

WEDNESDAY, JULY 9, 2025

AI in law: Optional tool or professional necessity?

AI excels at document review, summarization, and pattern recognition in legal work, but successful implementation requires understanding where these tools add value versus where human judgment remains essential.

By Benjamin Softness

Anyone following the rise and popularity of artificial intelligence is likely aware of stories about lawyers with egg on their face – lawyers who, say, used a popular chatbot to draft a brief containing case law that turned out not to exist. Those stories have served as cautionary tales about the dangers of AI, and as justification for strict limitations on its use by members of the bar. [Concerns about AI use](#) are valid. Lawyers have ethical duties of competence and diligence, among others, and the uncritical use of general-purpose AI to research case law could well violate those duties in some cases.

But AI can also make lawyers *better* at what they do. Large Language Models (LLMs) can summarize in seconds documents that a human might need hours just to read. They can isolate words, phrases, themes, or facts in large document sets faster and more accurately than human reviewers. And they don't get tired.

LLMs can summarize and re-express. Suppose you know what you want to say but you want to say it in half as many words. Or more professionally. Or in French. AI can do it.

To a limited extent, models can even think. A lawyer with nine-tenths of a good idea may be able to wrestle the last 10% to the ground by bouncing ideas off an LLM. (Note: this example is importantly different from having the model come up with the idea on its own.)

Indeed, AI helps more than just lawyers. Early studies suggest that, compared to human drivers, AI-



This art was created with the assistance of Shutterstock AI tools

powered driverless cars cause fewer deaths per mile. AI has solved microbiology's "protein-folding problem." And doctors use AI to help analyze imaging tests like CT scans and X-rays.

This all suggests that, soon, AI use may be effectively required. AI may in some situations add so much to the equation that *not* using it could raise a flag. It may mark a practitioner as behind the times, less-than-fully equipped, or wasteful

of human resources. As AI becomes better understood and more capable, boards will demand its use of their employees, including in-house counsel. (Name an in-house lawyer who hasn't been asked to "do more with less.") In-house lawyers will in turn demand it of their firms. The responsible use of AI could even come to be expected under the legal standard of care, as some have suggested in contexts outside the law.

How can you prepare for this shift? Deploying AI safely and responsibly in the legal context requires certain non-negotiable precautions. Consider these tips when deploying AI more broadly within your organization:

1. Understand your AI's connection to the outside world. One reason publicly available tools like ChatGPT and Google's Gemini get better over time is because of

what's called "reinforcement learning." Positive feedback from users reinforces prior model conduct. Every interaction is a chance to improve. This is fascinating from a computer science perspective but raises real confidentiality concerns in the context of providing legal advice because it means that prompts submitted to, and responses provided by, a public model may be used for future model training. That reality may destroy confidentiality and/or breach the attorney-client privilege.

To mitigate this risk, organizations can license and deploy AI models that are isolated from the outside world – confined within your organization's four walls. Those internal models can be configured not to train on your organization's (or anyone's) interactions with it. With a model like that, and assuming you take other confidentiality precautions, such as limiting access to the logs of a given user's interactions with the model, confidentiality is less elusive.

2. Think carefully about what AI is good and bad at. Most AI is trained, and operates, by processing language. So, it can find a needle in a haystack. It can untangle complex patterns and find structure in unstructured data. It can do Herculean amounts of work, it can do it well, and it can do it fast.

It does not have judgment. It does not know how to find the "right" answer to a judgment-based question that has many right and many wrong answers.

Consider these examples:

- **A legal question:** What are three persuasive in-circuit cases for the proposition that subject-matter jurisdiction is present in State A when Client Co. has sold goods into states neighboring State A with a high but uncertain expectation that the goods will eventually flow into State A's stream of commerce?

- **Discovery questions:** Assume you've uploaded 500,000 documents. You ask the model: How many times in these documents do employees at Example Corp. discuss the ABC Marketing Campaign? How many times do they mention the CEO, Jane Smith? How many of those references are positive? Negative?

- **Writing question:** You've written a two-page single-space update to your general counsel, but she prefers half a page of bullets. You ask the model to shorten and reformat the email.

In the legal example, we are first asking the model to identify a universe of legal authority; keep in mind that we have provided none. It then has to review it, understand it, and make judgments about which cases "persuasively" make the point we want to make about subject matter jurisdiction. Put another way, we're asking the model to know the law. That is a level of judgment and competence that LLMs generally do not yet have, and arguably are not even *designed* to have. (Some models may well be trained on case law. No doubt that helps with a question like this, but it does not substitute for a lawyer's judgment of the many ways in which a decision may or may not be persuasive.)

By contrast, consider the discovery questions. There, you're not asking the LLM to locate sources of authority nor decide whether those sources contain something as elusive as persuasiveness. Rather, you are telling the model exactly *what to review* – the documents that have already been supplied – and you're further telling it *exactly what to find* – documents that mention a particular marketing campaign, or a particular employee. This task is not totally devoid of judgment. But it is far more straightforward as a computing task. Responsibly deploying this capability can save meaningful time.

Ditto the writing question. You have already done the substance of the legal work; the challenge is to summarize that work into a more appropriate form factor. This is relatively low risk, and the AI is fit for purpose.

3. Be smart about compliance. Don't deploy AI in a vacuum. AI is powerful when used responsibly, and it's riskier when it's not, so subject your AI to internal governance. Here are a few examples:

If AI is good at one thing and bad at another (see above), consider formal limitations on how it may be used in your organization. Consider requiring a layer of human review of AI-generated solutions. As with anything else, set and enforce policies governing acceptable use, and then train people to follow those rules. Test your AI to understand its limitations and accuracy. Document that testing. Be able to demonstrate, if needed, that your deployment is a responsible one. And be transparent – with your board if you're in-house, or with firm management if you're in private practice – about the use of AI, and track the benefits it generates.

4. Consider your legal and ethical obligations. In many ways, think of AI as a junior colleague, or a research tool. It can support and improve and accelerate your work, but it does not relieve you of *responsibility for your work*. That mindset can prevent many of the worst outcomes from the uncritical use of AI. Consider, too, court rules and your obligation of candor to the tribunal: in some cases, AI use may be permissible if disclosed.

5. Measure success (and consider the baseline!). Any number of situations may present you with the opportunity – or obligation – to prove the utility of the AI you are using. Best case, you're telling your

board about the accuracy and efficiency gains you've realized through AI automation. Worst case, you're defending a client whose AI use is being scrutinized by an adversary or regulator. In either case, and in many cases in between, you'll be well-served to be able to quantify – whether in speed, cost, accuracy, or something else – AI's positive impact.

Any proper measurement should consider the pre-AI baseline. The fact that an AI made a mistake is not the end of the conversation and need not be a basis for legal liability. The question is whether the situation is better *than it would have been* without AI. This is more logical than measuring yourself against perfection – and it's better advocacy, too. Suppose doctors misdiagnosed cancer in a CT scan 10 times a year without AI, and five times a year when using AI. In the event of an AI-assisted misdiagnosis, one could reasonably argue against liability for using a tool that *reduces* the usual rate of error. Indeed, one could imagine a future plaintiff arguing that misdiagnosis made *without* AI falls short of the standard of care by practicing without a proven diagnostic aide. Standards and arguments could similarly evolve in the legal field.

Benjamin Softness is a partner in King & Spalding's business litigation practice and a member of the firm's Artificial Intelligence task force.

