

GUIDANCE ON RESPONSIBLE AI USE IN LITIGATION

U.S. Magistrate Judge Maritza Dominguez Braswell

October 2026

INTRODUCTION

In October 2025 the Court issued guidance on artificial intelligence. A lot has changed since then. Today, far more lawyers and self-represented parties use AI on a regular basis. The tools have also become more capable, which makes them more useful, but also harder to supervise.

This document offers general information about AI, a set of practices for responsible use, and a framework for addressing hallucinations. It is not comprehensive. Anyone who uses AI should do their own due diligence, understand the strengths and limitations of their chosen tool, and exercise good judgment.

Additionally, this document is offered only as *guidance*. It is not a standing order that limits AI use or requires the parties to certify that they did or did not use AI in their written submissions.

In this Court's estimation, a prohibition on AI constrains parties in a way that is impractical given the current state of AI. Moreover, a broad disclosure requirement does not meaningfully reduce risk. To the contrary, it can create a false sense of security, add burden that falls hardest on solo practitioners, small firms, and self-represented parties, and discourage responsible use that could help narrow the access-to-justice gap.¹

In this Court's view, the answer is education, not a blanket restriction. To that end, the Court offers the following guidance based on its own experience and work in this space.²

Because this is not an order, it does not create new rules or requirements, and nothing in it modifies or supersedes any law, rule, professional obligation, or order issued by this Court or another judge.

WHAT THESE SYSTEMS ARE (AND WHAT THEY ARE NOT)

As used here, "AI" refers to generative AI (GenAI): tools that produce text, images, audio, or video in response to a user's instructions, often given in everyday plain language. ChatGPT, Claude, and Gemini are examples of GenAI tools developed by frontier AI labs. It is important to know that AI is often built into document review software, word processors, phones, and more. So even if you never open a tool like ChatGPT or Claude, you may still be using AI.

Probabilistic vs. Deterministic Systems. GenAI systems do not work like a calculator or a traditional research database. A calculator follows fixed rules and gives the same correct answer every time. A database retrieves a document that actually exists. A GenAI system, on the other hand, was trained on enormous amounts of data and learned patterns from it. When you ask it something, it predicts a likely

¹ See Braswell, [Between hype and fear: Why I have not issued a standing order on AI | Thomson Reuters Institute](#) (2026).

² See [MDB Website](#) for a collection of writings, podcasts, and information about how this Court uses and thinks about AI.

response, piece by piece, based on those patterns. Because it is *probabilistic* rather than *deterministic*, the same question can produce a different answer each time you ask. The system generates plausible language rather than looking up verified facts. That means no answer is guaranteed to be accurate, no matter how confident the AI tool sounds.

Legal Research Tools. Some tools are designed to reduce the risk of inaccuracy. The AI features in legal research platforms like Westlaw and Lexis, for example, generate answers grounded in their own databases of legal sources and link to the documents they rely on. This reduces how often the tools invent authority, but it does not eliminate errors altogether. Researchers who tested leading legal AI research tools found they still produced incorrect answers sometimes.³

Agentic AI. Agents are AI tools that can take action on a user's behalf, such as browsing websites, filling out forms, editing files, or purchasing items on the web. They do this with less step-by-step direction, and they can access your systems, files, and networks if you give them permission. The more a tool does on its own, the more oversight it needs.

AI USE IN LITIGATION

The practices below are not legal advice, nor are they directives. Still, you might consider applying some or all of them in light of your existing litigation and/or professional obligations. Moreover, because AI is increasingly integrated into many everyday tools and can enter a case in several different ways, some of these principles and practices may be relevant to your work even if you do not use AI tools yourself.

- 1. Remember the rules.** One of the reasons the Court has not issued a standing order regarding AI is because the rules already require verification and candor to the Court. For example, Rule 11 applies to every filing, whether an attorney or a self-represented party signs it. By presenting a filing, the signer makes certain certifications concerning the supportability of legal and factual contentions. Rule 26(g) similarly requires verification, completeness, and more in the context of discovery. Attorneys also remain bound by the Colorado Rules of Professional Conduct and can be held accountable for a violation.⁴ In other words, the existing rules already allow the Court to hold someone accountable for irresponsible AI use.

Additional note for self-represented litigants: If you use AI to help you prepare legal arguments, you need to understand those arguments and be able to support them when the Court asks you about them. Rule 11 applies to you. If an AI tool made up a case, quote, or fact and you filed it, you are responsible for it. If what you filed is false, misleading, or frivolous, the Court may impose sanctions, including monetary sanctions or dismissal of your claims or defenses.

³Magesh, Surani, Dahl, Suzgun, Manning & Ho, *Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools*, *Journal of Empirical Legal Studies* (2025)

⁴ [Artificial Intelligence and Professional Conduct | Colorado Lawyer](#); [The 2026 AI-Related Amendments to the Colorado Rules of Professional Conduct | Colorado Lawyer](#)

2. **Stay in control.** AI works best as support for your judgment, not a replacement. A helpful rule of thumb is to sandwich AI in the middle of your process. Before you open an AI tool, read the materials, do your own thinking, and form your own views. Starting with AI can anchor a user to the tool’s framing, invite overreliance, and make it hard to separate your reasoning from the tool’s. After you’ve had an opportunity to choose your direction, AI can assist you in the middle of your process. It can find gaps, identify inconsistencies, suggest counterarguments, and more. However, before finalizing, verify. In other words, your own thinking comes first, verification comes last, and AI stays in the *middle*. This way, you stay in control and your submissions to the Court reflect your judgment.
3. **Verify outside of the AI tool.** Asking an AI tool whether its own answer is correct is not verification. Verification means confirming the case is real using a reliable legal source other than the AI tool.⁵
4. **Match the level of care to the level of risk.** Some uses are low risk, such as formatting or improving the clarity of your own notes for your own review. Others are high risk, such as using AI to prepare something you file with the Court. The higher the stakes, the more checking the work requires.
5. **Set the right guardrails.** AI tools can be constrained to some extent. Many offer features, under names that vary by product, that let you set “standing instructions” that apply to every session, give the tool reference documents to draw on, or create projects with custom instructions. It is worth learning which of these features are available and fit your work before relying on a tool. Often the tool can explain these features to you in plain language and guide the set up. They do not eliminate risk altogether, and a tool may not always follow your instructions precisely, but leveraging these features can help you set guardrails.
6. **Reduce AI noise.** AI makes it easy to produce long, repetitive text and file more. More words and more filings do not make a case stronger. They may slow the case down and pull the Court’s attention away from the issues that matter. The Court appreciates focused and concise filings.

Additional note for self-represented litigants: AI can help your case, but it can also hurt your case. If you decide to use AI, use it to help you say what happened in plain language. Do not try to sound like a lawyer or let it add facts, change what you mean, say things you don’t understand, or make your case sound stronger than it is. Additionally, if you make your legal arguments unnecessarily long, you risk burying the key points.

⁵ If you do not have access to tools like Westlaw and Lexis, there free resources, including: the Cornell Legal Information Institute (<https://www.law.cornell.edu/>), Justia (<https://www.justia.com/>), CourtListener (<https://www.courtlistener.com/>), and Google Scholar (<https://scholar.google.com/>).

- 7. Know what your AI tool does with what you give it.** Depending on the tool, its settings, and its terms of use, it may store what you enter or use it to train future versions of the tool. Free and consumer-grade versions often handle information differently than enterprise versions, and default settings can change over time. It helps to check them regularly.

If you would like to protect confidential documents you produce to the other side, it's best to confer early and try to agree on a protective order. You know your case and how you use AI far better than the Court does, so terms you work out together will usually fit your case better than terms the Court writes for you. The Court has addressed this issue in *Morgan v. V2X, Inc.*, No. 25-cv-01991-SKC-MDB, 2026 WL 864223 (D. Colo. Mar. 30, 2026) and *Dunn v. LexisNexis Risk Sols., Inc.*, No. 26-cv-01249-GPG-MDB, 2026 WL 2568346 (D. Colo. Aug. 31, 2026).

- 8. Remember your AI use is only part of the picture.** Associates and partners working on your matter may be using AI. Clients might use it to search, sort, or preserve their own information, and experts might use AI to assist them as they prepare reports. Talking with them early helps you understand how AI is being used and gives you an opportunity to address any problems on the front end. It also better prepares you to answer the Court's questions when they come up.
- 9. Consider AI when planning discovery.** AI may now affect how information is preserved, collected, reviewed, and produced. It may be most efficient to address AI early, at the Rule 26(f) conference, rather than after a dispute arises. For example, parties may wish to discuss whether they expect to use AI for document review or production, and if so, whether that needs to be accounted for in an ESI protocol. Parties may also wish to discuss upfront whether any AI-generated material might be subject to discovery, and if so, how the information will be preserved. Evidence from AI-enabled devices and systems may have several layers of data (e.g., raw data, data points inferred from the raw data, and conclusions about the user based on those data points). The Court prefers discovery plans that reflect robust conferral and the realities of the particular case, rather than something generic.
- 10. Raise legitimate concerns early.** AI makes it easier to fabricate convincing evidence and attempt to manipulate the process in other ways.⁶ If you have a reasonable basis to believe that the opposing party is using AI to improperly manipulate documents or information, raising it early through conferral and motion practice, if necessary, gives everyone time to investigate and resolve the issue before trial. However, keep in mind that making such a claim without a reasonable basis could leave you responsible for the expense of addressing the issues.
- 11. Be candid about your AI use.** Using AI is nothing to hide, and the Court welcomes the responsible use of these tools. The Court also appreciates when parties affirmatively disclose AI use that could impact how the other side or the Court evaluates something, such as an enhanced

⁶ See, e.g., *Elliott v. New York Bariatric Group, LLC*, No. AAN-CV-25-6066141-S (Milford Super. Ct., Ct. Aug. 6, 2026) (where a party embedded white text that was invisible to the human eye but expressly directed AI systems to generate favorable outputs).

exhibit or a demonstrative that is AI-generated. Sometimes the Court will ask a lawyer or party about their AI use. It may be for general learning purposes, or because the AI use may have an impact on the proceedings. If you have a legitimate concern about answering, such as protecting your mental impressions or litigation strategy, it is better to raise that concern directly than to give a partial answer.

- 12. Keep learning.** AI literacy does not require a sophisticated technical background. However, it does mean understanding how these tools generate information, where they tend to go wrong, and how to check their work. Additionally, the technology changes quickly, and what you learned a few months ago may be out of date. This means AI literacy is an ongoing effort.

HALLUCINATIONS

Just as with the information above, the information below is guidance. It is not an order. Every lawyer and self-represented litigant must use their own judgment in determining how to address errors in a document or court submission, regardless of how those errors came to be.

If you spot a hallucination in your own filing, it's best to promptly correct it and notify the Court and the other side. Prompt self-correction goes a long way with the Court, and prompt notice of the issue will help contain the hallucination.

If you spot a hallucination in an opponent's filing, the Court suggests the following:

- 1. Confer.** The purpose of conferral is to (a) confirm it is a hallucination; (b) discuss how the party responsible for it (the "Drafter") can correct or withdraw it without further harm; and (c) explore whether other remedial steps are appropriate, and if so, try to agree on them.
- 2. Jointly inform the Court.** Informing the Court will help avoid the perpetuation of inaccuracies.
- 3. Escalate, when necessary.** If the Drafter is unresponsive or will not engage in good faith, the opposing party may have no choice but to file a notice alerting the Court to the issue.

When conferral does not resolve the matter, or the Court otherwise becomes involved, the Court will consider the specific facts and circumstances, which may include:

- whether the Drafter engaged promptly and in good faith once alerted;
- whether the Drafter took responsibility and acted quickly to limit any harm;
- whether the Drafter took reasonable steps to verify accuracy or address other issues before filing;
- whether the Drafter follows written AI policies that reflect responsible use;
- whether the Drafter is self-represented, which may impact the Court's assessment of "reasonable steps";
- the nature and extent of the error or issue, which may show whether the AI use was reckless or malicious;
- whether the issues or errors caused meaningful prejudice or harm to anyone; and
- whether the issues or errors wasted significant judicial resources.

CLOSING

Beyond hallucinations, the legal profession will have to work through increasingly difficult questions: How do we detect and interrupt bias? How do we guard against fabricated evidence? How do we keep humans in control of critical decisions? How do we protect privacy as data collection increases? How do we equip the next generation of lawyers and judges to serve in an increasingly AI-integrated world? How do we ensure these tools expand, rather than narrow, access to justice?

We're all learning, and none of us has every answer, but we'll see greater success if we work through these issues together. As we do, we should offer each other a measure of grace. There is a real difference between malicious or reckless use of AI, on the one hand, and ordinary human error during a period of rapid change, on the other. An approach that pairs accountability with grace, rather than one built on fear and blame, will better serve the people we are meant to help.