

# Brainstorming Group Outline: Artificial Intelligence and Machine Learning in Discovery

*The Sedona Conference*

March 2024

Copyright 2024. All rights reserved.

---



# WG1 Artificial Intelligence and Machine Learning in Discovery Brainstorming Group Outline

March 2024

## Brainstorming Group Members

Steven Berrent	Jonathan Bryant
Xavier Diokno	Adam Holter
Mark Lanterman	Dave Lewis
Jason Lichter	Leeanne Mancari
Miriam Marks	Bim Onafuwa
Hon. Xavier Rodriguez	

## Team Leaders

Kelly Atherton	Amy E. Sellars
----------------	----------------

## Steering Committee Liaisons

Lea Bays	Tara S. Emory
Tessa K. Jacob	

## DISCUSSION TOPICS

The Brainstorming Group seeks feedback on the following topics during the midyear meeting:

1. From the court's perspective, what type of guidance from Sedona on AI in discovery would be helpful to the courts to encourage the acceptance and use of AI in discovery?
2. Should Sedona take a position on whether counsel should obtain consent before using GenAI for document review?
3. Should Sedona take a position on what validation should look like when using GenAI for document review?
4. Should Sedona take a position on whether we need judicial approval to use GenAI for document review?
5. Should Sedona take a position on whether court issued standing orders on AI should be reconsidered?
7. Did we miss any key content in our outline for the primer and jumpstart guide appendix? Did we get the organization right?
8. We considered the preservation of AI (prompts, outputs, and everything in between) and the issues it poses for discovery. We think it's a topic worth of discussion, but did not address in this outline. Should we have?

## SUMMARY AND RECOMMENDATION

- A. The WG1 Brainstorming Group on Artificial Intelligence and Machine Learning in Discovery (“BG”) was formed to consider whether guidance may be beneficial on the use of machine learning (“ML”) and artificial intelligence (“AI”), including Generative AI (“GenAI”), in discovery. Specifically, the BG was asked to consider the following topics:
1. When, and how, may AI and ML be appropriate and useful in the discovery process? When might they be less appropriate or useful?
  2. To what extent – now and potentially in the future – do AI and ML overlap with TAR or other technology-based tools and processes, and our prior guidance thereto?
  3. Are there recommendations around AI and ML that may help encourage their use and acceptance by courts and practitioners, if and when appropriate? For example, a code of conduct, guidelines, or a similar framework for use of AI and ML. Will looking back at our experience with the developments around law and practice for TAR be helpful to consider?
  4. AI and ML in relation to discovery topics, including those that may have been subjects of prior Working Group 1 guidance, such as Sedona Principle 6, FRCP required disclosures, and evidentiary issues.
  5. Ethical requirements of lawyers relating to AI and ML, including certification under FRCP 26(g) and ABA model code (and state bar) requirements on supervision and technology competence, and the use of AI and ML tools with client information.
  6. What format and type of guidance from Working Group 1 would be helpful to courts and practitioners? For example, a document providing guidelines, a primer, and/or educational outreach and resources? And how would WG1 maintain the relevance of such guidance in light of the rapidly evolving AI landscape?
- B. The BG believes guidance on the use of AI and ML, including GenAI, in discovery would be beneficial and makes the following recommendations:
1. To form a drafting team to:
    - a. Issue a primer on the use of AI in discovery.
    - b. Issue a supplementary jumpstart guide on GenAI in discovery that can more easily be updated and:
      - i. Focuses on the established and emerging use cases of GenAI in discovery.
      - ii. Provides a framework for crafting effective prompts for GenAI.
- The below outline sets forth the BG’s recommendation on the scope and guidance for the primer and jumpstart guide.

2. To prioritize speed of publication given the need for an organization like The Sedona Conference to issue a publication that educates courts and practitioners on AI in discovery, including GenAI.
3. To liaise with the Technology Resource Panel on the technical aspects of the primer and appendix.
4. To update The Sedona Conference Glossary: eDiscovery & Digital Information Management, Fifth Edition, particularly to include related GenAI terms.
5. To use consistent terms and definitions across all Sedona Conference publications regarding AI issues.

## **PRIMER INTRODUCTION**

### **A. Goals of the primer**

1. To analyze the appropriateness of established and emerging AI tools in the context of existing ethical rules and rules of practice, including applicable rules of civil procedure.
2. To learn from our experience with the developments around law and practice for TAR and to have a smoother integration around law and practice for GenAI and other future AI tools.
3. To encourage the responsible and defensible adoption of AI tools in discovery.
4. To encourage courts to consider whether disclosure on the use of AI is necessary, what disclosure would be recommended, and in what situations.
5. To educate the courts and practitioners on what GenAI is, how GenAI works, and how GenAI can be used in discovery.

- B. Practitioners have been using AI in discovery for years and recent developments in AI promise a new suite of tools that practitioners anticipate will be useful in the discovery process.

## **TERMS AND CONCEPTS**

### **A. Definitions:**

1. Clarify definition/scope of terms used in the primer, including high level explanations of how different AI models work, are similar and different from one another and can overlap, and to differentiate between technical and colloquial definitions.
2. Recommend adding a Venn diagram to illustrate the concepts.

## **WHEN AND HOW AI MAY BE APPROPRIATE AND USEFUL IN THE DISCOVERY PROCESS**

- A. A user should consider the output of an AI tool in deciding whether the tool is appropriate and useful.

1. Users can evaluate appropriateness and usefulness by considering both where the outputs can be placed in the Discovery process, such as is modeled by the EDRM; and what form the outputs take (groups of text, identification of portions of text, summaries, or answers to questions).
2. Discuss how newer forms of AI, including generative AI, can be applied—or introduce new risks/challenges—at each phase of EDRM (near-term v. medium-term v. long-term). This section discusses where we think AI may go in the future and the below applications may not be in use or may not have widespread adoption as of now.
3. Discuss ways in which AI has been part and parcel of technology-related tools and processes in use in discovery for years
  - a. TAR
  - b. Natural language-based queries and statistical ranked retrieval
  - c. Conceptual search and query reformulation
  - d. Clustering and other unsupervised learning methods
  - e. Sensitive content (PII and PHI) identification and masking
  - f. Sentiment analysis
  - g. Entity recognition
  - h. Machine translation
  - i. Other (subject to drafting team)
4. Discuss examples of proposed applications of recent AI developments to discovery, with varying degrees of actual use or implementation. These examples are not exhaustive and may not be practical or even permitted depending upon circumstance.
  - a. Identification
    - i. AI tools can summarize a record (and, increasingly, groups of records). Provided the summary is accurate, this might accelerate acquiring an understanding of the key points at issue, and which sets of documents need collection and further review. Several vendors have shown versions of tools for automated timeline generation and other higher-level perspectives on sets of documents.
    - ii. Aiding in generation of keyword queries.
  - b. Preservation
    - i. Using AI to preserve. AI tools may be able to review the records on a server or other data-storage device and implement a legal hold based on metadata such as custodian and created/modified dates or the content of records. Counsel should continue to supervise

the process and interview custodians to verify the appropriate records are preserved.

- ii. While AI may be used to find the set of data to be preserved, requesting and responding parties should note they have the same preservation obligations regardless of whether AI was used.
- iii. **[Drafting team:** We considered the preservation of AI (prompts, outputs, and everything in between) and the issues it poses for discovery. We are not addressing it in this outline, but think it's a topic worthy of discussion.]

c. Collection

- i. Using AI to collect (we are not aware of any tools at this time, but this may change and is subject to drafting team)

d. Processing

- i. Use of generative AI to produce new types of computed metadata, including assignment of categories, key points, and summaries.
- ii. If large language models are trained on (or tuned to) the collected materials for a case, consider what changes to processing are necessary to achieve high quality results.

e. Review and Analysis

- i. Review is the area in which AI currently offers the most benefit: document production, deposition preparation, affirmative story, key themes, interrogatory responses, etc.
- ii. Vendors are already fielding generative AI tools that promise to code for responsiveness, issues, confidentiality, and privilege, given appropriate prompting. To what extent these tools will be substituted or combined with traditional TAR, and to what extent they will be viewed differently by the legal system, is currently unclear.
- iii. Redactions (all types)
- iv. Summarize and extract key themes from documents and sets of documents. For example, analyzing incoming and outgoing productions and generating timelines. This may substantially change what is expected from manual reviewers.

f. Production

- i. Vendors have built initial generative AI tools for creating privilege log descriptions, though hallucinations and errors may still be a challenge.

g. Presentation (including Witness/Deposition Preparation)

- h. As at the Identification and Review stage, AI generated summaries, narratives, chronologies, and timelines could play a role. Stakes for correctness become higher if these are used in witness preparation or are presented in court.
  - i. Advocacy/Negotiation
  - ii. What needs to be disclosed (and/or discussed in ESI protocols) about the use of generative AI in the discovery process. Should generative AI declarations / affidavits be required beyond what is already required in the rules?
  - iii. Use of generative AI in drafting/analyzing ESI protocols, protective orders, etc.
- B. Considerations in deciding whether an AI tool is appropriate and useful.
  - 1. The line between appropriate and useful on one hand and neither appropriate nor useful on the other hand is likely to move. In addition to rapid developments in technology, user acceptance strongly affects whether a tool is appropriate and useful.
  - 2. It may be helpful for the user to consider tools, their outputs, and the quality of their outputs when deciding what tools are useful and appropriate. For example, a tool that drafts a privilege log may not be the most appropriate tool to use to conduct a search for collecting relevant records.
  - 3. Cost and volume may be factors when determining whether a tool is appropriate.
  - 4. Defensibility.
    - a. Appropriate and useful may or may not be “defensible.” Separate but related considerations.
    - b. Whether/what level of validation is needed will depend on the use case, i.e., internal knowledge management versus review and production.
    - c. How can the tool be validated? Technical validation methods (such as random sampling and comparison with expert standards) will need to be extended to handle new types of AI outputs.
    - d. What kind of disclosure may be required?
    - e. Judicial acceptance

## **WHAT WE LEARNED WITH TAR CAN BE LEVERAGED TO ASSESS WHETHER AND HOW GENAI MAY BE USEFUL IN DOCUMENT REVIEW**

- A. It was said on a *Sedona Conference Working Group 1 Virtual Town Hall on GenAI and eDiscovery* that the incorporation of GenAI to existing TAR workflows represents “new wine in an old bottle.” Of course, there are different ways in which the new wine might replace—or be blended with—the old, and some may prove to be more successful than others.
- B. Overview of similarities and differences between TAR and GenAI for document review.



1. Training
  2. Review
  3. Validation
- C. Considerations for practitioners to reference (and questions to answer) when determining how, if at all, to leverage GenAI in the document review context.
1. Is client consent required, whether pursuant to explicit outside counsel guidelines or otherwise? To the extent affirmative consent is not required, disclosure to clients of counsel's intent to use GenAI is certainly recommended.
  2. Absent a court order or explicit protocol, must counsel obtain consent from opposing counsel, courts, or regulators before using GenAI to facilitate document review? **[Drafting team:** Should Sedona advocate for a position one way or the other?]**]** The answer may depend on whether GenAI is being used in a manner that might limit which documents will ultimately be reviewed for production. If consent is required, opposing counsel, courts, or regulators may be less likely to grant that consent if they are unfamiliar or uncomfortable with GenAI. In those circumstances, what is the most effective and appropriate way to educate them on the nature and suitable uses of the technology and to facilitate consent?
  3. On the flip side, can a party be compelled to use GenAI over its own objection? **[Drafting team:** Consider both pre-collection in the party's environment and post-collection once the data set has been identified.]**]** To date, the weight of authority has come down against mandating a reluctant party to use traditional TAR, and courts may be hesitant to mandate the use of GenAI for similar reasons. Reference to Sedona Principle 6. But cost-shifting and proportionality/burden arguments may be unavailable to parties who does not adopt cost efficient processes.
  4. If consent is not required, should parties nonetheless disclose their intent to use GenAI during the meet-and-confer process and/or seek to define the manner of usage in an discovery protocol/order? **[Drafting team:** Consider both pre-collection in the party's environment and post-collection once the data set has been identified.]**]** Just as with many other aspects of discovery, it is important to consider leaving room for adjustment and refinement in any stipulation/protocol.
  5. Which flavor of GenAI is being used? GenAI is not a monolithic technology; there are many different implementations of GenAI, and a multitude of large-language models ("LLMs") that underlie those implementations. Adding to the complexity, specific LLMs are not necessarily static/fixed and vendors are continually refining their models, without necessarily alerting users or allowing users to opt out of the upgrades.
  6. GenAI models are inherently probabilistic/non-deterministic, meaning that they can produce different results when presented with the same inputs multiple times. This can present challenges given the iterative and evolving nature of document review, which takes place over a period of time. Practitioners may need to

determine how they are going to subsequently explain or defend document coding decisions that were informed by GenAI.

7. Most current applications of GenAI to document review require the up-front preparation of a description of the case and definitions of the contours of responsiveness and issue tags, akin to what is customarily included in a document review protocol used to guide first-pass review attorneys. But those definitions are often updated during the course of the review, with coding determinations subject to retroactive change. How will those refinements to the scope of responsiveness and issue designations be made if using GenAI to accelerate first-pass review?
8. Will GenAI be used in lieu of, or alongside, TAR 2.0? GenAI could theoretically be used *before* TAR 2.0, to identify a core set of highly responsive documents suitable to give the TAR 2.0 algorithm a jumpstart. GenAI could also be used *after* TAR 2.0, to summarize the documents receiving the highest TAR 2.0 responsiveness scores so that case teams can more efficiently triage which require more detailed examination. Or GenAI could replace TAR 2.0 altogether.
9. How do concerns about algorithmic bias in GenAI models apply in the document review context, and how can those concerns reasonably be addressed? Since a human reviewer could also be biased, should GenAI be held to a higher standard?
10. GenAI technologies also have the potential to hallucinate, or generate erroneous outputs. There are various ways to mitigate the risks associated with mistakes or hallucinations, including forcing the tool to explain its reasoning and include citations in its outputs, but none of these techniques is infallible. Of course, humans also make mistakes, as does traditional TAR. Are the recall, precision, elusion, and other statistical validation techniques widely used today in document review readily adaptable to GenAI, or must new validation rubrics be developed?
11. Whenever client data is shared with a third-party, confidentiality, privacy, and information security issues come to the fore. But LLM-powered GenAI tools present particularly acute concerns in the discovery context given their potential to ingest, learn from, and potentially leak highly confidential and/or privileged information to unrelated users. Service providers and platforms should be carefully evaluated, and contractual terms negotiated, to guard against these risks and ensure that counsel's duties of competence and confidentiality are assiduously maintained. These issues are also applicable where joint document repositories are contemplated.

#### D. TAR Experience – Lessons Learned for GenAI

1. We can learn from our experience with the developments around law and practice for TAR to have a smoother integration around law and practice for GenAI and other future AI tools. Some practitioners and clients are still reluctant to use TAR tools because of the potential added cost of negotiation and motion practice related to their use.
2. Challenges to acceptance

- a. Judicial acceptance. Does there need to be judicial scrutiny/approval to use GenAI for review? Or does *Rio Tinto*/TAR cover us?
  - b. Education and technological awareness
    - i. Courts and practitioners need familiarity with the tools and techniques in order to discuss AI and GenAI in Rule 26(f) conferences and reports.
    - ii. At the time of this outline, several courts have issued orders prohibiting the use of GenAI (or even AI altogether). Will prohibitions that courts have entered apply to discovery? This may chill innovation in discovery that could further the goals of Rule 1. **[Drafting team:** Consider whether Sedona should take a position on whether court issued standing orders on AI should be reconsidered.]
  - c. Concern over prompt discovery – are prompts protected work product? Do we have to retain all prompts used and the succession in which they were used? Will prompts become a source of contention? How do we measure the efficacy of prompts? How can we monitor improvements in results with successive prompts or better engineered prompts over time? **[Drafting team:** How would concern be shaped by human authored prompts as opposed to AI optimized prompts?]
3. Misconceptions of GenAI capabilities
    - a. Not a “magic button” – requires planning, knowledge, and skill. Often involves partnership with vendors and technologists. AI does not alleviate all of the burden or cost of identification, collection, and search.
    - b. When using any new tool, practitioners should allow the time necessary to evaluate the efficacy of the tool and be prepared to iterate as needed.
    - c. Concern that requesting parties will ask that companies search across “ALL data.” The Rules provide a framework for considering the relative cost burden and benefit of any approach.

## THE EXISTING RULES OF PRACTICE AND ETHICAL RULES PROVIDE A FRAMEWORK FOR USING AI IN DISCOVERY

- A. The Federal Rules of Civil Procedure provide the guidelines and framework for discovery, including discovery using established and emerging AI or other technology tools. **[Drafting team:** We expect that references to the existing Federal Rules of Civil Procedure and ABA Model Rules of Professional Conduct will be addressed in the various sections above.]
- B. Consider referencing other Sedona papers that review the application of the Federal Rules to TAR or discovery (e.g. Rule 34 papers).
  1. Sedona Principles 3 and 6

2. The Sedona Conference Cooperation Proclamation
  3. The Sedona Conference Best Practices Commentary on the Use of Search & Information Retrieval Methods in E-Discovery
- C. The existing ethical rules are sufficiently broad to apply to established and emerging AI technologies used in the discovery process.
1. ABA Model Rules of Professional Conduct
    - a. Confidentiality – Whether and how confidential client information may be used to train AI models or as prompts. Rule 1.6 and comments governs an attorney’s general duty of confidentiality, which requires an attorney to “make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating to the representation of a client.”
    - b. Competence. Rule 1.1 and comments require attorneys to provide competent representation.
      - i. Technological Competence. Comment 8 – keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology, such as AI tools.
    - c. Supervision. Rules 5.1, 5.3, 5.4, and 5.5 and comments hold attorneys responsible for the actions of others, including their attorney and non-attorney subordinates. Law firm and in-house attorneys must provide adequate oversight over any work done by litigation support and discovery service providers, including the use of AI tools in the discovery process (e.g., review or privilege logging).
    - d. Communication. Rule 1.4 governs an attorney’s duty to communicate with clients and requires an attorney to “reasonably consult with the client about the means by which the client’s objectives are to be accomplished,” including discussing with their client the decision to use (or not use) AI or why AI generates a particular outcome and the benefits and risks of such use.
    - e. Candor. Rule 3.3 imposes an ethical obligation to demonstrate candor to courts and other tribunals and prohibits “mak[ing] a false statement of fact or law . . . or fail[ing] to correct a false statement of material fact or law previously made.” Citing nonexistent case law or misrepresenting the holdings of a case based on output from GenAI is making a false statement to a court. Validate AI outputs and correct any incorrect outputs. Check if there are any requirements in the relevant jurisdiction that may require the disclosure of the use of GenAI in the discovery process.
    - f. Bias. Rule 8.4(g) on professional misconduct prohibits discrimination by attorneys against certain protected classes. Training data for AI models may be biased, potentially leading to biased outputs.

- g. Fees. Rule 1.5. If there is an opportunity to leverage AI in the discovery process to reduce billable hours, does any attorney have a duty to? Fee agreement should explain fees and costs associated with the use of AI.

## 2. State Bar Associations

- a. Competence. Ethics committees in [39] states have agreed this duty includes a requirement that attorneys stay informed and up-to-date on current technology and how that technology can be leveraged to provide the best results for their clients.
- b. California Bar – On 11/16/23, approved the [Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law](#).
- c. Florida:
  - i. FL Bar issued an advisory ethics opinion on using GenAI in the practice of law on 1/19/24. [Florida Bar Ethics Opinion 24-1](#).
    - 1. Lawyers remain responsible for their work product and professional judgment and must develop policies and practices to verify that the use of GenAI is consistent with ethical obligations.
    - 2. Lawyers should be mindful of the duty to maintain technological competence and educate themselves regarding the risks and benefits of new technology.
  - ii. The Florida Bar also formed a [Special Committee on AI Tools & Resources](#) to evaluate legal AI tools and their best uses in ways compliant with ethical duties.
  - iii. Task force formed in Illinois, [New York](#), Texas, New Jersey, Kentucky, Minnesota

## APPENDIX – GENAI JUMPSTART GUIDE

- A. Goals of the GenAI jumpstart guide
  - 1. To provide actionable information on how to use GenAI in discovery now and potentially in the near future.
  - 2. To encourage experimenting with the technology and incorporating the technology into one's discovery practice, including highlighting low risk use cases and providing current best practices to validate the output based on the use case.

[Drafting team: The BG anticipates the appendix may quickly become outdated and will need to be updated.]
- B. Established and emerging use cases for GenAI in discovery, including risks involved and current best practices to validate the output.
  - 1. Summarize information
    - a. ECA

- b. Assess incoming and outgoing productions
  - c. Deposition prep
- 2. Generate timelines
- 3. Generate deposition questions
- 4. Generate initial drafts of discovery requests, responses, and objections
- 5. Document review
- 6. Privilege identification
- C. Framework for crafting effective prompts

DRAFT