies (a) Definitions As used in this rule, the following definitions apply: 1. “Artificial intelligence” or “AI” means technology that enables computers and machines to reason, learn, and act in a way that would typically require human intelligence. 2. “Court staff” means all employees, contractors, volunteers, and any other persons working for or on behalf of the court. 3. “Generative AI” means artificial intelligence trained on an existing set of data (which can include text, images, audio, or video) with the intent to “generate” new data objects when prompted by a user. Generative AI creates new data objects contextually in response to user prompts based only on the data it has already been trained on. 4. “Judicial officer” means all judges, all justices of the Courts of Appeal and the Supreme Court, all temporary and assigned judges, and all subordinate judicial officers. 5. “Public AI system” means a system that is publicly available or that allows information submitted by users to be accessed by anyone other than judicial officers or court staff, including access for the purpose of training or improving the system.	
			1. Comment re “(a) Definitions” The definitions for “Artificial Intelligence,” “Generative AI,” and “Public AI System” in rule 10.430(a) merit closer examination.	Please see previous responses to the Orange County Bar Association’s comments regarding these definitions.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			a. Use of a Non-Standard Definition for Artificial Intelligence The proposal’s definition—“technology that enables computers and machines to reason, learn, and act in a way that would typically require human intelligence”—is overly broad and risks classifying simpler technologies (e.g., basic automation or statistical software, or even calculators) as AI. Recent California legislation defines “artificial intelligence” more precisely as an engineered or machine-based system that, varying in its level of autonomy, can infer from inputs how to generate outputs that may influence physical or virtual environments. Adopting that statutory language would align court rules with broader state policy and reduce confusion for practitioners. b. Clarifying “Generative AI” The rule currently frames “Generative AI” as technology “trained on an existing set of data” to “generate” new objects in response to user prompts. However, the phrase “based only on the data it has already been trained on” could inadvertently exclude systems that incorporate supplemental, user-provided data during an interactive session or that fine-tune outputs after receiving new inputs. By contrast, California’s statutory definition of “generative artificial intelligence system” or “GenAI system” references the “creation of derived synthetic content” (text, images, video, etc.) that emulates the structure and characteristics of the system’s training data. This language better captures how large language models actually produce new text or media, even if they integrate user-provided material at runtime. c. Scope and Impact of “Public AI System”	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			The proposal’s definition, which hinges on a system being “publicly available” or allowing user submissions to be “accessed by anyone other than judicial officers or court staff,” does not fully account for the reality that many AI tools—free or paid—may partially use user data to refine or retrain models. A more robust definition would explicitly address whether user-input data is retained, used, sold, or shared beyond the immediate generation of outputs requested by the user. This approach captures both openly accessible systems (e.g., consumer-facing chatbots) and those that, though credentialed or licensed, still aggregate user data for optimization, thereby posing similar confidentiality and privacy risks. d. Need for Alignment with Existing Statutory and Local Authority Courts operate within a complex legislative environment that includes federal, state, and local privacy or anti-discrimination mandates. To foster consistency, the definitions of AI-related terms in rule 10.430 should align with relevant California statutes where possible. Explicitly referencing these legislative definitions (e.g., Civil Code sections on AI training data transparency) would clarify that courts must uphold evolving legal standards and further ensure that no narrower or conflicting definitions undercut the rule’s stated goals of confidentiality, safety, and due process. e. Practical Implications Overly broad or vague definitions can inadvertently chill beneficial innovation or open the door to unregulated AI deployments that compromise court users’ data. Narrow, precise definitions, consistent with statutory language, help courts differentiate high-risk “Generative AI” from routine, rules-based automation tools,	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			thus allowing more targeted safeguards and clearer directives for staff and judicial officers. In sum, the definitions in the current proposal would benefit from (1) adopting or closely mirroring established legislative language, (2) explicitly addressing any supplemental or user-provided training data, and (3) expanding “public AI system” to encompass all systems that retain or exploit user data for retraining or sharing with third parties. Such refinements provide clearer guardrails against confidential-information leaks, reduce ambiguity about which tools are regulated, and better align with the broader California legal framework.	
			2. Suggested Revised Language for “(a) Definitions” (a) Definitions As used in this rule, the following definitions apply: 1. “Artificial intelligence” or “AI” means an engineered or machine-based system that varies in its level of autonomy and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments. 2. “Court staff” means all employees, contractors, volunteers, and any other persons working for or on behalf of the court. 3. “Generative artificial intelligence” or “Generative AI” means an artificial intelligence that can generate derived synthetic content—including text, images, video, or audio—that emulates the structure and characteristics of the system’s training data.	Please see the previous responses to the Orange County Bar Association’s comments regarding these definitions.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			4. “Judicial officer” means all judges, justices of the Courts of Appeal and the Supreme Court, all temporary and assigned judges, and all subordinate judicial officers. 5. “Public AI system” means any artificial intelligence platform, model, or service that: (A) Is accessible to the general public—with or without cost—and does not require specialized credentials or licenses beyond ordinary consumer terms of service; or (B) Retains, uses, sells, or shares user-input data for additional training, optimization, or any other purpose beyond the immediate creation of outputs requested by the user.	
			B. Section “(b) Generative AI use policies” - Original Text (b) Generative AI use policies If a superior court, Court of Appeal, or the Supreme Court permits the use of generative AI by court staff or judicial officers, that court must adopt a generative AI use policy. 1. Comment re “(b) Generative AI use policies” Even for tasks that are ostensibly administrative or non-adjudicative, serious constitutional due process concerns can arise if AI-generated outputs inadvertently shape, inform, or otherwise influence judicial decision-making. Court staff who rely on Generative AI for research, drafting, or data processing may unknowingly introduce biased or extrarecord material into the workflow, potentially undermining the impartiality required by both state and federal constitutions. As such, any policy permitting staff to use AI must include strict controls and oversight to prevent undue influence on a litigant’s right to a fair hearing.	The task force agrees that rule 10.430, as originally proposed, inadvertently excluded courts that do not take a position on use of generative AI. The task force recommends that rule 10.430(b) read as follows: “Any court that does not prohibit the use of generative AI by court staff or judicial officers must adopt a generative AI use policy. This rule applies to the superior courts, the Courts of Appeal, and the Supreme Court.”

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			2. Suggested Revised Language for “(b) Generative AI use policies” (b) Generative AI Use Policies (1) A superior court, Court of Appeal, or the Supreme Court shall not permit the use of generative AI by court staff or judicial officers unless and until that court adopts a written Generative AI Use Policy. (2) All uses of generative AI are prohibited except as expressly authorized by the Generative AI Use Policy adopted under subdivision (b)(1).	
			C. Section “(d) Policy requirements” - Original Text (d) Policy requirements Each court’s generative AI use policy must: 1. Prohibit the entry of confidential, personal identifying, or other nonpublic information into a public generative AI system. Personal identifying information includes driver license numbers; dates of birth; social security numbers; Criminal Identification and Information, and National Crime Information numbers; addresses and phone numbers of parties, victims, witnesses, and court personnel; medical or psychiatric information; financial information; account numbers; and any other content sealed by court order or deemed confidential by court rule or statute. 2. Prohibit the use of generative AI to unlawfully discriminate against or disparately impact individuals or communities based on age, ancestry, color, ethnicity, gender, gender expression, gender identity, genetic information, marital status, medical condition, military or veteran status, national origin, physical or mental disability,	No response required. (This portion of the comment appears to copy the proposed rule verbatim.)

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			political affiliation, race, religion, sex, sexual orientation, socioeconomic status, and any other classification protected by federal or state law. 3. Require court staff and judicial officers who generate or use generative AI material to review the material for accuracy and completeness, and for potentially erroneous, incomplete, or hallucinated output. 4. Require court staff and judicial officers who generate or use generative AI material to review the material for biased, offensive, or harmful output. 5. Require disclosure of the use or reliance on generative AI if generative AI outputs constitute a substantial portion of the content used in the final version of a written or visual work provided to the public. 6. Require compliance with all applicable laws, court policies, and ethical and professional conduct rules, codes, and policies when using generative AI.	
			1. Comment re “(d) Policy requirements” The listed requirements represent an important foundation for safe, non-discriminatory AI usage, but each provision could be bolstered to reflect the higher stakes of introducing generative AI in court operations.	No response required.
			First, limiting disclosure to situations where “a substantial portion” of the final content originates from AI may be too narrow, as even minimal AI-generated or AI-informed material can significantly influence court documents or communications; a more effective standard would require disclosure of any material reliance on AI.	Please see previous response to the Orange County Bar Association’s comments regarding the disclosure requirement.
			Second, provisions requiring review of “erroneous,” “incomplete,” or “biased” outputs should make explicit that court staff and judicial officers are personally responsible for verifying and correcting all AI-generated content, rather than treating AI as an authoritative source.	Please see previous response to the Orange County Bar Association’s comments regarding the task force’s recommended revisions to rule

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
				10.430(d)(3) and (4) and standard 10.80(b)(3) and (4).
			Lastly, restricting the entry of confidential data and prohibiting disparate impacts are necessary safeguards, but the policy should clarify strict compliance with local as well as state and federal rules, ensuring that all relevant privacy, anti-discrimination, and ethical standards are upheld.	The task force is not recommending revisions in response to this suggestion. The rule requires compliance with “all applicable laws,” which would include local laws.
			2. Suggested Revised Language for “(d) Policy requirements” (d) Policy requirements Each court’s generative AI use policy must: 1. Prohibit the entry of confidential, personal identifying, or other nonpublic information into a public generative AI system. Personal identifying information includes driver license numbers; dates of birth; social security numbers; Criminal Identification and Information, and National Crime Information numbers; addresses and phone numbers of parties, victims, witnesses, and court personnel; medical or psychiatric information; financial information; account numbers; and any other content sealed by court order or deemed confidential or personally identifiable information by court rule, statute, county or municipal ordinance, or other applicable law. 2. Prohibit the use of generative AI to unlawfully discriminate against or disparately impact individuals or communities based on age, ancestry, color, ethnicity, gender, gender expression, gender identity, genetic information, marital status, medical condition, military or veteran status, national origin, physical or mental disability, political affiliation, race, religion, sex, sexual orientation, socioeconomic status, and any other classification protected by federal, state, or local or other applicable law.	Please see previous responses to the Orange County Bar Association’s suggested revisions to rule 10.430(d).

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			3. Require court staff and judicial officers who generate or use generative AI material to review the material for accuracy and completeness, and for potentially erroneous, incomplete, or hallucinated output. 4. Require court staff and judicial officers who generate or use generative AI material to review the material for biased, offensive, or harmful output. 5. Require compliance with all applicable laws, court policies, and ethical and professional conduct rules, codes, and policies when using generative AI. 6. If any part of a written or visual work is derived from, informed by, or relies upon generative AI, or if generative AI was used to support research, logic, or reasoning incorporated into the final version of that work for public or official court use, the policy must require a clear, prominent statement (in the document itself or by separate accompanying note) that generative AI was utilized, including a brief description of how it was used; and Court staff or judicial officers who used generative AI must certify, in a verifiable manner, that they have reviewed, verified, and validated all AI-generated content or logic.	
			VI. RULE 10.80. USE OF GENERATIVE ARTIFICIAL INTELLIGENCE BY JUDICIAL OFFICERS A. Section “(a) Definitions” - Original Text Title 10. Standards for Judicial Administration Standard 10.80. Use of generative artificial intelligence by judicial officers (a) Definitions As used in this standard, the following definitions apply:	No response required. (This portion of the comment appears to copy the proposed standard verbatim.)

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			1. “Artificial intelligence” or “AI” means technology that enables computers and machines to reason, learn, and act in a way that would typically require human intelligence. 2. “Court staff” means all employees, contractors, volunteers, and any other persons working for or on behalf of the court. 3. “Generative AI” means artificial intelligence trained on an existing set of data (which can include text, images, audio, or video) with the intent to “generate” new data objects when prompted by a user. Generative AI creates new data objects contextually in response to user prompts based only on the data it has already been trained on. 4. “Judicial officer” means all judges, all justices of the Courts of Appeal and the Supreme Court, all temporary and assigned judges, and all subordinate judicial officers. 5. “Public AI system” means a system that is publicly available or that allows information submitted by users to be accessed by anyone other than judicial officers or court staff, including access for the purpose of training or improving the system.	
			1. Comment re “(a) Definitions” The definitions set forth in Rule 10.430 regarding “Artificial Intelligence,” “Generative AI,” and “Public AI System” are incorporated here for consistency and clarity, avoiding discrepancies between the two provisions. These definitions carry heightened importance in Standard 10.80 because judicial officers’ adjudicative responsibilities implicate core due process concerns. Precisely defining key terms ensures that any technology deemed “AI” or “Generative AI” does not inadvertently erode the impartiality, accuracy, and transparency required of judges, particularly when public AI systems might embed	Please see previous responses to the Orange County Bar Association’s comments regarding these definitions.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			biases or introduce extrarecord data. By aligning these definitions, courts reinforce a unified approach to regulating AI within both administrative and adjudicative contexts, laying the groundwork for the stricter scrutiny that follows in the subsequent sections on judicial officers’ use of AI. 2. Suggested Revised Language for “(a) Definitions” (a) Definitions As used in this rule, the following definitions apply: 1. “Artificial intelligence” or “AI” means an engineered or machine-based system that varies in its level of autonomy and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments. 2. “Court staff” means all employees, contractors, volunteers, and any other persons working for or on behalf of the court. 3. “Generative artificial intelligence” or “Generative AI” means an artificial intelligence that can generate derived synthetic content—including text, images, video, or audio—that emulates the structure and characteristics of the system’s training data. 4. “Judicial officer” means all judges, justices of the Courts of Appeal and the Supreme Court, all temporary and assigned judges, and all subordinate judicial officers. 5. “Public AI system” means any artificial intelligence platform, model, or service that:	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			(A) Is accessible to the general public—with or without cost—and does not require specialized credentials or licenses beyond ordinary consumer terms of service; or (B) Retains, uses, sells, or shares user-input data for additional training, optimization, or any other purpose beyond the immediate creation of outputs requested by the user.	
			B. Section “(b) Use of generative artificial intelligence” - Original Text (b) Use of generative artificial intelligence A judicial officer using generative AI for any task within their adjudicative role: 1. Should not enter confidential, personal identifying, or other nonpublic information into a public generative AI system. Personal identifying information includes driver license numbers; dates of birth; social security numbers; Criminal Identification and Information, and National Crime Information numbers; addresses and phone numbers of parties, victims, witnesses, and court personnel; medical or psychiatric information; financial information; account numbers; and any other content sealed by court order or deemed confidential by court rule or statute. (2) Should not use generative AI to unlawfully discriminate against or disparately impact individuals or communities based on age, ancestry, color, ethnicity, gender, gender expression, gender identity, genetic information, marital status, medical condition, military or veteran status, national origin, physical or mental disability, political affiliation, race, religion, sex, sexual orientation, socioeconomic status, and any other classification protected by federal or state law.	No response required.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			(3) Should review generative AI material, including any materials prepared on their behalf by others, for accuracy and completeness, and for potentially erroneous, incomplete, or hallucinated output. (4) Should review generative AI material, including any materials prepared on their behalf by others, for biased, offensive, or harmful output. (5) Should consider whether to disclose the use of generative AI if it is used to create content provided to the public.	
			1. Comment re “(b) Use of generative artificial intelligence” The revised text in subdivisions (b)(1) through (b)(6) incorporates essential safeguards that address fundamental due process concerns and enhance transparency in adjudicative procedures. Below is a section-by-section commentary: a. Balancing Innovation and Safeguards Although emergent AI platforms—like ChatGPT, Claude, or Google Gemini—offer efficiencies in research and drafting, these open, continuously trained systems can be highly susceptible to external manipulation. Introducing them into case determinations without robust oversight could, over time, distort legal interpretations and place judicial officers at risk of unknowingly relying on skewed or extrarecord content. Initially restricting AI usage to “purely administrative or ministerial” tasks, or for preliminary legal research with mandatory verification, prevents generative AI from displacing the judge’s personal assessment of facts or law. This two-pronged limitation acknowledges that certain low-risk functions—like routine scheduling or preliminary research—may benefit from AI’s efficiency, provided that the technology does not encroach on the judge’s responsibility to fully and independently analyze the record.	Please see previous response regarding the scope of the rule and standard and the suggestion that the task force prohibit specific uses of generative AI.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			b. Why Strong Oversight Is Needed in the Context of Judicial Use Unlike traditional legal tools that remain static, Generative AI models learn from user inputs and evolving data sets. Malicious actors could exploit this by “training” AI engines to favor certain legal outcomes or embed subtle biases into the model’s outputs. If a judicial officer consults such a model for case analysis or precedent, the tool may—intentionally or not—push misleading or incomplete interpretations. This raises significant due process concerns, as parties have a constitutional right to a judge’s independent, impartial assessment based solely on the record and applicable law. By limiting judicial officers from using or relying on Generative AI for any task that may affect the merits of a case, this provision ensures that core judicial decision-making remains the product of independent, human deliberation, free from opaque or extrarecord influences. This strict prohibition reflects the constitutional imperative that litigants have their cases decided by a judge rather than delegated to an AI system, especially given the high risk of biased or inaccurate outputs. c. Current Use: Limited and Cautious For now, limiting AI to purely administrative or preliminary research tasks—while prohibiting its direct influence on adjudicative rulings—helps preserve the integrity of the judicial process. If and when courts develop specialized, secured AI systems or more comprehensive training protocols, the door remains open for controlled expansion of AI usage. By adopting a “go slow” philosophy, courts minimize risks of systemic distortion while still exploring the potential benefits of carefully managed AI tools. The proposed limitations address concerns that third-party AI platforms could retain, train on, or inadvertently expose sensitive data—potentially creating ex parte channels of information or violating court orders. By absolutely banning such	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			disclosures, this provision reinforces the judiciary’s duty to uphold confidentiality and preserve the integrity of court records. d. Alternative Venues for AI-Adjudication Parties who wish to incorporate AI more extensively into dispute resolution are free to do so in private forums such as mediation or arbitration, where participants can consensually structure how and when AI is used. This allows experimentation and cost-efficiency within a self-governed framework, without compromising the stricter due process protections that public courts must uphold. e. Due Process as the Paramount Concern Above all, due process demands that judges derive outcomes from evidence properly admitted and law thoroughly vetted—requirements that stand at the core of public trust in the judiciary. Any reliance on AI in adjudicative roles must respect these constitutional and statutory principles, ensuring no hidden biases or manipulated outputs can sidestep adversarial testing or informed judicial scrutiny. A duty to disclose whenever AI contributes to a publicly provided document strengthens transparency and mitigates concerns that litigants or the public might be misled into believing the text is wholly the result of judicial authorship. Requiring a brief explanation of how AI was used empowers parties to challenge or probe whether reliance on AI has potentially introduced extraneous or biased materials, thus upholding the adversarial process and promoting informed review of judicial actions.	
			2. Suggested Revised Language for “(b) Use of generative artificial intelligence” (b) Use of generative artificial intelligence	Please see previous responses to the Orange County Bar Association’s comments.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			1. Except as expressly allowed below, a judicial officer must not use or rely on generative AI for any task that may affect the substance of an adjudicative decision, including but not limited to drafting orders, rulings, or opinions; analyzing evidence; or making factual or legal determinations. 2. A judicial officer may use generative AI only if: (A) The task is purely administrative or ministerial and does not implicate the merits of a case, the due process rights of any party, or the judge’s independent decision-making; or (B) It is for preliminary legal research and the judicial officer independently verifies any output (including citations, quotations, or summaries) before relying on it in an official decision or document. 3. When using any form of generative AI, a judicial officer must not provide or upload: (A) Any confidential, personal identifying, or sealed information; (B) Any nonpublic details about a case or proceeding; or (C) Any other information whose disclosure is restricted by statute, court rule, or court order. 4. A judicial officer must not use, rely on, or distribute generative AI outputs that unlawfully discriminate against or disparately impact individuals or groups on the basis of any protected classification under federal or state law. 5. Before relying on any material produced by generative AI—even for administrative or research purposes—a judicial officer must:	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			(A) Review the output for accuracy, completeness, and potential hallucinations or omissions; and (B) Assess the output for possible bias, offensive content, or other harmful language. 6. If generative AI is used for any part of a written or visual work provided to the public (including memoranda, reports, official notices, or other court-related documents), the judicial officer must disclose that generative AI was utilized and briefly explain how it was used.	
			CONCLUSION We appreciate the Artificial Intelligence Task Force’s initiative and thoughtfulness in regulating the use of Generative AI. The proposal addresses many of the confidentiality, security, and bias concerns associated with this technology. We believe, however, that the language could be strengthened to more effectively protect due process, ensure consistent disclosures, and clarify strict limitations on adjudicative tasks. With these refinements, the proposal would better maintain the public’s trust in the judicial process and uphold the courts’ integrity in an era of expanding AI capabilities. At this time, the nascency of Generative AI and its application in legal practice underscored that no rule of court should currently grant broad discretion to judicial officers or staff to adopt Generative AI unless a robust policy is in place. Such a policy should initially prohibit all uses except those narrowly authorized—particularly for judges, who must always remain the ultimate arbiters of legal and factual questions. Additionally, “public AI systems” need a clearer definition emphasizing whether data inputs may be used for further training or be exposed to third parties. By	Please see previous responses to the Orange County Bar Association’s comments.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			refining these key areas, the Judicial Council will better ensure the proposed rule and standard fulfill their purpose without inadvertently undermining core legal protections. Thank you for considering these comments. We welcome further discussion and stand ready to support the Judicial Council’s continued efforts on this critical matter.	
			Model Policy for Use of Generative Artificial Intelligence I. PURPOSE AND SCOPE [REVISED] The emergence of generative artificial intelligence (generative AI) technologies has prompted the court to develop this policy governing the responsible use of AI for court-related work. While generative AI can enhance efficiency, its open-ended and continuously trained nature raises significant concerns about confidentiality, data security, bias, and due process. a. These requirements apply to any use of generative AI systems by court staff for any purpose and by judicial officers for any task outside their adjudicative role. Generative AI systems referenced herein include well-known applications (e.g., ChatGPT, Claude, Dall-E2, Microsoft’s Copilot, Google’s Gemini, Westlaw Precision, Lexis+ AI, Grammarly) and any features in non-AI applications (e.g., Adobe Acrobat, Google search) that generate new content from user prompts. Where reference is made to “public” generative AI systems, it includes any system that is freely or widely accessible without specialized credentials, or that retains or uses user data for further training or dissemination. b. Under no circumstances does this policy authorize judicial officers to rely on generative AI in substantive adjudicative matters unless expressly permitted by separate rule or standard.	The task force will be updating the model policy to conform with changes made to the rule and standard in response to the public comments (such as changes to defined terms). The task force will consider other suggested revisions to the model policy as time and resources permit.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			II. DEFINITIONS [REVISED] For purposes of this policy only, the following definitions apply and should be interpreted consistently with other applicable court rules or standards: a. “Artificial intelligence” or “AI” means an engineered or machine-based system, varying in its level of autonomy, that can infer from inputs how to generate outputs capable of influencing physical or virtual environments. b. “Court staff” means all employees, contractors, volunteers, and any other persons working for or on behalf of the court. c. “Generative AI” or “Generative artificial intelligence” means an AI system capable of creating derived synthetic content (including text, images, audio, or video) based on the structure and characteristics of its training data, in response to user prompts. d. “Judicial officer” means all judges, justices of the Courts of Appeal and the Supreme Court, all temporary and assigned judges, and all subordinate judicial officers. e. “Public generative AI system” means any AI platform, model, or service that is accessible to the general public without specialized credentials or licenses, or that retains, sells, or shares user-input data for any purpose beyond the immediate generation of requested outputs (including further training or optimization). f. “User” means any person to whom this policy applies, including both court staff and judicial officers acting in a non-adjudicative capacity. III. CONFIDENTIALITY AND PRIVACY	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			a. Users must not submit confidential, personal identifying, or other nonpublic information to a public generative AI system. Personal identifying information includes driver license numbers; dates of birth; social security numbers; Criminal Identification and Information, and National Crime Information numbers; addresses and phone numbers of parties, victims, witnesses, and court personnel; medical or psychiatric information; financial information; account numbers; and any other content sealed by court order or deemed confidential by court rule or statute. [REVISED TO MAKE THESE PROVISIONS NON-OPTIONAL] b. If a document has been filed or submitted for filing in a case before the court, users must not submit it to a public generative AI system, even if the document is publicly available. c. Before submitting any information to a public generative AI system, the user must determine whether the submission is permissible under this policy. If it is unclear whether the submission is permissible, the user must obtain approval from [court leadership/their supervisor] before submitting the information to the system. [Courts adopting this provision should consider how to define “court leadership” if approval is to be given by the presiding judge, clerk/executive officer, court executive officer, or chief information officer, or other member of court leadership. Courts requiring approval by court leadership should also consider whether to include a provision allowing leadership to delegate approval authority to others.] d. [REVISED] When using a public generative AI system, users must disable or opt out of any data collection by the system. If it is not feasible to do so, or if the platform does not offer a version that forgoes collecting or training on user-submitted data, that system must not be used. Where the platform provides a licensed or paid version that refrains from retaining or using user-input data for training, users must use that version in lieu of a free, data-collecting service.	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			IV. SUPERVISION AND ACCOUNTABILITY a. Generative AI systems sometimes “hallucinate,” meaning they provide false or misleading information presented as fact. Generative AI outputs can also be faulty in other ways, such as outputs that are inaccurate, incomplete, or uncited. Users must review their generative AI material for accuracy and completeness, and for potentially erroneous, incomplete, or hallucinated output. Any use of generative AI outputs is ultimately the responsibility of the person who authorizes or uses it. [REVISED TO MAKE THESE PROVISIONS NON-OPTIONAL] b. Users must obtain approval from [specify which office, department, division, or individual will be responsible for approval] before using a public generative AI system. c. Public generative AI systems may be used only if they have been approved by the court [specify which office, department, division, or individual will be responsible for approval]. V. AVOIDANCE OF BIAS AND DISCRIMINATION a. Generative AI must not be used to unlawfully discriminate against or disparately impact individuals or communities based on age, ancestry, color, ethnicity, gender, gender expression, gender identity, genetic information, marital status, medical condition, military or veteran status, national origin, physical or mental disability, political affiliation, race, religion, sex, sexual orientation, socioeconomic status, and any other classification protected by federal or state law. b. Generative AI systems may be trained on material that reflects cultural, economic, racial, gender, and social biases, and content generated by these systems may contain biased or otherwise offensive or harmful material. Users must review their generative AI material for biased, offensive, or harmful output.	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			VI. TRANSPARENCY a. If generative AI outputs constitute a substantial portion of the content used in the final version of a written work or visual work that is provided to the public, the work must contain a disclaimer or watermark. b. Labels or watermarks used to disclose the use of generative AI should be easily visible and understandable, accurately informing the audience that generative AI has been used in the creation of the content and identifying the system used to generate it. VII. COMPLIANCE WITH APPLICABLE LAWS AND POLICIES a. When using generative AI, users must comply with all applicable laws, court policies, and ethical and professional conduct rules, codes, and policies. [Optional paragraph: b. Users should be aware that content produced by generative AI systems might include copyrighted material. If it is unclear whether the content produced includes copyrighted material, the user must consult [specify which office, department, division, or individual will be responsible for advice].] VIII. SAFETY AND SECURITY [REVISED TO BE MANDATORY] a. Users must use strong passwords when using AI platforms. Users must comply with the court's password requirements when creating passwords for generative AI platforms. b. When using generative AI systems to perform court-related work, users must use their court email address if the system requires users to provide an email address or create an account. Accounts created using a court email address must not be used	

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			for personal matters. [Courts should also consider whether to require staff to provide their supervisor or IT department with the username and password of any generative AI account created to do court-related work.] c. [REVISED] Public generative AI systems may only be used after the court's [specify which office, department, division, or individual is responsible for approval] has thoroughly reviewed and explicitly approved the system, confirming that it meets the confidentiality, security, data-collection opt-out, and other requirements outlined in this policy.	
13.	Public Counsel by Karla Chalif VP, COO, General Counsel Los Angeles	A	We support widespread adoption in accordance with the rules.	The task force appreciates the response.
14.	Peter Rundle Attorney - Arbitrator - Mediator Rundle Law Corporation Dana Point	N	Do not ever use generative AI to create juridical orders, judgments, opinions, etc. Judges and justices should take pride in the authorship of their writings -- start to finish. Happy to discuss all the reasons for this.	The task force appreciates the commenter's concerns but concluded that the question of whether and to what extent judicial officers may use generative AI to carry out their adjudicative duties is more appropriately addressed by the California Code of Judicial Ethics and related ethical guidance. The task force concluded that the canons likely prohibit judicial officers from having generative AI write their opinions for them, and that the Supreme Court's judicial ethics committees are the

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
				appropriate bodies to ask for guidance on this subject.
15.	SEIU California and TechEquity by Sandra Barreiro, Governmental Relations, Advocate SEIU California and Vinhcent Le, Vice President of AI Policy, Tech Equity	NI	On behalf of SEIU California and Tech Equity, we appreciate the opportunity to provide comments on the Artificial Intelligence Task Force’s model policy concerning the use of generative artificial intelligence (GenAI) in court-related work. Our organizations are deeply committed to ensuring the responsible development, deployment, and governance of AI systems, especially within the judicial system where the stakes for individuals’ lives and well-being are exceptionally high. To that end, we find the proposed model policy to be alarmingly inadequate and potentially dangerous if adopted in its current form. The policy’s superficial treatment of critical issues creates a false sense of security, suggesting that mere adoption of the model policy is enough to mitigate the significant risks of GenAI associated with specific uses of this tool in the courts. This approach could have severe consequences for the integrity of our justice system and the public it serves. Therefore, we submit the following comments, urging the task force to undertake a fundamental revision of the model policy to prevent the serious harms that could arise from the unaccountable and ill-considered use of GenAI in California courts.	Please see responses to SEIU California and Tech Equity’s specific suggestions below.
			II. The Model Policy Is Insufficiently Comprehensive The model policy is not comprehensive enough to promote responsible innovation and public trust. As-is, the model policy could greenlight the irresponsible use of GenAI by trial courts due to the lack of specificity and guidance on how to effectively implement the propose policy. To ensure that GenAI serves to advance, rather than undermine, trust within the court and our justice system, the following elements should be addressed:	No response required.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			A. Require Pre-Deployment Analysis In addition to the requirements set forth in Rule 10.430, each court’s GenAI use policy should be informed by a pre-deployment risk/benefit analysis tailored to particular use cases. A pre-deployment analysis is necessary to meaningfully inform each court’s individual policy on GenAI and should be mandatory before making the decision to deploy a GenAI system for an identified use-case. As the NIST AI Risk Management Framework (RMF) makes clear, effectively governing the use of AI requires courts to map the purpose, risks, and beneficial uses of GenAI and understand the potential impacts, benchmarks, and capabilities for each identified use.[1] Other resources, such as the State of California’s GenAI guidelines and the Center for Democracy & Technology’s (CDT) guide for public sector use of AI, echo the need for pre-deployment assessments and provide valuable frameworks and questions that courts should address before adopting GenAI.[2] [1] NIST’s AI RMF provides specific guidance for organizations to govern the use of AI by mapping, measuring and managing the risks of AI. For example, Map 1.1 requires that the “intended purpose, potentially beneficial uses, context-specific laws, norms and expectations, and prospective settings in which the AI system will be deployed are understood and documented. The framework notes that “the information gathered while carrying out the MAP function enables negative risk prevention and informs decisions for processes such as model management, as well as an initial decision about appropriateness or the need for an AI solution.” See National Institute of Standards and Technology. (January 2023). <i>NIST Artificial Risk Management Framework 1.0</i>, at pp. 24-26. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf [2] California Department of Technology, Department of General Services, Office of Data and Innovation, & Department of Human Resources. (2024, March). <i>State of California GenAI Guidelines for Public Sector Procurement, Uses and Training</i>. https://www.govops.ca.gov/wp-content/uploads/sites/11/2024/03/3.a-GenAI-	Revising the rule and standard to implement this suggestion would require further public comment because it is beyond the scope of issues presented in this invitation to comment. The task force may consider it as time and resources permit. Additionally, the task force will consider whether to implement this suggestion, and others that are beyond the scope of this invitation to comment, via other means including the model policy or other guidance documents.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			Guidelines.pdf; Srinivasan, S., & Laird, E. (2025, March). <i>To AI or Not to AI: A Practice Guide for Public Agencies to Decide Whether to Proceed with Artificial Intelligence</i>. Center for Democracy & Technology. https://cdt.org/insights/to-ai-or-not-to-ai-a-practice-guide-for-public-agencies-to-decide-whether-to-proceed-with-artificial-intelligence/. Before implementation, it is critical for trial courts to have a comprehensive understanding of the risks and benefits associated with each potential GenAI application. Predeployment evaluation would promote the goals of the model policy, namely responsible innovation and public trust, and is necessary to mitigate potential harms to privacy, safety, and security. Recommendation: The model policy should require courts to engage in pre-deployment analysis of GenAI uses to guide the development of their GenAI policies and to determine the appropriateness of GenAI for each use.	
			B. Require Pre and Post-Use Testing and Evaluation After the pre-deployment analysis, the model policy should require pre and post-use testing before GenAI is public-facing. The risks and benefits of each GenAI tool vary widely depending on the specific system and its intended use case. For example, if used for summarization, legal research, or translation, GenAI carries a high risk of “hallucinations,” which can, at a minimum, require substantial time to correct and, at worst, lead to inaccurate outputs that deny someone justice or their liberty. These risks may differ significantly from those associated with GenAI use in internal communications, drafting emails, or docket management. The NIST AI Risk Management Playbook emphasizes the importance of pre-use testing for identifying metrics and methods to assess risks identified in pre-deployment analysis and during operation, as well as establishing mechanisms for tracking identified AI risks over time.[3] This can include benchmarking GenAI systems for hallucinations in legal research, tracking the number of errors in data	As discussed above, this suggestion is beyond the scope of the current proposal, but the task force may consider it as time and resources permit.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			entry or docket management uses, evaluating prompts, or checking language accuracy. This type of pre-use testing can ground court staff in understanding the limitations of GenAI tools across different use cases. [3] NIST Measure 1.1 and 1.2 require that the: “approaches and metrics for measurement of AI risks enumerated during the Map function are selected for implementation starting with the most significant AI risks. The risks or trustworthiness characteristics that will not – or cannot – be measured are properly documented. [The] appropriateness of AI metrics and effectiveness of existing controls are regularly assessed and updated, including reports of errors and potential impacts on affected communities.” National Institute of Standards and Technology. (n.d.). NIST AI RMF Playbook “Measure”. https://airc.nist.gov/airmf-resources/playbook/measure/. Recommendation: The model policy should require courts to develop and implement mechanisms and metrics tailored to their specific uses of GenAI. As discussed below, if a court permits the use of GenAI, this evaluation should continue under the supervision of court staff, particularly for uses with higher risks as identified by pre-deployment analysis and testing.[4] [4] “AI systems should be tested before their deployment and regularly while in operation. AI risk measurements include documenting aspects of systems’ functionality and trustworthiness.” NIST, <i>supra</i> n.1, at 28. C. Require Clear Assignment of Responsibility and Ongoing Management The model policy should include requirements for the clear assignment of responsibility and ongoing management of GenAI systems. This includes designating specific staff or establishing a working group responsible for continuous monitoring of the effectiveness of GenAI tools and ensuring compliance with established policies. • Responsibilities should include:	
				As discussed above, this suggestion is beyond the scope of the current proposal, but the task force may consider it as time and resources permit.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			○ Ongoing evaluation and monitoring of prompts and outputs. ○ Development of best practices and refinement of court policies related to GenAI. ○ Ensuring the protection of privacy and data security. ○ Overseeing validation and review processes to ensure accuracy and reliability. ○ Establishing mechanisms for detecting and addressing the misuse of GenAI. ○ Receiving and responding to feedback from the public and court staff. As the NIST guidelines highlight,[5] effective AI risk management requires courts to assign personnel responsible for managing the use of GenAI who can regularly monitor its use to drive continual improvement, minimize harms and facilitate effective responses to errors and incidents. [5] National Institute of Standards and Technology. (n.d.). NIST AI RMF Playbook “Manage”. https://airc.nist.gov/airmf-resources/playbook/manage/. Recommendation: The model policy should require the assignment of staff oversight responsibilities to manage risk in the court’s use of GenAI.	
			D. Define Acceptable and Prohibited Uses The model policy should provide specific guidance on defining acceptable and prohibited uses of GenAI in the courts. Certain applications, such as external-facing chatbots or translation services, may present unacceptable risks where the costs of remediation, validation, and review outweigh the potential benefits. For instance, an external-facing GenAI chatbot used in New York for self-help routinely provided inaccurate or misleading legal information to its users.[6] Similarly, GenAI transcription and summarization tools have been shown to invent information as much as half of the time,[7] posing unacceptable risks to litigants if these outputs are relied upon in adjudication. Legal research tools powered by GenAI also present challenges related to the reliability and accuracy of their outputs with a Stanford	The task force is not recommending revisions in response to this suggestion. The task force concluded that courts and judicial officers are in the best position to identify acceptable uses of generative AI to meet their specific needs. It would be extremely difficult for the task force to create a list of acceptable tools and uses, and such a list would likely be both

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			study showing these systems would hallucinate answers 1 out of 6 times.[8] These examples show the potential harms that the unaccountable use of GenAI can pose to the court and public trust. [6] Lecher, C. (2024, March 29). <i>NYC's AI Chatbot Tells Businesses to Break the Law</i>. The Markup. https://themarkup.org/news/2024/03/29/nycs-ai-chatbot-tells-businesses-to-break-the-law. [7] Burke, G., & Schellmann, H. (2024, October 26). <i>Researchers say an AI-powered transcription tool used in hospitals invents things no one ever said</i>. AP News. https://apnews.com/article/ai-artificial-intelligence-health-business-90020cdf5fa16c79ca2e5b6c4c9bbb145. [8] Ho, F. S. E., Surani, F., & Ho, D. E. (2024, May 23). <i>AI on Trial: Legal Models Hallucinate in 1 out of 6 (or more) Benchmarking Queries</i>. Stanford HAI. https://hai.stanford.edu/news/ai-trial-legal-models-hallucinate-1-out-6-or-more-benchmarking-queries. Recommendation: The model policy should provide use-case-specific guidance on GenAI, with more stringent prohibitions, requirements, and oversight for use cases with greater risk of harm.	under- and overinclusive because the task force would have to speculate about how specific tools work or how courts might use them. Additionally, putting such a list in a rule of court would make it difficult to keep up with technological advancements. For these reasons, the task force recommends that the rule and standard address specific risks of generative AI rather than specific generative AI tools or uses.
			E. Define Minimum Elements of GenAI Training Programs: The model policy should mandate training for court staff and judicial officers on the responsible use of GenAI. Without detailed guidance and training, the policy risks being implemented inconsistently and ineffectively. At a minimum, these training programs should cover: ○ All elements of the court's GenAI policy, including acceptable and prohibited uses. ○ The limitations of GenAI systems, including the potential for hallucinations, biases, and errors. ○ Strategies for developing effective and unbiased prompts.	As discussed above, this suggestion is beyond the scope of the current proposal, but the task force may consider it as time and resources permit.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			○ Techniques for validating and verifying the accuracy of GenAI outputs. ○ Relevant legal, policy, and ethical rules governing the use of GenAI in the courts. ○ Strategies for addressing the risk of de-skilling due to reliance on GenAI, particularly for junior staff. Recommendation: The task force should ensure that the adoption of GenAI in California’s courts is both responsible and accountable by requiring courts to address these additional elements in its model policy and training programs.	
			III. The Model Policy Should Set Equivalent Standards for Judicial Officers. To maintain public trust in the integrity of the judicial process, we recommend that judicial officers, even in their adjudicative roles, be held to the same or substantially similar standards as court staff when using GenAI. The current discrepancy between rule 10.430 and standard 10.80 creates a double standard, which could erode public confidence in the courts. The use of GenAI to support adjudicative decision-making can pose a greater risk to the fairness of proceedings and the rights of parties within the court system. Therefore, the stronger, binding requirements of rule 10.430 should apply to judicial officers using GenAI in key adjudicative functions. Recommendation: Apply the requirements of rule 10.430 to judicial officers, given the unique risks posed by GenAI in adjudicative functions, which necessitate a clearly defined floor for responsible use.	The task force is not recommending revisions in response to this suggestion. The task force is recommending standard 10.80, which is discretionary, because it concluded that the question of whether and to what extent judicial officers may use generative AI to carry out their adjudicative duties is more appropriately addressed by the California Code of Judicial Ethics and related ethical guidance.
			IV. The Model Policy Should Provide Additional Guidance on Disclosure Requirements The model policy should require courts to create specific guidance and definitions around disclosure requirements for court documents produced with GenAI. Currently, rule 10.430(d)(5) and standard 10.80(b)(5) call for disclosure if GenAI outputs constitute a “substantial portion” of the content. This standard is too vague	In light of all the comments received on this issue, the task force has recommended a revised version of rule 10.430(d)(5).

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			and subjective and can lead to substantially different policies and therefore differential access to due process across court systems. Furthermore, the policy should provide examples of the type, placement, and level of disclosure required, with an understanding that these requirements may need to vary depending on the type of court document and the AI's role in its creation. Detailed and specific guidance can ensure that the use of GenAI in the creation of court documents is transparent, accountable, and consistent across all California courts, maintaining public trust in the integrity of the judicial process. Recommendation: The policy should define clear thresholds and examples for when disclosure is required including for pre-use notice.	
			V. The Model Policy Should Be Consistent with California Law Recommendation: To ensure clarity and consistency, the model policy should align its definitions of “Generative artificial intelligence system” and “Artificial Intelligence” with those used in existing California law, such as in CA Civ Code § 3110 (2024).[9] [9] (a) “Artificial intelligence” means an engineered or machine-based system that varies in its level of autonomy and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments. (c) “Generative artificial intelligence” means artificial intelligence that can generate derived synthetic content, such as text, images, video, and audio, that emulates the structure and characteristics of the artificial intelligence’s training data. CA Civ Code § 3110 (2024).	In light of all the comments received on this issue, the task force is recommending revised definitions of “generative AI” and “public generative AI system” in the rule and standard. The task force recommends deleting the definition of “artificial intelligence” in the rule and standard because it is not needed due to the proposed revisions to the definition of “generative AI.” However, although the task force agrees that consistency with statutory definitions can be beneficial in some circumstances, the existing statutory definitions, such as those in Civil Code section 3110, are not a good fit

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01

Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
				for the rule and standard. Those definitions are part of statutory schemes for regulating AI providers and use terminology that laypeople might find confusing, such as the reference in section 3110(a) to “explicit and implicit objectives.”
			Conclusion: We ask that the task force adopt the above recommendations to ensure that the use of GenAI in California’s courts is consistent with the best practices on responsible AI deployment. We urge the task force to adopt these changes to safeguard public trust, promote a justice system that serves all Californians equitably, and ensure that the deployment of GenAI has proper guardrails and oversight, particularly when it comes to decisions that impact access to justice.	No response required.
16.	Superior Court of California, County of Los Angeles by Stephanie Kuo	AM	The following comments are representative of the Superior Court of California, County of Los Angeles, and do not represent or promote the viewpoint of any particular judicial officer or employee. In response to the Judicial Council of California’s ITC SP25-01 Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work, the Superior Court of California, County of Los Angeles (Court), agrees with the proposed additions if the following modifications are incorporated.	Please see responses to the Superior Court of California, County of Los Angeles’s specific suggestions below.
			The Court commends the thoughtful approach reflected in much of the rule, as it seeks to balance the benefits of Artificial Intelligence (AI) tools with the need for accountability, transparency, and integrity in judicial administration. Our organization is broadly supportive of the framework outlined; however, we wish to	No response required.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			express specific concerns regarding subsection 10.430(d)(5) and offer recommendations for its refinement.	
			Subsection 10.430(d)(5) mandates that staff and judicial officers (exercised outside of their adjudicative role) disclose their use of or reliance on generative AI when a substantial portion of the content in the final version of a written or visual work—intended for public dissemination—is produced by such technology. While we recognize the intent to promote transparency, we respectfully submit that this disclosure requirement extends beyond what is necessary and introduces practical challenges. We propose that the language in Standard 10.80, which applies to judicial officers in their adjudicative functions, offers a more balanced approach. Standard 10.80 encourages judicial officers to exercise discretion in determining whether disclosure is appropriate on a case-by-case basis, rather than imposing a blanket obligation. We believe a similar standard would better serve the objectives of Rule 10.430. One of the Court’s highest concerns is that—provided all other safeguards such as mandatory review, verification of accuracy, and elimination of bias—are adhered to, an additional disclosure requirement appears redundant. For example, consider a scenario in which AI is used to initiate a draft message for a Bar publication or to draft routine correspondence like an email response. Requiring disclosure in such instances could impose an undue administrative burden without meaningfully enhancing public trust or accountability and would likely discourage the use of generative AI in instances when it is appropriate. These uses of AI are analogous to the use of other commonly used tools, such as image generation, word processing software or legal research databases, none of which trigger similar disclosure obligations. Moreover, the term “substantial” in subsection 10.430(d)(5) lacks clear definition, rendering the provision susceptible to inconsistent application. Without explicit guidance on what constitutes a “substantial portion” of AI-generated content, staff and judicial officers may face uncertainty, potentially leading to over-reporting or	In light of all the comments received on this issue, the task force has recommended a revised version of rule 10.430(d)(5). The task force understands the concern that requiring disclosure might limit courts’ flexibility but concluded that mandatory disclosure is necessary in some circumstances to maintain public trust in the judicial branch.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			inadvertent non-compliance. To address this, we recommend either replacing the term with a more precise threshold or adopting the discretionary framework of Standard 10.80, which avoids such ambiguity altogether.	
			Regarding the implementation timeline, six months is more reasonable when considering the different stages at which courts may be with their AI usage. Also, it would give the courts more time to implement tools to enforce the policy, which would be needed. Two months is not enough time to implement the tools.	The task force recommends revising rule 10.430(b) to require courts to adopt use policies by December 15, 2025, in order to give courts more time to prepare their policies and any tools necessary for implementation.
			It is difficult to assess how this proposal would be developed and implemented in courts of different sizes. We believe it would be more difficult for larger courts to implement this proposal, simply because of the level of development and use of AI-related tools and applications and, as a result, the larger number of staff currently using AI in larger courts. Larger courts will need tools to enforce the policy whereas smaller courts may not have enough activity to warrant the purchase of these tools.	The task force appreciates the response.
			In conclusion, while the Court supports the majority of Proposed Rule 10.430 and its commitment to responsible AI use, we urge reconsideration of subsection 10.430(d)(5). Aligning this provision with the discretionary approach of Standard 10.80 would enhance the rule's clarity, practicality, and effectiveness.	No response required.
17.	Superior Court of California, County of Placer by Naslie Rezaei, Court Services Analyst	AM	Thank you for the opportunity to comment on the proposed rule and model policy concerning generative AI. With the rapid use and integration of generative AI, in both the courts and daily life, we appreciate the Committee's dedication in addressing this emerging policy area. The Placer Superior Court, Court Administration (Court Administration or we) largely agrees with the Proposal, but we would like to offer a few suggestions for the Committee's consideration and one request for amendment.	Please see responses to specific suggestions below.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			Model Use Policy, Section VI. Transparency: We understand the Committee's intention to promote transparency using watermarks or disclaimers on AI-generated photos, videos, and audio clips. While Court Administration agrees differentiating between deepfakes is important, we ask that the Committee consider the integration of generative AI tools in word processing applications. As generative AI tools increasingly integrate with word processing applications (such as the launch of CoPilot for the Government Cloud), these tools will likely be used to simplify the creation of templates, documents, or memoranda. For instance, a court might generate a memorandum template in response to a citizen complaint using generative AI. While the initial draft of this memorandum may be AI-generated, a subsequent review would be conducted by court staff to ensure accuracy and completeness (as required by proposed Rule 10.430). With this additional review, it may not be necessary to indicate, to the reader, that generative AI was used to refine or speed development of written material. This is differentiated, perhaps, from legal writing that is submitted to the court where the submitting party may not be subject to rules or policies that require human review of written material. Ultimately, we ask that the committee consider revising this section to require disclosure of AI-generated photos, videos, and audio clips but remove the requirement on written publications that must be reviewed by the author and can be directly modified.	In light of all the comments received on this issue, the task force has recommended a revised version of rule 10.430(d)(5). The task force understands the concern that requiring disclosure might limit courts' flexibility but concluded that mandatory disclosure is necessary in some circumstances to maintain public trust in the judicial branch.
			Generative AI FAQs (in response to the Committee's question about issues courts would like addressed in the upcoming Generative AI FAQs): We appreciate the Committee's willingness to create resources to aid courts in developing and deploying generative AI. In addition to FAQs that address generative AI, we believe it would be helpful for the Committee to define other traditional forms of AI—such as advanced automation, machine learning, and natural language process models that do not independently generate text. This clarification may reduce confusion as courts move to implement other forms of AI, beyond generative AI, that are not covered by these new rules.	The task force will consider these suggested revisions to the FAQs as time and resources permit.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			Implementation Deadline (in response to the Committee's question about a two-month implementation deadline): Placer Court Administration agrees that a two-month deadline to adopt a generative AI use policy is sufficient. However, it may be helpful to include a grace period for compliance with adopted policies to allow courts time to coordinate with their vendors. Some software or software-as-a-service providers have already integrated AI into their applications. A six-month grace period, for example, would allow courts time to review these applications and work with their vendors to conform their solutions to this new policy.	The task force recommends revising rule 10.430(b) to require courts to adopt use policies by December 15, 2025, in order to give courts more time to prepare their policies and any tools necessary for implementation.
18.	Superior Court of California, County of San Francisco by Michael Corriere, Chief Data Officer	AM	Rule 10.430 (b)(1) and Model Policy III. a. (also seen in Standard 10.80 (b)(1)) A. Replace “nonpublic” with “non-disclosable.” Nonpublic information is not necessarily confidential. Some courts with data classification systems identify a level of data sensitivity that is neither public nor confidential and is disclosable (i.e., “internal” in San Francisco). Information that is nonpublic but not confidential should be a permissible input.	The task force appreciates the commenter’s concerns but is not recommending revisions in response to this suggestion. The task force expects that inputting nonpublic information into public generative AI systems may be problematic even if the nonpublic information is not confidential. However, the task force is recommending a revised definition of “public generative AI system,” which might resolve the commenter’s concern.
			B. Allow exception to input of confidential information into a public generative A.I. system when the court has a service-level agreement that: a. prohibits the public generative A.I. system provider from 1) allowing court data to be accessed by anyone outside of the court, 2) using court data to train A.I. systems, or 3) allowing others outside the court to use court data to train A.I. systems, and b. contains language ensuring that this prohibition survives the termination of the agreement indefinitely.	In light of all the comments received on this issue, the task force is recommending revised definitions of “generative AI” and “public generative AI system” in the rule and standard.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
			Rule 10.430 (d)(5) and Model Policy VI. A. (also seen in Standard 10.80 (b)(5)) “Substantial” will require some clarification/interpretation, either within the rule/model policy or at the local level. Including a discussion of this in the FAQ would also be helpful.	In light of all the comments received on this issue, the task force has recommended a revised version of rule 10.430(d)(5).
			Disclosure of the use of Gen A.I. in creating materials presented to the public will not be possible without tracking the use of this technology from the materials’ creation, as content gets cycled through various drafts and is repurposed over time. Disclosure by contractors to court staff will also be necessary to comply with this section of the rule, as the court would have no way of knowing what content it must identify as created with the help of A.I. if the contractor does not disclose. Recommend striking “provided to the public” as this will have to be done anyway to make this policy workable.	The task force understands the concern that requiring disclosure might limit courts’ flexibility but concluded that mandatory disclosure is necessary in some circumstances to maintain public trust in the judicial branch.
			For the FAQ, we suggest that it include guidance on how to determine that a Gen A.I. output is not a hallucination.	The task force will consider this suggested revision to the FAQs as time and resources permit.
			Rule 10.430 Judicial leadership is a bit concerned with the notion of judges notifying parties of their use of AI for research or any other purpose. Standard 10.80 Judicial leadership has the same concern mentioned above for judges under the standard.	The task force’s recommended revisions to rule 10.430(d)(5) should exclude typical legal research activity from the rule’s mandatory disclosure requirement.
			As for judges, recommend that we adopt the policy rule and standard with the limitations mentioned above as a local rule. We should also determine whether we want to add anything further to these model rules as they do seem a bit general.	No response required.
19.	Trial Court Presiding Judges Advisory Committee (TCPJAC) and the Court	AM	The Joint Rules Subcommittee (JRS) of the Trial Court Presiding Judges Advisory Committee and Court Executive Advisory Committee submits the following comments in response to the invitation to comment (ITC) on SP25-01 Judicial Branch Administration: Rule and Standard for Use of Generative Artificial	Please see responses to TCPJAC/CEAC Joint Rules Subcommittee’s specific suggestions below.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
	Executives Advisory Committee (CEAC) (TCPJAC/CEAC Joint Rules Subcommittee)		Intelligence in Court-Related Work (Adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)	
			JRS agrees generally with the proposed if modified. The subcommittee proposes the following changes:	
			<u>[1] Change the disclosure requirement in CRC, rule 10.430 (d)(5)</u> Instead of being required to disclose the use of generative AI if the work product constitutes a substantial portion of the content, the rule should only require judicial officers and court staff to consider whether disclosure should occur. With necessary edits due to differences in phrasing in the rule, the proposed language in rule 10.430 (d)(5) should be replaced with the proposed language in Standard 10.80 (b)(5). This change would allow for a more appropriately flexible approach, accommodating various contexts in which AI might be used. If mandatory disclosure is kept in this section of the rule, “substantial” needs to be defined as it is ambiguous and susceptible to various interpretations.	In light of all the comments received on this issue, the task force has recommended a revised version of rule 10.430(d)(5). The task force understands the concern that requiring disclosure might limit courts’ flexibility but concluded that mandatory disclosure is necessary in some circumstances to maintain public trust in the judicial branch.
			<u>[2] Remove the word “unlawfully” from CRC, rule 10.430 (d)(2)</u> The word "unlawfully" should be removed from the rule language because it is redundant and unnecessary. This section of the rule already states that generative AI may not be used to discriminate based on any classification protected by federal or state law. Further, the word "unlawfully" does not add any additional value and including it might imply that forms of discrimination that are not explicitly unlawful are permissible.	The task force appreciates the commenter’s concerns but is not recommending revisions in response to this suggestion. The task force acknowledges that it is not strictly necessary to prohibit the use of generative AI to unlawfully discriminate because such discrimination is already prohibited. The task force included this provision in the rule primarily to ensure courts are aware of the risk that generative

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.

SP25-01**Judicial Branch Administration: Rule and Standard for Use of Generative Artificial Intelligence in Court-Related Work** (adopt Cal. Rules of Court, rule 10.430; adopt Cal. Standards of Judicial Administration, standard 10.80)

All comments are verbatim unless indicated in a footnote.

	Commenter	Position	Comment	Committee Response
				AI systems can produce biased or discriminatory outputs.
			<u>[3] Add a definition for “adjudicative role” in CRC, rule 10.430 and Standard 10.80</u> Some judicial officers may consider most tasks they perform as part of their adjudicative role, while others might only consider tasks directly related to making judicial decisions in cases as part of their adjudicative role. The proposal as drafted assumes a common understanding or definition exists, which is not the case. This ambiguity would inhibit the correct application of the rule and standard, as it leaves it unclear when the requirements in the rule for generative AI use would apply to judicial officers as opposed to the suggestions in the standard. The ambiguity would also inhibit the ability of presiding and supervising judges to enforce the rule with other judicial officers.	The task force acknowledges that the term “adjudicative role” may seem vague but is not recommending revisions in response to this suggestion. The task force concluded that it is appropriate to leave the term undefined to avoid potential conflicts with other rules or guidance that use similar terms. Additionally, judicial officers have discretion and are best situated to determine whether a task is within their adjudicative role.
			The Joint Rules Subcommittee thanks the Artificial Intelligence Task Force, and staff for the opportunity to review and provide commentary on this proposal.	No response required.

Positions: A = Agree; AM = Agree if modified; N = Do not agree; NI = Not indicated.