Artificial Intelligence

Guidance for Use of AI and Generative AI in Courts

August 7, 2024

from the **AI Rapid Response Team** *at the National Center for State Courts*

COSCA
Conference of State Court Administrators

CONFERENCE OF CHIEF JUSTICES 1949

NCSC
National Center for State Courts

Table of Contents

Introduction and Background

The Artificial Intelligence Rapid Response Team (AI RRT) is a project of the Conference of Chief Justices (CCJ) and the Conference of State Court Administrators (COSCA), and supported by the National Center for State Courts (NCSC).

The AI RRT was established to help courts plan for the impact that Generative Artificial Intelligence (GenAI) may have on the courts. GenAI is rapidly evolving and has the potential to change the practice of law and how courts operate. As with many new technologies, it is imperative that the courts become informed consumers of GenAI. The AI RRT has spent the past eight months examining this issue. As part of its work, the AI RRT has published seven (7) interim guidance documents for the courts and created a **resource center** that includes a landscape of court orders, rules, guidance and other initiatives of the state court community and the federal courts regarding AI or GenAI. The AI RRT conducted a survey and follow-up survey of state activities and published the results on **NCSC's AI website** (ncsc.org/ai). The information provided in this document is intended to help get courts started on their GenAI journey. State Court leaders are encouraged, if they have not already done so, to establish an internal work group to examine the impact of AI and GenAI on their courts and establish a plan moving forward.

Artificial Intelligence (AI) is an umbrella term and GenAI is a type of AI technology that is one of the most recognized by the public today. The term AI is used to refer to something as simple as spell check, predictive typing or asking Siri or Alexa the temperature, or as complex as computer based legal research, projections, facial recognition, or generating documents, videos, or audio.

AI Rapid Response Team Members

Chief Judge Anna Blackburne-Rigsby
Chief Judge, D.C. Court of Appeals/CCJ Chair
AI RRT Co-Chair

Justin Forkner
Chief Administrative Officer
Indiana Supreme Court/COSCA
AI RRT Co-Chair

Justice Beth Walker
Justice, Supreme Court of West Virginia

Chief Justice Matthew Fader
Chief Justice, Supreme Court of Maryland

Chief Justice Michael P. Boggs
Chief Justice, Supreme Court of Georgia

Judge Joseph A. Zayas
Chief Administrative Judge,
New York State Unified Court System/COSCA

Stacey Marz
Administrative Director,
Alaska Court System/JTC/COSCA

Sara Omundson
Administrative Director of the Courts,
Idaho Supreme Court/COSCA

NCSC Staff

Shay Cleary Cathy Zacharias David Sachar Miguel Trujillo Andrea Miller

Definitions

Artificial Intelligence (AI)

"A machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments."[1]

AI Models

- **Open AI Models**

  Open AI is a type of model that is publicly accessible which anyone can modify. Open AI models are designed to be more flexible and adaptable, capable of learning and evolving over time. They are trained using publicly available data from across the internet, such as text articles, images, and videos. Open AI models have the source code openly shared so that people are encouraged to voluntarily improve its design and function. However, Open AI models can pose a severe data security risk. Data input into an Open AI model can be accessible by anyone seeking that data. Confidential, personal, and/or sensitive data should never be input into an Open AI model. As such, careful consideration and responsible usage are necessary to mitigate potential risks.

- **Closed AI Model**

  Closed AI is a type of model that is not publicly accessible although it may be trained using publicly available data. Closed AI models typically prioritize data security and confidentiality, maintaining strict control over internal data access and usage to safeguard sensitive information. This distinction ensures a higher level of data privacy, particularly concerning personally identifiable or confidential data.

Chatbot

A computer program that simulates a conversation to assist an end user with a task.

Generative Artificial Intelligence (GenAI)

Artificial intelligence that is capable of generating new content (such as images or text) in response to a submitted prompt (such as a query) by learning from a large reference database of examples.[2]

Hallucination

A misleading, inaccurate, or fictitious result produced in response to a GenAI prompt.

---

1 National Artificial Intelligence Act of 2020

2 "Generative AI." Merriam-Webster.com Dictionary, Merriam-Webster, https://www.merriam-webster.com/dictionary/generative%20AI. Accessed 16 Apr. 2024.

Definitions

Large Language Model (LLM)
Deep learning algorithms using natural language processing that can recognize, summarize, translate, predict, and generate content using very large datasets.

Machine Learning
A branch of computer science and AI that uses data and algorithms to enable AI to imitate the way that humans learn, gradually improving its accuracy.

Natural Language Processing (NLP)
A branch of computer science and AI that uses machine learning to enable computers to understand and communicate using human language.

Prompt
The user input that directs AI content generation.

Training
Refers to the process of teaching an AI model to properly interpret data and learn from it to perform a task with accuracy. This involves feeding the model massive amounts of data, examining the results, and tweaking the model output to increase accuracy and efficiency. The AI training process typically has three key stages: 1) Training: an AI model is given a set of training data and asked to make decisions based on that information; 2) Validation: in this phase, assumptions are validated to determine how well the AI will perform using a new set of data; and 3) Testing: the AI model is given an unstructured dataset.

Guidance for Using AI in Courts

Public Trust and Confidence

Public trust and confidence in the courts is integral to the credibility of the judicial branch. Courts and judicial officers are responsible to ensure that the use of GenAI and other AI tools does not erode the public's trust and confidence in courts due to errors or biases.

Court Rules: Courts should review their rules to determine whether they are sufficient to address expectations of lawyers and litigants concerning the responsible use of GenAI in court filings and proceedings or whether changes may be appropriate to clarify those expectations.

Ethical Guidelines: Education on the applicability of current ethical guidelines is vital to ensure that GenAI is used ethically by lawyers, litigants, and the courts. Courts should review their rules and comments to the rules to determine if they should be updated to clarify their applicability to new technological tools.

Education: Courts must ensure that judicial officer and court staff are educated on the benefits and risks of AI. Courts will need to be aware of how GenAI is used to create content that looks real, sometimes referred to as deepfakes, which will increasingly impact discovery and evidentiary issues in legal proceedings.

Understanding GenAI — What Courts Should Know

AI has the potential to streamline tasks within the courts, increasing efficiency and allowing staff to work on higher level tasks. AI also has the potential to be used to help create resources for self-represented litigants, expanding access to justice. But like any technology, AI is not infallible or without risks.

Limitations

With the proliferation of new GenAI tools being developed for the courts, lawyers, and self-represented litigants and the ease of their use, courts need be aware of the capabilities and potential limitations of GenAI tools such as ChatGPT, Gemini, and CoPilot (popular GenAI tools at the publication of this document).

GenAI is not a traditional search engine and most GenAI platforms are not designed to provide legal authority. The purpose of GenAI is to create content. Lawyers and self-represented litigants are already using GenAI in drafting legal documents and performing legal research, and courts must understand the capabilities of the tools they are using. This includes the benefits of time saving legal research, drafting assistance, and organizing large volumes of information. There are also significant concerns about lack of accuracy, bias, GenAI-enhanced evidence, and deepfakes. As discussed below, judicial officers should be aware of certain indicators that a document filed with a court was generated with GenAI.

Accuracy

Early GenAI tools have been known to create hallucinations, which means generating inaccurate or fictitious content, such as case citations to cases that do not exist. Multiple courts have now issued sanctions for lawyers submitting filings with fictitious citations generated by GenAI tools.

Attorneys and self-represented litigants are using these tools to create legal documents. Westlaw and Lexis now provide the capability of using GenAI for legal research. However, a recent Stanford paper revealed inaccuracies in the output generated by these legal research tools, despite the fact that they use closed training systems.[3] Courts should be aware of these issues with accuracy when reviewing legal documents.

The following are indications that GenAI may have been used to create a document:

- References to cases that do not sound familiar, cannot be found through traditional legal research, or have unfamiliar citation formats.
- At first read, AI text may sound impressive and well written, but there are often structural issues. AI content tends to be overly formulaic and lacks natural transitions between topics. Once you strike out all the words that are meaningless filler, there may not be a lot of substance left. AI is also not mindful of grammar rules or basic punctuation although that is improving.
- AI is designed to recognize patterns and replicate them as accurately as possible so look for repetitive patterns in the writing. Perhaps the most obvious sign of AI-generated content is the use of repeated words, phrases, or the same sentence structure used regularly in different paragraphs within the same document.
- Often AI generated content is written in the general sense, glossing over facts and figures and may be lacking details, unnatural phrasing, lack of natural transitions between topics, or errors that a human is less likely to make. It often uses alliteration to articulate an appealing word arrangement.
- The absence of relevant very recent on-point case citations may indicate the use of AI generated content. OpenAI models are trained on massive data sets that are not continually updated so if recent relevant cases are not cited, it may be due to the AI being trained on an earlier dataset.
- Humans use idioms and slang frequently. AI often uses these phrases and words incorrectly. If you spot an idiom that feels a bit off and seems forced into the text it is likely a sign it was created with GenAI.

---

[3] *Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools*, preprint study, https://dho.stanford.edu/wp-content/uploads/Legal_RAG_Hallucinations.pdf

Bias

Courts need to be aware of potential bias in the content produced by GenAI. It is important to understand the datasets used in training the model because if they are not diverse or contain incorrect data, the results could be biased or inaccurate. Examples include the initial version of Google's Gemini chatbot that created images of people who did not match the historical ethnic backgrounds, such as creating images of people of color wearing Nazi uniforms[4] and AirCanada's chatbot that gave wrong information about the policy on bereavement travel.[5]

Confidentiality

Any information entered into an open GenAI platform, including through a basic prompt, could become visible to the company operating the platform and other users. Court personnel should be educated to not enter confidential, sensitive or privileged information in a chatbot or GenAI system that uses an open training model. Open systems use the information entered to train the database and will retain the information in the system unless the terms of use for the system explicitly specify that it does not retain the information. If using a chatbot, disable the chat history if possible. Judicial officers and law clerks must avoid inputting confidential or non-public information, including draft decisions and opinions, when using tools that use open models.

Ethics

Judges and court staff need to learn to use GenAI ethically and responsibly and be aware of applicable ethical obligations under the judicial canons and rules of professional responsibility. *See* Ethics section below.

Security

Courts should continue to follow best practices related to cybersecurity in connection with GenAI. When using GenAI that is authorized by the court, court personnel should use court issued equipment and email software so that appropriate security protocols are in place.

---

4 https://tech.co/news/list-ai-failures-mistakes-errors

5 Id.

Deepfakes and Other Evidentiary Issues

Judges are increasingly grappling with evidentiary issues, particularly authentication, related to digitally enhanced evidence as well as the emergence of deepfakes (convincing false pictures, videos, audio, and other digital information) generated by AI. AI advances make it easier and cheaper to create enhanced digital evidence and deepfakes.

Digitally Enhanced Evidence

Digitally enhanced evidence is audio, video, or image evidence that have been enhanced by AI software. The purpose is generally to improve the quality of the audio, videos, or images. This differs from past uses, such as zooming in on an image, speeding up or slowing down a video, or separating a voice from background noise, in that AI may fill in pixels or other data in an image with what the software "thinks" should be in the image, thus altering it from the original.

This technology was recently at the center of a criminal trial in Washington state where digitally enhanced video was not admitted into evidence. The court based its decision on the testimony of an expert witness who testified that "the AI tool(s) utilized ... added approximately sixteen times the number of pixels, compared to the number of pixels in the original images to enhance each video frame, utilizing an algorithm and enhancement method unknown to and unreviewed by any forensic video expert." The court found that the expert "demonstrated that the AI method created false image detail and that process is not acceptable to the forensic video community because it has the effect of changing the meaning of portions of the video."

Courts should consider whether the rules of evidence are sufficient to address issues presented by the emergence of digitally enhanced evidence. In the meantime, in individual cases in which evidentiary questions related to the actual or possible use of GenAI are presented, judges may need to consider requiring expert testimony to determine the authenticity and reliability of audio, videos, and images that are challenged.

What are Deepfakes?

"Deepfake" refers to fabricated or altered but realistic audio, videos, or images made using AI software, for example, by embedding another person's likeness into an image or video. Deepfakes have become very sophisticated in recent years, and it may not be easy for an average person to identify the audio, video, or image as fake.

Deepfakes and the Courts

The issue of deepfakes can arise in any court proceeding in which a party presents evidence in the form of an image, video, or audio. Fabricated evidence could be submitted as authentic evidence or authentic evidence could be challenged as fabricated evidence. When a party alleges that digital evidence has been fabricated, expert testimony may be needed to authenticate the challenged evidence. This could result in a battle between the experts that could increase litigation costs for all parties and consequently could widen the access to justice gap.[6]

Deepfakes present a special concern because of the considerable impact that visual evidence has on fact finders. According to studies referenced in a recent law journal article, as compared to jurors who hear just oral testimony, "jurors who hear oral testimony along with video testimony are 650% more likely to retain the information."[7] Once jurors have seen video evidence, it is very hard for the impact to be undone, even with admonishments from the court. Another study, published in 2021 by the Center for Humans and Machines at the Max Planck Institute for Human Development and the University of Amsterdam School of Economics, demonstrates the difficulty of identifying deepfakes. The study found that the participants could not reliably detect deepfakes, that people are biased towards identifying deepfakes as authentic (not vice versa), and that people overestimate their own abilities to detect deepfakes even after being instructed on how to detect deepfakes.[8] Moreover, as the general population becomes more aware of the existence of deepfakes and the difficulty of detecting them, it is possible that jurors will become increasingly skeptical of all digital evidence that is challenged.[9]

Current Evidentiary Rules

The existing Federal Rules of Evidence and the various state rules of evidence require that any evidence submitted must be real and that the party submitting the evidence has the obligation to authenticate it. Judicial officers have an obligation to determine whether the probative value of the evidence submitted outweighs the possible unfair prejudice, confusion of the issues, or misleading of the jury that would result from its admission.

---

6 Delfino, Rebecca, Pay-to-play: Access to Justice in the Era of AI and Deepfakes (February 10, 2024). Loyola Law School, Los Angeles Legal Studies Research Paper No. 2024-08.

7 Rebecca A. Delfino, Deepfakes on Trial: A Call To Expand the Trial Judge's Gatekeeping Role To Protect Legal Proceedings from Technological Fakery, 74 HASTINGS L.J. 293 (2023).

8 Köbis NC, Doležalová B, Soraperra I. Fooled twice: People cannot detect deepfakes but think they can. iScience. 2021 Oct 29;24(11):103364. doi: 10.1016/j.isci.2021.103364. PMID: 34820608; PMCID: PMC8602050.

9 Rebecca A. Delfino, Deepfakes on Trial: A Call To Expand the Trial Judge's Gatekeeping Role To Protect Legal Proceedings from Technological Fakery, 74 HASTINGS L.J. 293 (2023).

Are the Current Rules Sufficient?

Prior to the advent of deepfakes, the rules of evidence have been sufficient to adapt to technology changes. Courts may eventually conclude that laws and rules of evidence addressing deepfakes lag behind the technology. At present, tools to detect deepfakes are not as sophisticated as the tools to create deepfakes. To mitigate the disruptive effect of deepfakes on litigation and jurors, judicial officers should identify AI-related evidentiary issues and rule on those prior to trial and outside the presence of the jury, if possible.

The legal community is having ongoing discussions about the need for changes to the rules of evidence. It will be important for the courts to address the potential for harm to the legal process that deepfakes pose, and to evaluate whether more stringent rules should be adopted for the admission of audio, video, and image evidence. In addition, for case types with high rates of self-representation, relying on the parties to challenge the authenticity of evidence, which the current adversarial process requires, may be unrealistic. If deepfakes proliferate, courts may need to reconsider who is responsible for determining whether evidence is authentic, especially if reliable technology tools become available that would enable courts to identify deepfakes.

Ethics

Competence in Technology is an Ethical Requirement

Judicial officers and lawyers have a basic duty to be competent in technology relevant to their profession.

Model Code of Judicial Conduct (MCJC) 2.5 imposes a duty of competence on judicial officers in performing judicial and administrative duties, which could include an obligation to keep current with technology and to know the benefits and risks associated with all types of technology relevant to service as a judicial officer. Model Rules of Professional Conduct (MRPC) 1.1 states that lawyers must provide competent representation to their clients which includes technical competence.

Judicial officers and lawyers should:

- Have a basic understanding of AI, including GenAI, and its capabilities. This includes knowledge of the terms of use and how data will be used by the AI tool, as well as general familiarity with machine learning algorithms, natural language processing, and other AI techniques relevant to legal tasks.
- Analyze the risks associated with using AI for research and drafting, such as bias or hallucinations.
- Determine which areas of practice or processes can be improved with AI.
- Determine where AI may not be appropriate for use in the legal profession or the judicial system.
- Learn how to optimize prompts to get better results when using GenAI models.
- Identify which issues may require new policies or rules for AI use in the court system.

Ethical Standards for Consideration

Judicial officers should be aware of the potential ethical issues arising from AI usage and keep the following rules in mind when using or considering AI.

Ex Parte Communication (MCJC 2.9)

The Rule prohibiting ex parte communication also prohibits considering "other communications made to the judge outside the presence of the parties or their lawyers" (MCJC 2.9[A]). Under certain circumstances, material generated by GenAI could arguably be viewed as outside information that is improperly introduced into the judicial decision-making process. Relying on such information could also result in a violation of the Rule's provision barring independent investigation (MCJC 2.9[C]). External influences on judicial conduct (MCJC 2.4) could also be an issue when a judge relies on an AI program that sets forth an opinion on legal policy.

Confidentiality

Judicial officers have a duty of confidentiality, and they must be cognizant of whether they — or their clerks or staff — are entering confidential or sensitive information, such as information included in a draft opinion, into an open AI system, and how that information is being retained and used by the AI technology. In an open system, it is possible that the AI tool will use the shared information to train the model, potentially breaching confidentiality. This is also true for lawyers per MRPC 1.6.

Impartiality and Fairness (MCJC 2.2) and *Bias, Prejudice, and Harassment* (MCJC 2.3)

The Rule requiring judicial officers to perform their duties fairly and impartially could be violated if a judicial officer is influenced by an AI tool that produces results infected by bias or prejudice.

Judicial officers need to be aware of the potential bias or prejudice inherent in certain AI technologies and that using it could violate the Rule against acting with bias or prejudice if the AI tool has biased data in its algorithm or training data.

Hiring and Administrative Appointments (MCJC 2.13)

Judicial officers should be aware of the risks of bias or discrimination if AI tools are used to help screen prospective clerks or other staff or to otherwise assist in the hiring process. If the algorithmic recruiting program is biased, it could produce results or recommendations based on discriminatory information, which could violate the rule requiring judges to make appointments impartially and on the basis of merit, as well as Title VII.

Duty to Supervise (MCJC 2.12)

Judicial officers have a duty to supervise staff and to make sure they are aware of the obligation to use AI technologies appropriately.

Candor towards the Tribunal (MRPC 3.3)

Attorneys have an obligation of candor to the tribunal.

Understanding AI's capabilities and risks, especially regarding bias and confidentiality, is a necessity for technological competence. Court professionals must stay up to date on developments in AI and the potential ethical implications of using it.

Developing an Internal AI Use Policy

Court leaders should establish policy that enables their organization to experiment and benefit from AI technologies while at the same time minimizes risk.

Establish an AI Governance Working Group

Establish a working group to oversee the acceptable use, development, and management of AI technologies and policies, consistent with the court's mission and values. The group should consist of representatives from all relevant stakeholders, including court leadership, business process, legal, and technology.

Assess the Court's Needs

Assess current processes, identify the court's goals and needs and determine whether AI technology furthers them. When drafting an AI use policy, be sure to think broadly about a wide range of use cases. Identify the use cases that could benefit from AI tools, such as automating repetitive functions, data analysis, summarizing, drafting, and other tasks.

Assess the Risks

Assess the risks associated with implementing an AI tool, in areas such as hallucinations, data security, bias, copyright infringement, and staff concerns about job replacement. When drafting an AI use policy, think broadly about potential and perceived risks and address ways to mitigate them. Ensure that any new technology complies with existing technology or security policies and technology infrastructure standards.

Considerations in Developing a Policy

When developing an AI use policy, consider including:

- the policy's purpose and scope: to whom it applies, to what technologies it applies, how it can be used, such as requiring the use of secure and encrypted networks when accessing or transmitting data through AI tools, and requirements about the use of court data for training AI tools;
- acceptable uses of AI that are responsible and ethical and comply with all applicable laws, regulations, and policies (*See Kentucky's and Utah's policies*);
- prohibited uses of AI that would jeopardize the court's network or potentially disclose confidential information;
- staff should not access, collect, use, or disclose personal or sensitive information beyond what is necessary for authorized business purposes;
- what data protection laws, regulations, or policies apply to the use of personally identifiable information and the data privacy and security measures that should be implemented or that employees should follow to protect the court's data;
- how to ensure that AI-generated content is not biased and does not reflect discrimination based upon race, ethnicity, gender, age, or other protected classes;
- when to update and patch AI tools to protect against vulnerabilities and security risks, if not already covered in another security policy;
- mechanisms to monitor whether the policy is being followed, and plans for what to do if the policy is violated (security and HR).

Implement, Review, and Update the Policy

After adoption, communicate the policy to staff, educating them on how to responsibly and ethically use AI tools. Schedule regular reviews of the policy and update it as necessary.

AI Platform Use and Procurement Considerations

As courts experiment with and use various AI tools, it is important that leaders understand how information and data may be utilized by the AI technology.

Understand the Technology and Contract Terms and Develop Procurement Requirements

Before implementing or purchasing any AI technology, understand what generative AI and other AI technologies are, how the technology will be used, and the vendor's terms of use, and then develop applicable procurement requirements.

Data Governance Plus Applies

AI tools are similar to other technologies in that it is critical to understand the sensitivity of data that will be entered, who will have access to it, and what will happen with it. The same considerations apply for any new data generated by the AI.

New Terms but Basic Contracting Principles Still Apply

As with any technology it is important to carefully review and understand all contractual terms and conditions. Be sure to also review terms and conditions buried in click-through agreements.

Key considerations also include whether any prompts or generated content will be available to other users of the product, the technology provider, or any third parties; how such data is stored; and if any such data will be utilized to train and fine-tune the model.

What is acceptable or not will depend on the sensitivity of the data for the specific task as well as how the AI technology was developed (for example AI built for specific legal use).

Take a Team-Based Approach

Given the novelty, complexity, and rapid pace of innovation, it is recommended to take a team-based approach that includes representatives from IT, Legal Counsel, The Bench, Business Operations, and those with diverse ethical perspectives. to evaluate AI technologies from all perspectives and understand how they will be used. It should be clear who in the organization has authority to agree to any terms and conditions.

Also Be on the Lookout For

Terms of embedded/required services

Some Generative AI technologies, such as Google Gemini, require that users have a Gmail account to access the technology, requiring evaluation of additional terms and conditions. Those terms may not be consistent with the court's security policies.

AI-related changes to terms and conditions

As existing technologies commonly utilized by courts (e.g., Zoom or Adobe) incorporate generative AI into their products, they may modify terms and conditions. These terms and conditions should be continually reevaluated, and any long-term costs of a free trial or preview should be understood.

Marketing Hype or Embedded AI

The label "AI-enabled" may be used loosely in marketing to sell a product or conversely AI may be buried in a product and not disclosed. Rigorous evaluation is required to discern genuine AI capabilities.

AI and How to Get Started

Courts that want to start using AI tools can limit the risks involved by considering the following approach.

Decide Whether to Use Open or Closed AI Models

Readily available GenAI tools like ChatGPT, Gemini, and Co-Pilot use open training models, often utilizing the entire internet's content for its training. Benefits of open models include that they are free, accessible, and easy to use. Downsides of such models include potential bias in the training data and that the information included in prompts may be used to train the models so users risk sharing confidential or nonpublic information and data if it is included in a prompt. Closed AI models are those created using specified datasets, so they are typically more secure and do not share prompts or results beyond the intended system. Courts should determine their comfort level in using AI tools that use open versus closed training models, considering intended use of the tool, type of information and data that may be shared, and available financial and personnel resources to develop, manage, and support a closed AI tool.

Ensure Permission and Understand the Terms of Use

Before using any generative AI technology, ensure that the organization and policy makers are comfortable with the tasks it will be used for and can accept any terms and conditions that are attached to the use of the technology (e.g. data being sent back to the model). If one does not already exist, consider creating an internal policy that provides guidelines for the use of AI technology.

Select a Few Simple "Low Risk" Tasks

Select tasks to be performed with the assistance of GenAI tools that exclusively utilize public data or nonconfidential information and are easily verified for accuracy. Internal facing examples include drafting internal communications and policies, drafting performance evaluations and improvement plans (not including identifying information), drafting training plans for different positions, and conducting basic research. Outward facing examples include summarizing published Supreme Court opinions, drafting press releases about upcoming public events, or drafting committee meeting agendas and minutes. Get comfortable with using the different GenAI tools by starting with internal facing tasks and documents before using AI tools on external facing items.

Use a "Human-in-the-Loop" Approach

GenAI technologies and the use of them in courts are new, and therefore AI-generated output should not be relied upon until it has been reviewed by a human subject matter expert an approach called "Human-in-the-Loop". Presume the output will contain errors and likely bias. Carefully review AI-generated documents and output for accuracy, bias, and completeness. Once more comfortable with the technology (and depending on the task), and its reliability in terms of desired results, accuracy and bias, reevaluate to determine whether the documents and output can be periodically spot-checked by a human to ensure accuracy, instead of checking every document.

Note that the approach may vary with a closed model AI tool. from a reputable vendor having a model that was developed/trained for a specific purpose versus free or low-cost public tools.

Train Staff and Judges on AI Systems

To effectively utilize generative AI technologies, provide training and education to staff and judges on those technologies approved for court use. This helps them understand how to navigate the AI tool, interpret and successfully generate outputs, and effectively review and validate the AI-generated documents or results.

Prepare for Advanced Tasks

As court personnel become more comfortable with utilizing GenAI for basic tasks, consider how it can be used for more advanced tasks, such as data extraction and entry, external facing chatbots for customer service using court self-help and website content, or automated drafting of orders. Conduct pilot projects to test the feasibility and effectiveness of the technology in each specific context. This allows for a controlled testing environment where the technology's impact, benefits, and risks can be assessed.

Engage in Knowledge Sharing

Share what is learned with other courts that are also experimenting with GenAI. This allows for the exchange of experiences, best practices, and lessons learned, enabling courts to make informed decisions and avoid potential pitfalls.

Possible Uses of AI in the Courts

Potentially Useful Tasks

- AI tools are capable of summarizing large amounts of text. As with any summary, care needs to be taken to ensure the summary is accurate.
- AI tools can organize a large amount of information as directed.
- AI tools can find specific information in a large volume of data.
- AI tools can do an acceptable job of creating a first draft of something – a contract, a speech or remarks on a specified topic, job interview questions, position descriptions, performance evaluations, or policy provisions. However, it is essential to review, check, and refine the output and not treat it as a final product. Be aware that different prompts, even with only slightly different wording, will produce different results, so try several prompts to get closer to your desired result.
- AI tools can be used in writing presentations, e.g., to provide suggestions for topics to cover.
- Administrative tasks like composing emails and memoranda can be performed by AI.
- Generating images for presentations. Images often contain hallucinations or accuracies so make sure to closely review to make sure there aren't oddities included. Multiple prompts may be needed to get the desired outcome.

COSCA
Conference of State Court Administrators