AI as a Force Multiplier for Appellate Practice: A Step-by-Step Guide

June 9, 2026

PRACTICES Appellate, AI and Technology, Litigation

AI is reshaping legal practice — appellate work included. The lawyers who master AI tools will not be replaced by them; they will be hired over lawyers who do not. Careful AI use enhances both written and oral advocacy while creating opportunities for meaningful, impressive efficiency gains.

The technology was breathtaking when I ran my first AI query in August 2025 and has rapidly improved since. I now use it almost daily, while always rigorously checking the output. My firm uses Harvey AI, which can run queries across multiple leading models, including Claude, Gemini and GPT. But no matter which tool you use, AI fluency is increasingly important to modern, efficient practice.

Here are some practical lessons I have learned that are aimed at appellate practitioners but applicable to any litigation practice.

First Steps

Before you query anything, organize your materials. In Harvey AI, for example, you can set up a "vault" — an internal shared workspace — and invite your full team so everyone works from the same materials and can build on each other's queries. This setup pays dividends throughout the engagement.

I typically start by creating three top-level folders with subfolders:

- First, relevant authorities (cases, statutes and secondary sources), with subfolders organized by issue — if there are four key issues in the case, create four subfolders.
- Second, the filed briefs — yours, your opponent's and any amicus submissions.
- Third, the underlying record, divided among subfolders for what we in Texas call the reporter's record (trial and hearing transcripts), the clerk's record (pleadings, motions and orders) and key exhibits (like the contract in a breach case).

That structure matters. You can aim each query at the right subfolder so that the model reads what matters rather than rummaging through everything. As a result, outputs are sharper and more reliable. By the time you sit down to write, your organizational work is already done.

Record Review

AI's most obvious value for an appellate lawyer is helping you digest complex and lengthy trial records. Aim the model at targeted portions of the reporter's and clerk's records and ask for witness-specific or issue-specific summaries tied to the elements you care about. You can build first-pass chronologies and timelines much faster. The model extracts dates, events and responsible actors from transcripts and exhibits, then aligns the entries to your theories of

reversible error. It will not replace your record review and issue-spotting — but it gives you a major head start.

Then pressure test your factual narrative. When contemplating a legal-insufficiency (no-evidence) challenge, use the model to anticipate what evidence your adversary will marshal in response. Ask AI to identify the most troublesome adverse evidence and brainstorm how to neutralize it in your brief. The best appellate lawyers have always done this instinctively; AI simply accelerates the process.

Case Research

Treat the model as a research accelerator, not an oracle. Use it to think through the law: frame the issue, ask for competing rules, insist on the best arguments for each side and probe how standards of review shape the outcome. You will surface lines of attack and defense faster and see where your case sits on the doctrinal spectrum.

Then build a case table you can scan at a glance — columns for citation, court and year, procedural posture, material facts, holding, reasoning and why it matters to your brief. Have your AI tool draft the first pass at the table, then refine and verify each row. Ask it to flag weak support or holdings that do not match the proposition cited. With a little cleanup, you now have a working document ready for a case appendix, internal strategy session or oral argument review.

When you need to go deeper, combine this synthesis with your traditional tools. Run targeted follow-ups in your research platform to confirm and expand on your findings, then use your Vault to organize your key authorities and let subsequent prompts draw only from those curated sources.

Needless to say, always verify every case citation, quotation and holding that AI generates. Models can hallucinate authorities, misattribute quotes or subtly misstate holdings — errors that can undermine your credibility with the court (or worse).

Brief Writing

If I am working on an appellee's brief or an appellant's reply brief, I always use AI to analyze my opponent's brief before drafting a word of my own. Ask AI to do a surgical read: identify logical gaps, unsupported assertions and leaps from premise to conclusion. What comes back is an initial roadmap for your response. You still must verify each point, but you start with a structured list of likely gaps rather than reading with a highlighter and hoping something catches your eye.

Next, check your opponent's citations. Miscitation is common — not always intentional, but usually exploitable. A case cited for a broad rule that actually turned on a narrow procedural fact. A parenthetical that omits a critical qualifier. A string cite where the second and third authorities vary from the first. The model flags these problems quickly.

The same logic applies to the record. Ask AI to analyze whether the factual assertions are accurate. Does the brief overstate testimony? Does it omit the next sentence that provides crucial context? Does it mischaracterize an exhibit? This is not a shortcut around close reading — it makes close reading faster and more thorough.

For your own writing, use AI to brainstorm and outline. Dictate your themes and ask for a clean structure with issue framing, standards of review and placeholders for authority. Then, if you have an AI program embedded in Word (like Harvey for Word), shift to that and work inside the

document. Generate tracked-change edits for clarity and tone while preserving citations. Ask the model to explain every revision so you can accept the gains and reject the rest — nothing changes without your review and approval.

For lawyers developing their own voice, targeted prompts can help. Try this: “Review this argument section and suggest edits that would make the prose more direct and persuasive. Prioritize short, declarative topic sentences; active voice over passive; concrete language over abstraction; and varied sentence rhythm. Explain each suggested change so I can learn from the feedback.” The explanations turn editing into a learning exercise while also sharpening your work product.

I have previously written and spoken on the value of incorporating visuals in briefs. AI can help with that too. Ask your model to: “Identify any opportunities to enhance the brief’s advocacy through visual elements. Specifically, analyze the brief’s arguments, factual background, and key data points to assess whether adding images, graphics, tables, timelines, charts or other visuals would strengthen the persuasive impact.”

As you get close to the finish line, use AI to help you put the finishing touches on your brief by asking for stylistic edits that would “eliminate redundancy, tighten sentence structure, improve clarity and flow.” (My preferred prompt is a bit more elaborate, but you get the picture.)

Oral Argument Preparation

Once you begin preparing for argument, turn the AI model into a moot panel. Ask for the hardest questions on your weakest points. Answer them and ask for candid feedback. Run another round. Sharpen each response as the argument date approaches. This is similar to the iterative process that elite appellate teams use in moot courts — now available whenever you need it.

If you are the appellee, ask the AI model to simulate what the argument might look like for the appellant — presenting your opponent’s best version of the case and identifying where your answers need sharpening.

Then ask the model to act as a judge and draft an opinion on each issue based on the record and briefing in your vault. Where you lose, shore up. Where you win, tighten. Test whether your arguments move the needle. You will get far better results analyzing one issue at a time rather than the entire case at once.

Guardrails and Oversight

Two rules govern everything above. First, verify *every* authority and fact. Second, never outsource your professional judgment. The safest workflows keep you inside Word, where you see tracked changes and can click into sources you have provided. Tools like Harvey ground responses in the materials you select and expose them for inspection. That supports the habit appellate specialists already live by: rigorous verification of every record cite and case citation.

But AI verification requires a higher level of scrutiny than you would apply to a human colleague’s work. AI errors can be uniquely egregious and insidious. The model does not know what it does not know, and it will not flag its own uncertainty. I have seen AI cite “testimony” that was actually an inadmissible offer of proof or lawyer argument in a closing statement. I have seen it invent a nonexistent judge — a phantom “Judge Smith” who never sat on the case. I have seen it confidently misplace a geographic landmark, inexplicably relocating I-45 to Beaumont. These are not the kinds of mistakes a competent associate would make. Treat AI as an associate who does

brilliant work but is diabolically trying to make you look bad. You must verify not only that the analysis is sound, but that every underlying assumption and cited fact is correct. Rigorously check every record and case citation as if your reputation depended on it — because it does.

Ethics point the same direction. Competence now includes understanding the benefits and risks of relevant technology, and courts expect meaningful human oversight. Your AI model is a collaborator; you remain the lawyer of record.

Be aware that some courts and judges now require disclosure of AI use in filings (although, as of this writing, no federal circuit court does). The Northern District of Texas, for example, has adopted standing orders mandating that parties certify whether generative AI was used in preparing submissions. Similar requirements are emerging in other jurisdictions. Before filing, check the local rules and any judge-specific orders that may apply.

There is another reason to stay current: The judiciary is also studying and, in some contexts, experimenting with AI. Understanding how AI works — its capabilities and its failure modes — helps you anticipate how an AI-assisted judge might engage with your arguments. If your audience is using AI, you should understand it too.

Client Approval and Confidentiality

Should clients be informed that you are using AI on their matter? Whether client notice or consent is required depends on the tool, the information being used, client guidelines, the terms of your engagement, court rules, protective orders and applicable ethics guidance. Some practitioners take the position that AI use is now so widespread that clients should be presumed to understand it is part of modern legal practice — much like research databases or document review software. Others believe clients should be given an express opportunity to opt out. The right answer depends on the client, the matter and the sensitivity of the information involved.

My firm has adopted a formal generative AI policy that governs how we use these tools. Only firm-approved models may be used for client work. Users must complete mandatory training before using any approved tool. And the policy is clear: AI output must be verified by a lawyer with appropriate experience, and AI may not replace a lawyer's professional judgment or legal analysis. Some clients have opted out entirely, and we maintain a list of those matters. These guardrails reassure clients that AI is being used responsibly and under close supervision.

My firm's experience has been positive. We find clients are most comfortable when firms clearly explain the controls. In our firm-approved configuration, approved AI tools are governed by contractual, technical and policy controls designed to protect client data. For example, client materials are not used to train public models, access is controlled and matter materials are kept within the workspace or matter context selected by the user. We have encountered few objections. When clients do ask questions, they are usually satisfied once they understand how the tools work.

Conclusion

Technological competence is now a part of professional competence. Used thoughtfully, with verification, grounded sources and human oversight, AI tools let you spend more time on what matters: thorough, accurate and persuasive advocacy that respects the record, the law and the court's time.

This article was originally published by *The Texas Lawbook*.