ALM | LAW.COM

New York Law Journal

AI-Native Appeals: Promises and Pitfalls

By Robert Niles-Weed

August 17, 2026

When we look back on this era of appellate practice, our most consequential judges may turn out not to be the ones on the bench publishing opinions but the ones on our computers. Artificial intelligence is already capable of reading briefs, reviewing a record, ingesting case law and other authority, and producing well-reasoned, well-enough-written opinions. A plausible AI-generated opinion is no substitute for the outcome of actual human decision-making. But AI's ability to simulate appellate decision-making is transformative: a simulated panel can show where an appeal is stable and where sensitive, and can support better advocacy and analysis at every stage.

Some judges already seem to recognize as much. A random-sample survey published this spring in The Sedona Conference Journal found that more than sixty percent of responding federal judges had used at least one listed AI tool in their judicial work—1.8% reported using AI to *make* decisions, 4.5 to *inform* them.

Appeals and AI: A Natural Affinity

There are good reasons to think that appellate practice is particularly well suited for AI enhancement. The universe on appeal is largely closed: the record in the lower court is set, the briefs and the cases they cite are known, and the issues for decision are often narrow and well framed. There are no credibility assessments, no disputed facts to be found, and the vicissitudes of trial have been liquidated into a cold transcript. All that remains for the court is to synthesize these materials and write an opinion resolving the case. An appeal is text in, text out. This makes it distinctively well suited for a tool conceived in text whose greatest skill is predicting what output is most likely to result from a given set of inputs.

Courtesy photo

Robert Niles-Weed of Weil Gotshal & Manges.

Appellate standards of review only strengthen the fit. Deference is strongest where AI is weakest—as appellate courts give wide berth to trial courts when it comes to the fact-bound discretionary analyses calling for human judgment. And the pure legal or textual questions where AI is strongest receive de novo review.

Appeals are not solvable by AI, however. Appellate adjudication is not a completely closed world. Courts may consult authorities the parties overlooked, raise issues not fully ventilated in the briefing (within limits of party presentation), or announce a rule no party asked for. And humans

are unpredictable, and sometimes driven by considerations an AI would never expect or appropriately weigh.

Where AI Can Help

Briefing: Start with the brief. Appellate lawyers have always mooted oral argument with a panel of lawyers acting as judges. But rarely will they convene a moot panel to review a brief. An AI-simulated panel can read the record and a draft brief (and an adversary's brief, if it has been filed), state the apparent rule of decision, and identify arguments that are decisive, compelling, unconvincing, or unanswered. AI is no substitute for the human drafter, but it is a rich and tireless interlocutor in the editing process. If successive "panels" keep bypassing the merits to resolve the case on a threshold issue, that may be a valuable signal to reallocate real estate to that point. Or if they repeatedly find the case hinges on a point buried on page forty, the table of contents may be telling the wrong story. Or if AI keeps tripping on a particular turn of phrase, it may be time to kill that darling.

Oral Argument: The same discipline applies to oral argument. AI can generate difficult questions in volume and pressure-test proposed answers. To mitigate potential panel effects, counsel can retrieve the various judges' relevant opinions and argument transcripts and generate questions keyed to their actual treatment of the issue and argument style. This is particularly valuable on the fly, as some courts make panel announcements as late as the day of the argument, and AI can help guide last-minute preparation and refine forecasted outcomes based on a known panel. Of course, judge-specific simulation should not become judicial astrology. The corpus may be small; a judge's views may have evolved; and a three-judge panel is an institution, not the sum of three personas.

Valuation Exercises: AI simulation can also inform valuation and risk assessment. Litigants want to know what the expected value of their case is and the probabilities of various dispositions, whether for potential settlement, litigation reserves, or otherwise. While the future remains hard to predict, AI can help by forecasting outcomes, even by doing so repeatedly or by assigning probabilities to various potential dispositions. Repeated simulations can allow counsel to gauge the stability of the disposition and reasoning.

To be sure, running a model one hundred (or one thousand) times does not yield a perfect estimate of what a real court will do. Ninety affirmances out of a hundred may reflect a robust case—or a consistent blind spot. But repeated runs do yield a distribution of that model's answers that can support a quantitatively grounded sensitivity analysis. AI offers evidence to investigate and incorporate, not an ironclad prediction.

***Mediation or Arbitration*:** Appellate decision-making can be slow, and parties may prefer an immediate simulated outcome to a long-delayed and uncertainly timed real one. A stable simulated outcome, generated from a controlled model that neither side commissioned, can anchor mediation and settlement discussions in a way dueling advocacy cannot. Parties can even agree to be bound by the outcome, on appropriate terms. (The American Arbitration Association has launched an "AI Arbitrator" for a narrow class of cases, with human review built in—a promising first step.)

In each case, AI is not a substitute for the appellate lawyer's taste and judgment, but an enhancement.

Building a Serious and Safe AI Panel

A serious AI panel requires more than a simple prompt. For one thing, controlled inputs will increase the reliability of the outputs and avoid the risk of hallucination. Every factual assertion should link to the record; every legal proposition should be independently verified. An AI decision

based on non-existent authority is worse than no decision at all. Human review of AI output remains critical.

Bias is another risk. Blinding the AI panel to the side being represented to the extent possible will help avoid the known risk of being told what you want to hear. Groupthink is enough of a risk among humans—AI should mitigate it rather than exacerbate it. Counsel must also guard against model updates that silently bias results, points of view subtly embedded in uploaded materials, and the human tendency to trust a fluent answer.

Perhaps the most important guardrails for any use of AI in this context involve concerns around information governance and professional responsibility. AI-enhanced decision-making is hard to justify if the cost is busting privilege. Judge Rakoff's decision in *United States v. Heppner*—finding that materials generated by a criminal defendant using a consumer AI were not protected—has made waves. And for good reason: casual client use of AI tools without attorney supervision is risky. But controlled use within a closed enterprise-based system is safer, and there are further precautions to minimize risk as this area of the law evolves.

Privilege is only one concern. Counsel must separately examine confidentiality duties, retention and deletion obligations, and provider access—including whether client inputs will train anyone's model.

These risks are real, but controllable—reason for careful design, not for avoiding AI altogether.

The Human Panel Endures

None of this displaces human appellate judgment. We commit consequential disputes to human lawyers arguing to human courts because we want accountable and public judgment that will carry weight in future cases. Candor, practical wisdom, and the trust of the court cannot be automated. And AI will have difficulty weighing complex competing considerations and business objectives to reach the best approach for actual clients. Taste matters.

But rigorous and embedded use of AI is already enhancing appellate practice. The best appellate lawyers of the next generation will not be simply human or simply AI, but rather the strongest human lawyers collaborating with the best AI tools. And AI has the capability to raise the bar at every level.

The moot court was once a new technology; it has now become standard practice, and it improves the quality of appellate advocacy across the board. So, too, for electronic legal research. Now comes AI. An AI appellate panel takes both of these innovations and supercharges them. A moot can be convened on a moment's notice, over and over again. Briefing can be reviewed at a previously unheard-of level of detail and precision. Prediction and valuation can become more science and less art.

Counsel needs to understand that simulation is not prophecy and rhetorical fluency is not necessarily judgment. But careful use of thoughtfully designed AI panels can allow appellate lawyers to have a bench on their side long before they enter the courtroom.

Robert Niles-Weed, *a partner with Weil, Gotshal & Manges, is co-head of the firm's appeals and strategic counseling practice and a member of the complex commercial litigation practice. He focuses on various commercial disputes at all phases of trial court litigation, as well as appeals in state and federal courts, including the U.S. Supreme Court.*