



Janet Gusdorff

GUSDORFF LAW, P.C.

The cold-eye advantage

USING AI TO STRESS-TEST TRIAL STRATEGY, EVIDENCE, AND APPELLATE RISK

Trial lawyers spend months – sometimes years – living with their cases. By the time a matter reaches dispositive motions, mediation, or trial, the file is no longer just information. It becomes familiar terrain: which facts matter, which ones don't, what arguments land, and which feel like distractions. That familiarity is a strength. It allows a team to move quickly, speak confidently, and present a coherent story.

It is also a predictable source of blind spots.

The longer a team lives with a case, the easier it becomes to normalize its weak points. Facts that once felt sharp fade into the background. Evidentiary gaps become “explainable.” Early strategic decisions – about themes, witness order, or emphasis – quietly harden into assumptions. Over time, the team stops noticing where the story relies on shared context rather than proof.

This article is about using AI not to write better, but to *see better*. Properly constrained, AI can function as a cold-eye diagnostic – a way to surface risk early, while there is still time to adjust strategy, proof, and framing.

Why AI works as a cold-eye tool

Trial lawyers already have feedback mechanisms: associates who review drafts, co-counsel who debate strategy, consultants who weigh in on persuasion, and, when budgets allow, mock juries. Those tools matter. But they are shaped by human dynamics – hierarchy, time pressure, relationships, and a natural reluctance to reopen decisions that feel settled.

AI adds something different: detached skepticism at scale.

AI does not know how much time a team has invested in a theory. It does not care that an argument “usually works.” It will not soften criticism to preserve morale or avoid friction. When prompted to critique, it does so consistently, repeatedly, and without fatigue. And because it can examine the same facts through different

lenses – a skeptical juror, eager defense counsel, a trial judge, an appellate panel – it can surface vulnerabilities the team has stopped noticing.

Used well, AI functions as a structured skeptic. It pressures lawyers to articulate assumptions, close gaps, and confront how the case looks without the benefit of shared context. Used poorly – through leading prompts or broad, open-ended requests – it simply mirrors existing beliefs faster.

The difference is discipline. Narrow inputs. Clear tasks. Skepticism-oriented prompts. Often, a short factual summary plus a few key excerpts produces more insight than uploading everything and asking for “thoughts.” The goal is not completeness. It is identifying what stands out when the case is viewed without your team's internal logic.

AI, for all its diagnostic usefulness, also has a gift for sounding completely authoritative while being completely wrong. That is not a bug unique to AI – it describes most confident people in a deposition. But it does mean every output deserves the same skepticism you are

asking AI to apply to your case.

Prompt for skepticism, not comfort

The difference between AI as a diagnostic tool and AI as a faster echo chamber comes down almost entirely to how you ask the questions. Notice the pattern in the examples below: Effective prompts assign a specific lens (skeptical juror, defense counsel, appellate panel), narrow the task, and explicitly invite criticism. AI given a validation task will validate.

Confidentiality: Using AI without getting burned

Before any of this is useful, there is a threshold question: What can you safely share with an AI tool?

The answer depends on the platform, your firm's policies, and your state bar's ethics guidance. Many consumer AI tools have terms of service that allow your inputs to be used for model training, which creates obvious confidentiality concerns. Enterprise versions of those same tools, or legal-specific platforms built on them, are typically safer.

The good news is that the cold-eye



exercises described in this article can often be run on summaries rather than raw documents. A one-paragraph factual narrative stripped of client-identifying information can generate surprisingly useful diagnostic output. When in doubt, anonymize.

Check your state bar's formal ethics opinions on AI use. Several have issued guidance, and more are on the way. This is not a reason to avoid the tool. It is a reason to use it thoughtfully.

Pressure-testing the narrative

A strong theme does real work. It organizes facts, guides witness preparation, shapes motion strategy, and gives jurors a framework for decision-making. It becomes the connective tissue of the case.

That central role is also what makes themes fragile.

Themes are usually developed early and reinforced repeatedly through briefing, mediation, and internal discussions. Once a theme feels intuitive and emotionally satisfying to the trial team, it tends to stop being questioned. Over time, confidence in the theme can grow even if the evidentiary foundation has not meaningfully strengthened. Facts that support the theme are emphasized; facts that complicate it are explained internally and then set aside.

The theme accomplishes what the evidence cannot yet do.

AI is useful here because it forces a separation between narrative strength and evidentiary strength. By asking AI to restate the case using only admissible proof – without inferred intent or rhetorical framing – trial teams can see where the story depends on assumptions rather than record support. Themes built around motive or moral blame are especially vulnerable when key documents may be excluded, limited, or ambiguous.

AI's literalism is an asset. If the connection between conduct and harm is implied rather than shown, it flags that gap – much like a skeptical juror or judge would.

This process also exposes assumptions teams stop noticing: What jurors will find intuitive, how much context they will supply on their own, or whether moral

wrongdoing automatically translates into legal causation or damages. These assumptions may ultimately prove correct – but if they are not proven, they are points of fragility.

In practice, this rarely leads to abandoning a theme. More often, it leads to tightening. Language narrows. Overstatement falls away. Proof is reordered so factual foundations precede moral conclusions. Sometimes it reveals the need for a secondary framing that can carry the case if key evidence is limited (or excluded pre-trial). Those adjustments show up in court in ways that matter: Openings become cleaner, direct examinations more disciplined, and credibility stronger.

A theme that survives a cold-eye review is more likely to feel grounded, fair, and durable. The goal is not to make the theme smaller. It is to make it sturdier.

Sample prompts:

- *“Restate the following case narrative using only facts that would be admissible at trial, without any inferred intent or rhetorical framing. Flag every place where the story depends on an assumption rather than record support.”*
- *“You are a skeptical juror hearing this case theme for the first time. Identify every claim that invites confusion, resistance, or follow-up questions, and explain what additional proof you would need before accepting it.”*
- *“Assume our key document on motive is excluded. Restate our theme without it. How does the story hold up?”*

When bad facts become invisible

Every case has bad facts. What creates risk is not their existence, but how quickly trial teams become desensitized to them. Over time, bad facts are discussed, explained, and contextualized internally until they stop feeling sharp. They become “handled.” Once that happens, they quietly drop down the risk hierarchy – even if they remain among the first things a juror or judge will notice.

Trial teams tend to ask defensive questions: *Can we explain this? Can our expert address it? Can we neutralize it on redirect?* Those questions matter, but they miss something more fundamental: How the fact functions for someone encounter-

ing the case for the first time.

AI helps shift that lens. By evaluating which facts are most likely to dominate credibility, causation, or damages perceptions, it helps teams recalibrate their instincts. A short, unfiltered factual summary – bad facts included, without justification – often reveals mismatches between what the team worries about and what actually stands out.

AI is also effective at identifying compounding effects. A treatment gap alone may be manageable. Combined with a prior injury and delayed reporting, it may tell a story the team has underweighted. Humans evaluate facts discretely; AI is better at spotting patterns.

Tone matters as well. Explanations developed internally can drift into defensiveness. AI can flag when an explanation sounds minimizing or argumentative, even if it is accurate. That insight often leads to better direct examinations: addressing a fact affirmatively, briefly, and neutrally rather than defensively or at length.

Once the risk hierarchy is recalibrated, strategy follows. Some facts should be addressed earlier to control first impressions. Some require expert context. Others should be acknowledged and moved past. Just as importantly, some uncomfortable facts do not deserve disproportionate airtime.

The goal is perspective, not panic.

Sample prompts:

- *“Here is a one-paragraph summary of a plaintiff's case, including the bad facts. Do not read any prior context from me. Identify the three facts most likely to cause a reasonable juror to rule against this plaintiff, and rank them by impact.”*
- *“Review the following explanation of [bad fact]. Does this explanation sound defensive, minimizing, or argumentative to someone hearing the case for the first time? Suggest a more neutral framing.”*
- *“Taken together, do the following facts – [list them] – compound each other in a way that tells a damaging story beyond their individual weight? Explain how a defense lawyer would use them together.”*

Before you file that motion

Motion practice does more than seek relief. It frames the case for the court and

signals which issues deserve attention.

The cold-eye question asks: *What does this motion cause the court to focus on – and what does it push to the background?*

AI helps simulate that effect. By asking AI to describe the case as it understands it after reading a draft motion – or even just the headings – counsel can see which facts are elevated, which assumptions are exposed, and whether the motion reframes the case in unintended ways.

Asking AI to identify the shortest path to denial is particularly useful. Judges often resolve motions narrowly. If a motion requires accepting multiple premises or weighing credibility disputes, AI will flag where the court is likely to stop.

This analysis is especially valuable for motions in limine, where protective instincts can backfire by spotlighting evidence the defense had not emphasized.

Sample prompts:

- *“After reading this motion, describe the case as you understand it. Which facts are most prominent? Which assumptions does the motion ask the court to accept without resolving?”*
- *“What is the shortest path to denial of this motion? Identify the single most likely reason a judge would stop reading and rule against us, without reaching the merits.”*
- *“We are considering filing a motion in limine to exclude [evidence]. What is the risk that this motion draws more attention to that evidence than if we had not filed it?”*

Reality-checking the case (and the client)

Mediation is not just a negotiation with the defense. It is often the first serious confrontation between a client’s internal narrative and the risks that will actually drive resolution. For plaintiffs’ lawyers, that confrontation can be difficult – especially when expectations are misaligned with the record.

By the time a case reaches mediation, clients have usually anchored their expectations. Sometimes that anchor is a verdict they heard about. Sometimes it is a number discussed early in the case. Sometimes it is a belief about what

“should” happen based on fairness rather than proof.

That anchoring can cut both ways.

Some clients overvalue their cases.

They focus on liability strength and moral clarity while discounting credibility issues, causation gaps, or evidentiary limits that will matter to a neutral. Others undervalue their cases – often because they are risk-averse, financially strained, or simply exhausted – and assume the defense holds more leverage than it actually does. Both dynamics can undermine mediation.

Rather than positioning the lawyer as the sole voice of recalibration, AI can be used to model how a neutral decision-maker might assess the case when encountering it for the first time based on the current record. The goal is not to assign value. It is to identify pressure points – where credibility matters most, where causation may break down, where damages are well-supported, and where they are vulnerable.

For clients who overestimate value, this exercise reveals where confidence rests on assumptions rather than proof. It reframes the conversation away from disagreement and toward perspective: *This is what a neutral is likely to focus on.*

For clients who underestimate value, the same process highlights where the record is stronger than they believe – where liability is clearer, where damages are supported, and where feared defense arguments are unlikely to carry decisive weight. That clarity can help distinguish real risk from emotional fatigue.

The benefit is not just analytical. It improves the quality of the attorney-client conversation. Clients are more receptive to recalibration when it is grounded in how others will view the case, rather than in the lawyer’s subjective judgment. Used well, AI aligns expectations, reduces friction at mediation, and allows counsel to advocate effectively without having to overcome the client’s internal narrative at the same time.

Sample prompts:

- *“You are a neutral mediator evaluating this case for the first time based only on the following summary. Identify the three issues most likely to reduce the plaintiff’s settlement value in your view, and the*

three issues that most support a higher number.”

- *“Our client believes this case is worth \$X based on [reason]. Based solely on the following case summary, does that expectation appear to be supported by the record, or does it rest on assumptions a neutral is unlikely to share? Be specific.”*
- *“Our client is inclined to accept a low offer because she fears [specific defense argument]. Based on the following summary, how likely is that argument to carry decisive weight with a neutral mediator or jury? Are there record strengths she may be underweighting?”*

Prepare for the judge you’ve actually got

Oral hearings are where internal narratives collide with judicial priorities. The risk is not under-preparation, but preparing for the wrong conversation entirely.

AI helps shift preparation from advocacy to anticipation. By asking AI to identify the narrowest grounds on which a judge could rule against you – or the questions a skeptical judge is most likely to ask – counsel can focus on pressure points rather than polishing talking points.

Stress-testing proposed answers is equally important. AI can flag responses that sound persuasive internally but read as evasive, overbroad, or concessionary when stripped of context. This is particularly effective when paired with second-chair review.

AI can also help with preparing for judge-specific advocacy. With proper prompts and source material, it can synthesize prior rulings by a particular judge and identify patterns that suggest what priorities and values appear to guide that judge in deciding similar legal issues. It can help you tailor your rebuttal points to resonate with those priorities – or avoid taking a futile approach.

The result is disciplined advocacy: cleaner answers, fewer unintended concessions, and a stronger record.

Sample prompts:

- *“You are a skeptical federal district court judge. Based on the following motion and opposition, what are the three questions*

you are most likely to ask plaintiff's counsel at oral argument?"

- *"Here is my proposed answer to the question: [question]. Does this answer sound evasive, overbroad, or concessionary when read by someone without our internal case knowledge? Suggest a cleaner version."*
- *"Based on the attached opinions by Judge [X], identify any recurring priorities or values that appear to guide this judge's rulings in [type of case]. How should we frame our argument to resonate with those priorities?"*

What the jury actually hears

Trial teams prepare witnesses for the questions they expect. Defense counsel has been preparing for the questions your witness will not answer well.

AI helps close that gap. By feeding a witness's deposition transcript – or a summary of expected testimony – and asking AI to identify the most damaging lines of cross-examination defense counsel is likely to pursue, you often get a list that is uncomfortably accurate. Not because AI has read the other side's mind, but because it is reading the same record, disinterested.

This exercise is particularly valuable for plaintiffs. Clients who have lived with their injury for years – who have told their story dozens of times – often develop blind spots about how their own testimony sounds to a stranger. A client who answers questions about gaps in treatment with "I couldn't afford it" may be telling the complete truth. But AI will flag that the answer, unaccompanied by context, can be read as minimizing the injury's severity. That is not something the client will notice, and it is not something the lawyer, after the tenth prep session, will notice either.

Ask AI to play skeptical (not hostile) juror – the one in the back row who hasn't decided yet, who is listening for inconsistency and waiting to see if the plaintiff is credible. You cannot afford to lose that juror. This is also where the "what sounds defensive" analysis is most valuable: Tone that feels earnest in prep can read as rehearsed at trial. AI, reading a transcript, catches the rehearsed cadence that the room misses.

One caveat: Do not use AI to generate witness scripts. The goal is diagnosis, not dialogue coaching. The difference matters both ethically and practically – over-prepared witnesses often sound worse than under-prepared ones.

Sample prompts:

- *"You are defense counsel. Based on the following deposition transcript excerpt, identify the five most damaging cross-examination questions you would ask this witness at trial to undermine her credibility on [issue]."*
- *"Read this plaintiff's expected trial testimony summary. You are a skeptical juror who has not yet decided the case. Identify every answer that sounds rehearsed, inconsistent, or evasive to you, and explain why."*
- *"The plaintiff is likely to be asked about [gap/bad fact] on cross. Review the following planned response. Does this answer sound defensive or minimizing to someone encountering the case for the first time? Suggest a more neutral framing."*

Fix the record before it's fixed

Preservation errors rarely feel like errors in the moment. They become obvious only later, on appeal.

AI helps by treating the record literally. Reviewing hearing transcripts or dailies through a cold-eye lens surfaces ambiguity that felt harmless in real time. It flags thin offers of proof, unclear rulings, and objections that assumed understanding rather than stating grounds.

Used post-hearing – not at counsel table – AI enables targeted fixes while there is still time: clarification requests, supplemental offers, or standardized objection language going forward.

Preservation failures are rarely about ignorance. They are about overload. AI creates space to notice what advocacy-focused minds miss.

Sample prompts:

- *"Review the following hearing transcript excerpt. Identify every objection where the grounds were unclear, unstated, or assumed rather than articulated on the record. Flag any ruling that is ambiguous enough to create appellate risk."*
- *"We made the following offer of proof during yesterday's hearing: [text]. Would*

an appellate court reading this cold have a clear understanding of what evidence was being offered, why it was admissible, and what it was offered to prove? If not, identify the gaps."

Appellate hindsight before trial

Most trial teams think about appeal after something goes wrong. A damaging ruling. A surprise verdict. An instruction that didn't quite fit. By then, the record has likely been cemented.

Anticipating appellate risk is most useful as a trial tool. An appellate court will read the record the way AI reads a document: without the benefit of courtroom atmosphere, witness demeanor, or the trial team's running internal commentary. It will ask whether the legal foundation was laid correctly, whether objections were properly preserved, whether rulings were adequate to the issues raised, and whether the verdict is supported by what was actually in evidence – not what everyone in the room understood to be true.

One of the most valuable uses is stress-testing the sufficiency of your evidence on each element as the case develops. Trial teams often feel a case is "in" on an element because witnesses covered the relevant ground. AI, reading a summary of the testimony, may flag that no witness actually said the operative thing – that the conclusion was implied, not proven. That gap is manageable at trial. On appeal, it becomes a sufficiency argument the other side gets to exploit.

The same analysis applies to how key rulings are framed for the record. An evidentiary ruling that feels like a win may carry hidden appellate exposure if the underlying legal theory was not clearly articulated. AI can review the relevant transcript passages and flag whether the ruling would be reviewable, on what standard, and whether the grounds preserved are the grounds most likely to succeed on appeal.

Damages are another underappreciated appellate vulnerability. Jury awards that feel supported in the courtroom sometimes look different when mapped against the specific evidence admitted at trial. Running an AI review of admitted



damages evidence against the verdict form – before closing argument, not after – can surface excessiveness risks while there is still time to adjust how damages are presented.

This is also where jury instruction framing matters most. Instructions that felt like a win during negotiations can create appellate problems if they do not precisely track the legal standard applicable to your facts. AI can compare proposed instructions against the governing standard and flag language that may give the other side a reversible-error argument.

None of this replaces appellate counsel. If the stakes warrant it, bringing in an appellate specialist during trial – not after – is one of the highest-value investments a trial team can make. AI is not a substitute for that judgment. What it can do is help the trial team think with appellate discipline in real time, so that the record being built reflects an awareness of how it will be read by someone who was not in the room.

Sample prompts:

- “Review the following summary of testimony on [element]. Based solely on what witnesses actually said – not what was implied – is each element of [claim] proven? Identify any element where the proof is implied rather than stated, and flag what an appellate court applying *de novo* review might find insufficient.”
- “We received a favorable ruling on [issue] based on [grounds]. Review the following

transcript passage. Would an appellate court have a clear record of the legal theory we preserved? What is the likely standard of review, and are there gaps in the preservation that could limit our ability to defend this ruling on appeal?”

- “Compare the following admitted damages evidence against the verdict form. Identify any category of damages where the award could be characterized as unsupported by admitted evidence, and flag the argument the defense is most likely to make on a sufficiency or excessiveness challenge.”
- “Compare the following proposed jury instruction against [governing standard]. Identify any language that diverges from the precise legal standard and flag whether that divergence could support a reversible-error argument on appeal.”

Where verdicts go sideways

Verdict forms do more than guide deliberations; they structure decision-making. The order of questions, the grouping of elements, and the sequencing of claims can materially affect how jurors reason and what findings they reach. A verdict form that “tracks the causes of action” may still be logically flawed when read sequentially by a juror who does not share the team’s internal understanding of the case. Consider a form that asks jurors to assess damages before they have been explicitly asked to find liability on each theory. It happens. It creates problems.

AI is particularly well suited to this review because it reads the verdict form the way jurors and appellate courts do – literally, one question at a time, without filling in gaps. That perspective makes it easier to spot where questions assume findings that were never required, where sequencing allows damages without clear liability, or where overlapping claims invite inconsistent results.

Addressing those issues early forces alignment between what the law requires, what the evidence supports, and what the jury is actually asked to decide. It is not glamorous work. But it is often the difference between a defensible verdict and a post-trial problem.

Sample prompts:

- “Read the following proposed verdict form as if you are a juror with no prior knowledge of this case. Answer each question in sequence. Flag any place where the question assumes a finding that was never explicitly required, or where the sequencing could lead to an inconsistent result.”
- “We have the following causes of action: [list]. Review the proposed verdict form and identify whether any overlapping claims could produce logically inconsistent jury findings if answered independently.”

Don’t build a better echo chamber

The risk with any new tool is not misuse at the margins, but misplaced confidence at the center. AI can disrupt echo chambers, but it can also reinforce them if used without discipline.

Beware of asking leading questions, iterating until you get a desired answer, and treating that output as confirmation rather than diagnosis. In that form, AI becomes a faster, limitless echo chamber.

Prompts should invite skepticism, not validation. Inputs should be narrow and purposeful. Outputs should be treated as information to evaluate, not conclusions to adopt. When AI flags a problem, the task is to decide whether it identifies a risk that should be addressed, not to blindly accept it.

It is equally important to remember what AI cannot do: read witness demeanor, gauge a courtroom, connect with ju-

rors, or make real-time strategic tradeoffs. Those powers remain the province of experienced trial lawyers.

Used with these guardrails, AI becomes less about efficiency and more about clarity. It helps trial teams surface blind spots while there is still time to act – before strategy hardens, before records close, and before assumptions turn into avoidable risk.

The goal is not to replace judgment. It is to sharpen it.

The cold-eye advantage

The advantage of AI as a cold-eye collaborator is not efficiency. It is perspective. When carefully used as a skeptical diagnostic, AI helps plaintiffs' lawyers

view their cases the way jurors, judges, and appellate courts will – early enough to matter.

And in a practice where the outcome often turns on what the factfinder sees – and what the record shows – that shift in perspective can make all the difference.



Janet Gusdorff is a California Certified Appellate Law Specialist and the founder of Gusdorff Law, P.C. She represents plaintiffs in complex, high-stakes civil appeals and partners with trial lawyers to navigate post-trial strategy, preserve issues, and elevate written advocacy. She can be reached at janet@gusdorfflaw.com.