

Law Firm AI Training Isn't Working. Here's What Has to Change

By Patricia Libby

June 1, 2026


When a partner at a prestigious law firm is forced to issue a lengthy apology for submitting a brief that cited hallucinated cases, the instinct in the legal world is to reach for two familiar tools: better policies and more training.

Both are reasonable responses. But having both isn't enough if neither one is designed to actually change behavior.

A recent Law.com article by Jon Campisi, "As AI Blunders Pile Up, Law Firms Address Ethics Implications of New Technology," gets the starting point right: firms are no longer treating AI risk as theoretical, and training cannot stop at technical tool use. Lawyers need to understand their ethical duties in the context of AI-assisted legal work.

But identifying those duties is only the first step. The harder question is whether training is actually preparing lawyers to carry them out when the AI output looks polished, the deadline is real, and no one wants to slow down.

The Sullivan & Cromwell incident illustrates exactly why that question matters. What makes this particular incident worth examining isn't just that lawyers cited cases that don't exist. It's that they reportedly cited real cases incorrectly, misrepresenting what those cases said, or where they said it. That's a different kind of failure, and a more instructive one. It's not a technology failure. It's a judgment failure: the failure to read



Credit: IRStone/Adobe Stock

critically, to verify skeptically, to ask whether the source actually says what you need it to say. That kind of judgment doesn't come from knowing that AI can hallucinate. It comes from having practiced the discipline of verification enough times, under enough realistic pressure, that it becomes instinct.

Most law firm AI training does not build that instinct. And the gap between what firms think their training is doing and what it's actually doing is where the real risk is.

Policies and Training Are Not the Same Thing, But They're Failing the Same Way

It helps to separate two things that often get conflated: policies and training. Both matter. Both are currently falling short in similar ways.

A policy tells lawyers what they're supposed to do. Verify citations. Use AI as a starting point, not an ending point. Don't input confidential client information into public tools. These are necessary guidelines. But a policy document, however carefully drafted, can't teach a lawyer how to do those things in practice. Reading that you should verify AI output and knowing how to verify it efficiently, skeptically, and consistently under real deadline pressure are completely different things.

Training, in theory, is supposed to bridge that gap. But most AI training at law firms isn't doing that, because most of it is still passive. Picture the typical rollout: a lunch-and-learn when the AI tools go live, a PDF of usage guidelines, maybe a CLE webinar on the ethics rules, a compliance module that ends with a short quiz. These can build awareness. They cannot build skill. And the gap between awareness and skill is exactly where verification failures happen—where a lawyer reads an AI-generated case summary, thinks it sounds right, and moves on.

Awareness Is Not a Skill

Think about learning to play a musical instrument. You can watch video lessons, read about technique, listen to master musicians explain common mistakes, and understand intellectually what good playing sounds like. None of that prepares your hands to find the right notes under pressure, to recover when you lose your place, to make the quick judgments that experienced musicians make without thinking. That only happens by sitting down, playing the wrong notes, feeling the friction of getting it wrong, and doing it again. The discomfort isn't incidental; it's the mechanism. It signals that you're operating at the edge of your current ability, which is exactly where new habits and real judgment form. Comfortable training, by definition, isn't pushing you there.

I've seen this same principle play out closer to home. Two of my sons played competitive water polo. They did not get better at catching passes in game situations by watching their coach diagram plays on a whiteboard. They got better by getting in the pool, having balls thrown at them from bad angles, under fatigue, with defenders

crowding them. They failed repeatedly. They adapted. The friction of failing in a situation that felt real is what made the skill stick. This is deliberate practice. What it means practically is that you have to actually be in the situation or something close enough to it that your brain responds as if the stakes are real.

Responsible AI use in legal practice is that kind of skill. A lawyer who has watched a video about hallucinations is not the same as a lawyer who has been put in a realistic scenario, had to make a judgment call about whether to trust a piece of AI output, felt the discomfort of not being sure, and then received structured feedback on that decision. The first lawyer has information. The second lawyer has the beginning of judgment. Only one of them catches the case that says something different from what the AI claimed it said.

What Experiential Training Actually Looks Like

This is where training design matters most and where most firms are leaving the biggest gap.

Effective AI training isn't a choice between videos and scenarios. Both have a role. Instruction can efficiently introduce concepts and frame the relevant issues. But that's the setup, not the training itself. The real learning happens when lawyers have to apply those concepts in situations that feel like actual practice.

Here's a concrete example of what that looks like in the supervisory context. A partner is reviewing a research memo produced by an associate who used AI to identify and summarize supporting cases. The memo looks clean and well-organized. The citations seem solid. The deadline is tomorrow. What does the partner do? Not in the abstract, but specifically. Do they ask how the associate used AI? Do they pull up the cases themselves? Do they ask the associate to walk them through the verification process? What does adequate supervision actually look like here, and what does it look like when it's falling short in ways that won't be obvious until a judge is asking questions?

A scenario like that, built into training, forces the participant to make an actual decision and sit with the discomfort of not having a clear answer.

Then they see the range of possible responses—not just right or wrong, but what good looks like, what better looks like, and what best looks like. This is how professional judgment actually develops. You see the options, you weigh them, you choose, and you get feedback. Any one piece of that cycle alone doesn't change behavior. It's the combination that does, and most law firm AI training is only delivering the first part.

That's the gap effective training is built to close. Courses that work start with instruction because that foundation matters. But the instruction is the setup, not the point. The real work happens in the scenarios, where participants have to make actual decisions in situations that mirror what they'll face in practice, and in simulated assignments, where they apply what they've worked through in conversation with an expert who pushes their thinking. And that challenge is even sharper for the partners and managers responsible for overseeing AI-assisted work.

The Supervision Gap

Most conversations about AI risk at law firms focus on the lawyers doing the work. That's the wrong place to stop, because the partners and managers who supervise them need this kind of training too, and in some ways their need is greater, because the consequences of their failure are more significant.

The duty to supervise isn't new. What's new is that AI creates work product that looks polished, arrives fast, and can be wrong in ways that aren't easy to spot without knowing what to look for. A supervising partner who has been told to exercise reasonable oversight, but who has never actually practiced what that oversight looks like in an AI-assisted workflow, is operating on instinct shaped by a pre-AI world. That instinct isn't calibrated for catching the specific ways AI fails: the confident citation, the plausible-sounding summary, the case that's real but doesn't say what the AI said it does.

Supervisory attorneys need to practice asking: Did you actually read this case? How did you use AI to get here? Can you show me your verification process? Those questions don't come naturally in a pre-AI supervisory workflow. They have to be practiced.

For supervisors, the challenge runs deeper and the consequences are more significant. When partners and managers are put in scenarios they'll actually face—asked not just to recognize a supervisory failure but to articulate exactly what they would do differently—they may find it harder than they expected. Turning a general awareness of the duty to supervise into a specific set of practiced habits requires real work. That's the point.

The Path Forward

The question worth asking in the wake of any high-profile AI incident isn't whether the firm had policies and training in place. It's whether the training was designed to actually change what lawyers do when the deadline is real and the AI output looks right.

The lawyers who catch the case that doesn't say what the AI claimed it says aren't the ones who watched the most videos. They're the ones who have already been in that situation, even if only in a simulation, felt the friction of not being sure, and had to check anyway. That's what real training builds. And after the last few months, it's what the legal profession needs to start demanding from the programs it puts in place.

Patricia A. Libby is chief learning officer of AltaClaro, an experiential attorney training platform where she oversees the development of innovative legal education content. She writes and speaks on critical thinking and lawyer skills training, and appears at industry conferences and on podcasts. She was previously a litigator and a faculty member at USC's Gould School of Law.