Human Judgment and AI in Legal Practice

Where Machines End and Lawyers Begin

Colin S. Levy

2026

This document is for informational purposes only and does not constitute legal advice.

Contents

Part One: The New Reality

Generative AI is already part of legal practice. The research is clear on what it does and does not change. A 2024 randomized trial of 60 practicing attorneys found that lawyers using GPT-4 completed tasks significantly faster, but the quality of their work improved only slightly and inconsistently.[1] Speed went up. Judgment did not automatically follow. A 2025 follow-up testing more advanced systems, including RAG tools and reasoning models, showed both speed and quality gains of 50% to 130%, without introducing additional hallucinations.[2] Better-designed tools gave lawyers more room to exercise judgment, not less need for it.

That is the tension that defines this moment. AI can process information at a scale no lawyer can match, but processing is not understanding. Pattern matching is not reasoning. Every lawyer who has reviewed an AI-generated memo that reads beautifully but misses the actual issue knows this. Kate Barton, CEO of Dentons, put it plainly: lawyers must double down on judgment, persuasion, and the ability to build trust.[3] A 2025 Harvard study of AmLaw 100 firms found none anticipated reducing headcount, even as some tasks showed 100x productivity gains.[4] Deloitte projects roughly half of legal workflows will be automated or reengineered by the end of this decade.[5] The work is shifting, not disappearing.

This guide is for lawyers who are already using AI or about to start, and who want to do so without surrendering the thing that makes them valuable: their ability to think, evaluate, and decide.

Part Two: What AI Actually Does Well (and Where It Falls Short)

To use AI effectively, you need an honest accounting of what it can and cannot do. Not the marketing claims. Not the doomsday predictions. The actual, current state of the technology as it applies to legal work.

What AI Does Well

Large language models are, at their core, pattern-recognition engines trained on enormous volumes of text. In legal practice, this makes them useful for specific, well-defined tasks:

- **Document review and summarization.** AI can process hundreds of pages of contracts, briefs, or regulatory filings and produce structured summaries in minutes. For due diligence involving 500-page data rooms, AI can extract key provisions, flag unusual terms, and produce a structured summary that would take a junior associate days to compile manually. This is a genuine efficiency gain, provided someone with judgment reviews the output.
- **Legal research.** AI-powered research tools can identify relevant cases, statutes, and regulatory guidance quickly. RAG-based tools that ground outputs in verified legal databases (such as Lexis+ AI Protege or Westlaw's CoCounsel) perform meaningfully better than general-purpose chatbots because they retrieve from authoritative sources rather than generating from memory.[2] In the 2025 follow-up study, RAG tools improved both speed and quality without increasing hallucination rates.

- **First-draft generation.** AI can produce initial drafts of memos, motions, contracts, and correspondence. For routine documents like NDAs, engagement letters, or standard motions, these drafts can be 70-80% complete, saving significant time. For novel or complex matters, they serve as a structural starting point, not a substantive one.
- **Data extraction and organization.** Pulling key terms from contracts, identifying obligations and deadlines, organizing information into structured formats, and comparing provisions across multiple agreements are tasks where AI performs reliably and at scale.
- **Regulatory monitoring.** AI can continuously track regulatory changes across multiple jurisdictions, compare new requirements against existing compliance programs, and flag gaps. For multi-jurisdictional practices, this kind of automated surveillance would require dedicated headcount to replicate.

Where AI Falls Short

The limitations matter as much as the capabilities, and ignoring them is where lawyers get into real trouble.

- **Hallucination.** AI models generate text that sounds authoritative even when it is fabricated. A 2024 Stanford HAI study tested major legal research tools and found hallucination rates of 17% for Lexis+ AI, 33% for Westlaw AI, and 43% for GPT-4 when used for legal research.[6] These are not edge cases. They are baseline error rates. And the consequences are real: in *Mata v. Avianca* (S.D.N.Y. 2023), attorneys were sanctioned $5,000 for submitting a brief containing six entirely fabricated case citations generated by ChatGPT.[7] Since then, similar sanctions have been imposed in multiple jurisdictions, with one tracker identifying over 600 cases nationwide by mid-2025 where lawyers cited nonexistent authority generated by AI.[8]
- **Contextual reasoning.** AI does not understand the strategic context of a legal matter. It cannot weigh the practical implications of a clause against a client's business objectives, assess how a specific judge in a specific jurisdiction is likely to react to an argument, or factor in the relationship dynamics between the parties. Ask AI to draft a cease-and-desist letter and it will produce a competent template. It will not know that the recipient is your client's largest customer and that a softer approach is warranted.
- **Ethical judgment.** AI has no capacity for ethical reasoning. It cannot identify conflicts of interest, evaluate whether a course of action serves a client's genuine interests, or navigate the gray areas that define much of legal practice. It will not flag that a proposed deal structure, while technically legal, creates reputational risk for the client.
- **Nuance and ambiguity.** Legal work often involves situations where the right answer is not clear, where the law is unsettled, or where the facts are contested. AI tends to present confident answers regardless of underlying uncertainty. It will rarely tell you "this is a close call" or "reasonable lawyers could disagree here," even when that is the most honest and useful answer.
- **Multi-step reasoning under uncertainty.** Complex legal analysis often requires holding multiple uncertain variables in mind simultaneously and reasoning through their interactions. Tax planning, M&A structuring, and multi-party dispute resolution all involve this kind of thinking. AI can process each variable in isolation but struggles with the interconnected judgment calls that experienced lawyers make intuitively.
- **Cross-jurisdictional and cross-cultural reasoning.** AI models are predominantly trained on English-language legal texts, with heavy weighting toward U.S. and U.K. common law. Lawyers working in international transactions, immigration, trade compliance, or any matter that spans legal systems should be aware that AI is far less reliable when reasoning across civil law, common law, and

hybrid systems. It may confidently apply a common law concept like "consideration" in a jurisdiction where it has no meaning, or miss that a contractual obligation enforceable in New York is void under local law in the counterparty's home country. For multi-language matters, AI translation of legal terms often strips critical nuance: the French *obligation de moyens* and *obligation de resultat* distinction, for example, has no direct English equivalent, and AI routinely flattens it.

> *The Thomson Reuters 2025 Future of Professionals report estimates that AI will free up approximately three to five hours per week for the average legal professional.[9] The real question is what you do with those hours. If you use them to take on more volume without deepening your analysis, you are running faster without running better.*

Part Three: The Cognitive Traps

The most significant risk of AI in legal practice is not that the technology fails. It is that lawyers stop doing the mental work that produces good judgment. This is not hypothetical. Cognitive science research has documented specific patterns of human-AI interaction that degrade decision-making, and legal practice is particularly vulnerable to them.

Automation Bias

Automation bias is the tendency to favor information generated by automated systems over information from other sources, including your own analysis. In legal practice, this shows up when a lawyer reads an AI-generated memo and accepts its conclusions without independently evaluating the reasoning, or when a contract reviewer accepts AI's risk flags at face value without reading the underlying provisions.

Research published in *AI & Society* in 2025 found that professionals working with AI decision-support tools showed significant anchoring effects: they tended to adjust insufficiently from AI-generated starting points, even when those starting points contained errors.[10] The more confident the AI's output appeared, the less likely professionals were to question it. A related 2025 study in *Nature Scientific Reports* found that cognitive forcing functions, specifically interventions that require professionals to form their own assessment before seeing AI output, significantly reduced this overreliance.[11]

For lawyers, this is particularly dangerous. Legal analysis depends on skepticism, on testing arguments against counter-arguments, on identifying weaknesses in your own position. If AI gives you a polished first draft, the temptation is to polish it further rather than interrogate its foundations.

What this looks like in practice: A litigation associate receives an AI-generated summary of opposing counsel's motion. The summary identifies three key arguments and recommends a response strategy. The associate drafts the response based on the AI's framing without reading the original motion carefully. The AI missed a procedural argument buried on page 14 that turns out to be dispositive. The associate's work product is built on an incomplete foundation, and the missing argument goes unaddressed.

Deskilling

There is a legitimate concern that routine AI use will erode the skills lawyers develop through repetitive practice. Junior lawyers build legal judgment by doing legal research, by drafting and redrafting, by making mistakes and learning from feedback. If AI handles the initial drafting and research, where does that learning

happen?

Consider the specific skill of case synthesis: reading 15 opinions on a legal issue, identifying how the doctrine has evolved, and understanding which factors drive outcomes in close cases. A lawyer who has done this work dozens of times develops an intuition for how courts think. A lawyer who reads AI-generated summaries of those opinions acquires information but not the underlying pattern recognition that comes from engaging with primary sources.

This concern is not limited to junior lawyers. Experienced practitioners are also vulnerable. A partner who stops reading case opinions because AI summaries are faster may gradually lose the depth of knowledge that made them an authority in their field. A general counsel who delegates all contract review to AI and only skims the output may miss the subtle shifts in vendor behavior that signal a deteriorating relationship. Expertise is not a permanent acquisition. It is maintained through continued engagement with the underlying work. The more experienced you are, the more you have to lose, and the less likely you are to notice the erosion until it matters.

This is not an argument against using AI. It is an argument for being deliberate about how you use it. A junior lawyer who uses AI to produce a first draft and then critically evaluates every citation, tests every argument, and rewrites the analysis in their own voice is still learning. A junior lawyer who submits the AI output with minor formatting changes is not.

False Confidence

AI outputs are fluent and well-structured. They read like the work of a competent professional. This creates a confidence problem: the quality of the writing can mask deficiencies in the quality of the reasoning. This is especially true for audiences who lack the expertise to independently evaluate the substance, including clients, junior team members, and even lawyers working outside their core practice area.

The Stanford hallucination study is worth revisiting here. When legal research tools hallucinate, they do not produce obviously wrong results. They produce plausible-sounding citations to cases that do not exist, or accurate citations to cases that do not stand for the propositions attributed to them.[6] In the *Mata v. Avianca* case, the fabricated citations included realistic-sounding case names, correct court identifiers, and plausible legal reasoning. They looked exactly like real cases.[7]

Confirmation Bias Amplification

AI tools are designed to be helpful, which means they tend to support the framing of your question rather than challenge it. If you ask AI to "find cases supporting the argument that the contract is enforceable," it will do exactly that, even if the better answer is that enforceability is genuinely in doubt. This creates a feedback loop: the lawyer's initial hypothesis gets reinforced by AI output, and the critical evaluation that should follow never happens.

Practical countermeasure: After any AI-assisted research or analysis, explicitly ask the AI for the strongest counterarguments to your position. Compare the strength of those counterarguments against your original analysis. If you would not have identified those counterarguments on your own, that is a signal that you need to engage more deeply with the underlying material.

Cognitive Offloading

Cognitive offloading is the practice of delegating mental tasks to external tools rather than performing them internally. We all do it: writing a phone number down instead of memorizing it, using a calculator instead of doing math in our heads. In moderation, offloading is efficient and rational. The danger with AI is the scale and depth of offloading it enables. When a lawyer uses AI to research, analyze, draft, and organize, the lawyer may be offloading not just mechanical tasks but the cognitive processes that produce understanding and judgment.

A 2025 study published in *Societies* surveyed 666 participants across diverse age groups and professional backgrounds and found a significant negative correlation between frequent AI tool usage and critical thinking abilities, with cognitive offloading as the mediating mechanism.[14] The more participants relied on AI for thinking tasks, the less they exercised their own analytical capabilities, and the measurable decline was most pronounced among younger, less experienced users. A separate 2025 paper in *Nature Communications* framed the core tension: generative AI can function as either a mind-extending technology or a mind-replacing one, depending entirely on how the user engages with it.[15] That distinction matters for lawyers. When you use AI to pull together background research and then read, evaluate, and synthesize that research yourself, AI is extending your mind. When you use AI to produce an analysis and adopt its conclusions without doing the underlying thinking, AI is replacing your mind. The first approach builds judgment. The second erodes it.

The erosion goes deeper than analytical reasoning. A 2025 randomized controlled trial by Barcaui, published in the *International Journal of Information Management Data Insights*, found that students who used ChatGPT as a study tool scored significantly lower on knowledge retention tests 45 days later (57.5% correct) compared to those who studied using traditional methods (68.5% correct).[16] The mechanism is straightforward: when AI handles the effortful processing, the brain does not encode the information as deeply. For lawyers, this means that routinely delegating research and analysis to AI may weaken the very knowledge base that judgment depends on. You cannot exercise good judgment about contract law if you have not internalized the principles through repeated, effortful engagement.

What this looks like in practice: A transactional lawyer uses AI to generate a risk assessment for a vendor agreement. The AI identifies five risk areas and rates each as high, medium, or low. The lawyer copies the assessment into a memo, adjusts the formatting, and sends it to the client. The lawyer has not actually evaluated whether those risk ratings make sense for this particular client, this particular vendor relationship, or this particular deal structure. The cognitive work of risk assessment, the part that requires judgment, has been offloaded entirely. In a second scenario, a litigation associate uses AI to summarize 20 depositions. Six months later, when preparing for trial, the associate cannot recall key testimony because they never actually read the transcripts. The knowledge was never encoded; it was outsourced.

Practical countermeasures: Be intentional about which cognitive tasks you delegate and which you retain. As a rule, offload information gathering but retain analysis. Offload first-draft generation but retain evaluation and revision. Offload data organization but retain interpretation. For high-stakes matters, do your own analysis first and use AI as a check, not the other way around. If you find yourself spending more time formatting AI output than thinking about its substance, that is a warning sign. And periodically work without AI entirely, particularly on tasks that build the foundational knowledge your judgment depends on. The goal is to keep the cognitive muscles active, not to let them atrophy.

> *Practical check: If you cannot explain the reasoning behind an AI-generated argument without referring back to the AI's output, you do not understand it well enough to rely on it. If a supervising partner or judge asked you to walk through the analysis step by step, could you do so from memory? If not, the AI is doing your thinking for you.*

Part Four: Building a Human-AI Workflow

The goal is not to use AI less or to use it more. It is to use it in a way that preserves and strengthens your judgment rather than bypassing it. What follows is a concrete framework for doing that, with specific steps you can implement immediately.

Step 1: Define the Task Boundaries

Before using AI on any matter, decide which tasks are suitable for AI assistance and which require unassisted human work. Not every task in a matter should involve AI, and not every task needs to be done from scratch. The following framework maps the division of labor:

Task Type	AI Role	Human Role	Example
Initial research	Generate search queries, identify sources, summarize findings	Evaluate relevance, verify accuracy, assess applicability	AI pulls 40 potentially relevant cases; lawyer reads the 8 that matter and assesses their weight
Document review	Flag key terms, extract provisions, identify patterns	Assess significance, identify strategic implications, decide materiality	AI flags 23 indemnification clauses across a data room; lawyer evaluates which create actual exposure
Drafting	Produce first drafts of routine docs, suggest alternatives	Evaluate tone/strategy, align with client goals, apply judgment	AI drafts NDA; lawyer adjusts non-compete scope based on client's industry and deal dynamics
Analysis	Organize data, identify patterns, generate frameworks	Apply reasoning, weigh competing factors, reach conclusions	AI maps regulatory requirements across 5 states; lawyer prioritizes based on enforcement risk
Client advice	Organize talking points, draft outlines of complex advice	Decide what to communicate, tailor to context, manage expectations	AI prepares risk summary; lawyer decides how to frame the options given the board's risk tolerance

Step 2: Implement Verification Checkpoints

Every AI-assisted workflow needs explicit verification steps built into the process, not tacked on at the end. This is about maintaining the standard of care that legal practice demands.

- **Source verification.** Every citation, case reference, and statutory reference must be independently verified. Open the case. Read the holding. Confirm it says what the AI claims. Given the hallucination rates documented by Stanford (17-43%),[6] this is not a "nice to have." It is a professional obligation. The ABA's Formal Opinion 512, issued in July 2024, makes this explicit: lawyers must verify AI-generated outputs and cannot rely on AI tools as a substitute for their own competent analysis.[12]
- **Reasoning review.** Do not just check whether the AI cited the right cases. Evaluate whether the reasoning connecting those citations to the conclusion is sound. AI can assemble correct components into an incorrect argument. Test the logical chain: does the rule from Case A actually apply to the facts of your matter? Does the AI's analogy hold up under scrutiny?

- **Completeness check.** AI routinely omits relevant authorities, counterarguments, or considerations that a thorough human analysis would include. Ask yourself: What is the strongest argument against this conclusion? What jurisdictional variations might apply? What practical consequences does this analysis overlook?
- **Factual accuracy.** AI frequently gets basic facts wrong in ways that look correct on the surface. Check dates, dollar amounts, party names, regulatory thresholds, and statutory deadlines independently. AI may cite the right statute but give the wrong effective date, name the right agency but describe the wrong enforcement action, or produce a damages figure that conflates two different cases. These are not hallucinations in the traditional sense; they are factual errors embedded in otherwise accurate text, which makes them harder to catch.
- **Context alignment.** Confirm that the AI output accounts for the specific facts, jurisdiction, and strategic posture of your matter. A memo on California employment law is not useful if your client operates in Texas. An aggressive litigation posture may be appropriate in one matter and counterproductive in another.

Step 3: Preserve the Thinking

The most important part of legal work is often the reasoning process itself, not the final product. When you use AI, build habits that keep you actively engaged in that process.

- **Write your own issue analysis before consulting AI.** Spend 15 minutes outlining the key issues, your initial take on each, and the questions you need answered. This gives you a baseline against which to evaluate AI output. Without this, you have no independent anchor for your judgment, and you will default to whatever the AI produces.
- **Argue with the AI's conclusions.** Take the opposing position. Identify the weakest points in the AI-generated analysis. Ask the AI to steelman the counterargument. If you cannot find weaknesses in the analysis, you are not looking hard enough.
- **Document your reasoning.** When you modify, accept, or reject AI-generated work product, note why in a brief annotation. This creates a record of your professional judgment, reinforces the habit of active evaluation, and provides a defensible record if your work is later questioned.
- **Use the "explain it cold" test.** Before relying on any AI-assisted analysis, close the AI tool and explain the reasoning to yourself or a colleague from memory. If you cannot do this, you have adopted the AI's judgment rather than exercising your own.

Step 4: Build Feedback Loops

Track your experience with AI over time. Which tasks consistently produce reliable output? Where do errors tend to appear? What kinds of prompts yield better results? Maintain a running log, even a simple one. Over time, this becomes an invaluable guide to where AI adds genuine value in your specific practice and where it requires heavy intervention.

> *A useful framework: Think of AI-assisted legal work in three tiers. Tier 1 tasks (data extraction, document organization, initial research) are high-confidence AI territory. Tier 2 tasks (first drafts, comparative analysis, issue spotting) require active review. Tier 3 tasks (strategy, client counseling, ethical judgment) are human-primary and should use AI only for background preparation.*

Addressing the Verification Burden

There is a fair and increasingly vocal criticism of AI in legal practice that deserves a direct response: the verification burden. The argument, formalized by legal scholar Joshua Yuvaraj in his 2025 paper on what he calls the "verification-value paradox," goes like this: if AI output requires thorough verification to meet professional obligations, and if that verification takes significant time and effort, then the efficiency gains from using AI may be partially or entirely negated.[17]

This is a serious concern, and dismissing it does lawyers no favors. The verification burden is real. When AI hallucinates case citations at rates of 17-43%,[6] when over 600 cases of lawyers citing fabricated authority have been documented nationwide,[8] and when sanctions have reached $10,000 or more in individual cases, the cost of trusting without verifying is clear. And the time required to check every citation, evaluate every chain of reasoning, and confirm every factual claim can be substantial.

But the criticism, while valid as a caution, does not tell the whole story.

The verification burden is not the same across all tasks. Document organization, data extraction, contract provision mapping, and regulatory monitoring produce outputs that are relatively easy to verify and consistently reliable. The paradox bites hardest in legal research and novel legal analysis, where hallucination rates are highest and the consequences of error are most severe. Most legal work involves a mix of task types. Many of them sit well outside the danger zone.

The comparison baseline also matters. The question is not whether AI output requires verification. Of course it does. The question is whether the total time, generation plus verification, is less than doing the work from scratch. For a 200-page contract review, AI can extract and flag relevant provisions in minutes. Even with substantial verification, the total time is typically a fraction of a manual review. For a straightforward legal research question, the math may be different, particularly if the AI hallucinates and the lawyer has to redo the research anyway.

Better tools are also narrowing the gap. The 2025 follow-up study by Schwarcz and colleagues found that RAG-based tools achieved productivity gains of 50-130% without increasing hallucination rates.[2] Legal-specific AI grounded in verified databases is significantly more reliable than general-purpose chatbots. The verification burden you face with GPT-4 is not the same burden you face with purpose-built legal AI.

And here is the part that gets overlooked: verification is not wasted work. When you verify a citation, you are engaging with the case law. When you evaluate whether AI's reasoning holds, you are exercising judgment. The verification step is not a tax on productivity. It is the part of the work where the lawyering actually happens. The problem arises only when lawyers treat verification as a box-checking exercise rather than an analytical one.

Practical approach to managing the burden:

- **Match AI tools to task risk.** Use general-purpose AI for low-stakes tasks (brainstorming, outlining, internal summaries). Use legal-specific RAG tools for research and analysis where accuracy matters. Reserve manual work for the highest-stakes elements.
- **Verify in layers.** Not every element of AI output needs the same depth of review. Spot-check routine provisions. Deep-verify every citation in a court filing. Scale your verification to the risk profile of the task and the audience for the work product.
- **Build verification into the workflow, not after it.** If you treat verification as a separate step at the end, it feels burdensome. If you engage critically with AI output as it is generated, evaluating each section as you go, verification becomes part of the analytical process rather than an afterthought.

- **Track your ratios.** For each type of task, note how long AI generation takes, how long verification takes, and how long the task would take without AI. Over time, this data tells you exactly where AI delivers net value and where it does not. The answer will vary by practice area, task type, and the specific tools you use.

> *The verification burden is not a reason to avoid AI. It is a reason to use AI strategically, with clear-eyed awareness of which tasks benefit from AI assistance and which do not. The lawyers who figure this out will capture real efficiency gains. The lawyers who use AI indiscriminately will spend their time cleaning up after it.*

Part Five: Judgment in Practice

Abstract principles are useful, but lawyers need concrete examples. What follows is a detailed look at how human judgment and AI interact across several common practice scenarios, with specific guidance on where to draw the lines.

Contract Negotiation

AI can review a contract and identify provisions that deviate from standard market terms. It can flag risk allocation clauses, compare indemnification language against templates, generate redline suggestions, and even benchmark specific terms against market data.

What it cannot do is assess whether a particular risk allocation makes sense for this deal. A limitation of liability clause that looks aggressive on paper might be perfectly acceptable given the client's risk appetite, the size of the deal, the strength of the relationship, or the availability of insurance. A non-compete that AI flags as "overly broad" might be standard for the specific industry. That assessment requires human judgment informed by experience and knowledge of the client.

Practical workflow: (1) Run the contract through AI for an initial deviation analysis against your standard playbook. (2) Review each flag and categorize it: accept, reject, or requires further analysis. (3) For provisions requiring analysis, read the underlying language yourself and assess in context. (4) Draft your negotiation positions yourself, informed by but not dictated by the AI's flags. (5) Use AI to check your redline for internal consistency and missed cross-references.

Litigation Strategy

AI can analyze court filings, identify relevant precedent, summarize deposition transcripts, and surface patterns across large document sets. The Minnesota Law Review study showed that AI meaningfully sped up the research and drafting phases of litigation work.[1]

But litigation strategy involves reading the room. It involves assessing witness credibility, anticipating opposing counsel's moves, making judgment calls about what arguments to prioritize based on the specific judge and jurisdiction, and knowing when to settle versus when to push. These are decisions that depend on experience, intuition, and a kind of practical wisdom that AI does not possess.

Practical workflow: (1) Use AI for document review and initial research to build your factual foundation quickly. (2) Have AI summarize each deposition by witness and topic, but read the key testimony yourself, especially anything you might use at trial. (3) Use AI to identify the universe of potentially relevant

precedent, then read and assess the most important cases personally. (4) Develop your case theory, motion strategy, and settlement position through your own analysis. (5) Use the time AI saves you on the mechanical work to invest more deeply in witness preparation and trial strategy.

Regulatory Compliance

AI excels at monitoring regulatory changes, mapping requirements across jurisdictions, and identifying potential compliance gaps. For organizations operating under multiple regulatory frameworks (GDPR, CCPA, HIPAA, SOX), AI can maintain continuous surveillance that would require significant dedicated headcount to replicate manually.

The judgment component comes in prioritization and interpretation. Not every regulatory change requires the same response. A lawyer's value lies in assessing which changes are material for a specific business, how to interpret ambiguous requirements (regulators themselves often disagree), and how to balance compliance costs against business objectives.

Practical workflow: (1) Set up AI-powered monitoring for relevant regulatory sources and industry-specific alerts. (2) Have AI produce weekly or monthly summaries of changes, categorized by jurisdiction and subject matter. (3) Review the summaries and flag items that require action, using your knowledge of the business to assess materiality. (4) For complex or ambiguous requirements, conduct your own analysis rather than relying on AI interpretation. (5) Build compliance response plans that reflect practical business constraints, not just legal requirements in the abstract.

Client Counseling

This is the area where human judgment is most irreplaceable. Client counseling involves empathy, reading emotional and political dynamics, and delivering difficult advice in a way that maintains trust. It means knowing when a client needs to hear the hard truth and when they need reassurance. It means picking up on what the client is not saying.

As Kate Barton of Dentons noted, the skills that matter most for lawyers in an AI-enabled environment are judgment, persuasion, empathy, and the ability to build trust.[3] These are not soft skills or secondary capabilities. They are the core competencies of client service, and they are what differentiate a trusted advisor from a legal information provider.

Practical workflow: (1) Use AI to prepare thorough briefing materials before client meetings: relevant facts, risk assessment, option analysis, and potential talking points. (2) Review the materials and form your own recommendation before the meeting. (3) Walk into the conversation prepared but flexible. Listen first. Ask questions. (4) Adjust your advice based on what you hear, including the emotional and political dimensions that AI cannot capture. (5) Follow up with a written summary that reflects what was actually discussed, not what the AI predicted would be discussed.

M&A Due Diligence

AI transforms the efficiency of due diligence review but cannot replace the judgment calls that determine whether a deal should proceed. AI can extract and organize key provisions from hundreds of contracts in a data room, flag change-of-control clauses, identify unusual termination rights, and surface patterns in employment agreements.

Practical workflow: (1) Use AI to conduct a first-pass extraction of key terms across the entire data room. (2) Review the AI's output against a customized diligence checklist. (3) For each material finding, go back to the source document and read the relevant provisions in context. (4) Assess risk in light of the deal structure, the buyer's integration plans, and the commercial terms being negotiated. (5) Apply judgment to determine which findings are deal-breakers, which are price adjustments, and which are acceptable risks.

Intellectual Property

AI can accelerate prior art searches, compare claim language against existing patents, organize prosecution histories, and flag potential conflicts in trademark clearance. For patent portfolio analysis, AI can map claims across hundreds of patents and identify coverage gaps or overlap faster than any manual review.

Where AI falls short is in the judgment calls that define IP practice. Claim construction involves interpreting language in light of the prosecution history, the specification, and evolving case law. AI can surface the relevant materials, but it cannot assess whether a claim term will be construed broadly or narrowly by a particular judge applying Phillips or its progeny. In trademark clearance, AI can identify potentially conflicting marks, but evaluating likelihood of confusion requires weighing the DuPont factors in context: the strength of the senior mark, the similarity of the goods, the sophistication of the buyers, and the channels of trade. AI will give you a list. A lawyer gives you a risk assessment.

Practical workflow: (1) Use AI to conduct initial prior art or clearance searches and organize results by relevance. (2) Review the top results yourself, reading the actual patent claims or trademark registrations rather than AI summaries. (3) For patent drafting, use AI to generate a first pass of dependent claims based on the specification, then critically evaluate whether each claim adds meaningful scope. (4) For trademark clearance opinions, use AI to build the factual record but write the likelihood-of-confusion analysis yourself. (5) In prosecution, use AI to draft office action responses as a starting point, but ensure the arguments align with the specific examiner's prior rejections and the prosecution history.

Employment and Labor Law

AI is useful for reviewing employee handbooks, comparing policies against current regulations, analyzing wage and hour compliance across multiple jurisdictions, and drafting template employment agreements. For multi-state employers, AI can map the patchwork of state and local employment laws, from paid leave requirements to non-compete enforceability, much faster than manual research.

The judgment-intensive work in employment law involves factual nuance and human dynamics that AI cannot evaluate. Whether a termination decision creates meaningful litigation risk depends on the employee's protected characteristics, performance documentation quality, the manager's communication history, the company's past practices with similarly situated employees, and the enforcement climate in the relevant jurisdiction. AI can help you organize these factors, but it cannot weigh them the way an experienced employment lawyer does. Accommodation requests, harassment investigations, and reductions in force all involve sensitive human judgment that goes well beyond legal analysis.

Practical workflow: (1) Use AI to audit handbooks and policies against current federal, state, and local requirements, generating a gap analysis. (2) Review each flagged gap yourself, assessing whether it reflects a genuine compliance risk or a difference in policy drafting style. (3) For termination risk assessments, use AI to compile the documentation record and identify potential legal claims, but make the risk call yourself based on the full factual picture. (4) For wage and hour compliance, use AI to map requirements across jurisdictions but verify threshold calculations (overtime exemptions, minimum wage

rates, tip credits) against current published figures. (5) For investigations, use AI to organize witness statements and timelines, but conduct the credibility assessments and conclusions yourself.

Part Six: The Skills That Matter Now

If AI is absorbing an increasing share of the information-processing work in legal practice, the premium shifts to the skills that AI cannot replicate. These are not new skills. They are the skills that have always defined great lawyers. But the case for investing in them has never been stronger, and the consequences of neglecting them have never been higher.

Critical Thinking

The ability to evaluate arguments, identify logical flaws, and test conclusions against evidence is the foundation of legal judgment. AI cannot do this. It can generate arguments that follow patterns it has seen in training data, but it cannot assess whether those arguments are sound in a specific context.

How to strengthen it: Make it a habit to identify the three strongest counterarguments to any position before finalizing your analysis. When reviewing AI output, start by looking for what is wrong rather than confirming what looks right. Read opinions you disagree with. Argue the other side of issues you feel confident about. The goal is to train yourself to question effectively, because that is the skill AI lacks entirely.

Emotional Intelligence

Legal work is human work. Clients come to lawyers during some of the most stressful moments of their lives and businesses. The ability to read emotional dynamics, manage difficult conversations, and maintain trust under pressure is irreplaceable, and it is what clients consistently value most in their outside counsel relationships.

How to strengthen it: Practice active listening in every client interaction. Ask open-ended questions before jumping to solutions. Pay attention to what the client is not saying as much as what they are. After difficult conversations, reflect on what you learned about the client's concerns beyond the strictly legal issues. In negotiations, invest time understanding the other side's motivations and constraints before focusing on the terms.

Strategic Reasoning

Strategy requires weighing multiple uncertain variables against each other and making decisions under conditions of incomplete information. It requires understanding not just what the law says, but what the parties want, what they can live with, and what will happen if negotiations break down.

How to strengthen it: For every significant matter, map out the decision tree. What are the possible outcomes? What is the probability and impact of each? What are the second-order effects? AI can help you gather the inputs for this analysis, but the synthesis and judgment calls are yours. Practice scenario planning: what happens if we win this motion? What happens if we lose? What does the other side do in each case?

Communication

The ability to explain complex legal issues in plain language, to tailor advice to a specific audience, and to advocate persuasively remains central to legal practice. AI can generate technically correct explanations, but effective communication requires understanding your audience, reading their reactions, and adjusting in real time.

How to strengthen it: Practice explaining legal concepts to non-lawyers. Write client-facing summaries before generating AI drafts. When presenting options to a client, lead with what matters to them (cost, timing, risk), not with the legal analysis. Record yourself presenting arguments and review the recordings. The best communicators are the ones who practice the most.

Effective AI Prompting

This is the one genuinely new skill on the list. Getting useful output from AI requires knowing how to frame questions, provide context, and structure requests. Lawyers who develop strong prompting skills will get meaningfully better results from the same tools that their peers use poorly.

How to strengthen it: Be specific about the jurisdiction, legal standard, and type of analysis you need. Provide context: the type of client, the stage of the matter, the audience for the work product. Define quality standards: tell the AI to note uncertainty, flag assumptions, and identify counterarguments. Iterate: review the first output, identify what is missing or wrong, and refine your prompt accordingly. Treat AI as a capable but literal-minded associate who needs clear, detailed instructions and whose work always requires review.

What good prompting looks like varies by practice area. A litigator researching a motion to dismiss should specify the court, the procedural posture, the specific rule (e.g., 12(b)(6) vs. 56), the legal standard the court applies, and whether the goal is to find supporting or opposing authority. A transactional lawyer reviewing an indemnification clause should tell the AI the deal type, the client's position (buyer vs. seller), the industry, and whether to compare against market standards or a specific template. A regulatory lawyer should specify the agency, the applicable regulatory framework, the compliance timeline, and whether the analysis is for initial assessment or enforcement response. A patent practitioner should specify the technology field, the prosecution stage, and whether the task involves prior art analysis, claim construction, or freedom-to-operate. The more specific the context, the more useful the output, and the easier it is to verify.

Part Seven: Making It Work in Your Organization

Individual skill development matters, but so does organizational design. How a firm or legal department structures its AI use will determine whether AI amplifies collective judgment or gradually erodes it.

Training That Goes Beyond Button-Pressing

Most AI training programs for lawyers focus on tool mechanics: how to log in, how to enter a prompt, what buttons to press. That is necessary but grossly insufficient. Effective training must also address the harder skills:

- **Critical evaluation.** How to review AI output systematically. Specific exercises: give lawyers AI-generated memos with intentional errors and have them identify the problems. Review real examples of AI hallucinations from published court cases. Build the muscle of skepticism through repeated practice.

- **Task selection.** When to use AI and when not to. Create clear guidelines for your organization: which tasks are approved for AI assistance, which require human-only work, and which fall in a gray zone requiring senior approval.
- **Failure mode recognition.** What AI-generated errors actually look like, with concrete examples drawn from the growing body of cases where AI output has led to professional sanctions.[7] [8] Teach lawyers to recognize the telltale signs: overly confident language, suspiciously perfect case summaries, and citations that cannot be verified in standard databases.
- **Ethical obligations.** The ABA's Formal Opinion 512 identifies six areas of ethical concern for lawyers using generative AI: competence, confidentiality, client communication, candor toward tribunals, supervisory responsibilities, and fees.[12] Training should cover each area with practice-specific examples.

Junior Lawyer Development

This deserves special attention because the stakes are high and the solutions are not obvious. The traditional model of legal training depends on junior lawyers doing foundational work: the research, the first drafts, the document review. If AI handles that work, the development pipeline breaks. And unlike other business challenges, the consequences of a broken development pipeline do not show up immediately. They show up five to ten years later, when you discover that your mid-level associates lack the judgment that comes from having done the foundational work themselves.

Practical approaches:

- **"Do it first" rule.** Require junior lawyers to complete their own analysis before consulting AI. They can then compare their work against the AI's output, which becomes a powerful learning exercise. Where did they miss something the AI caught? Where did they identify something the AI missed?
- **AI as a teaching tool.** Use AI-generated work product in training sessions. Have junior lawyers identify errors, evaluate reasoning quality, and rewrite sections in their own voice. Finding AI mistakes is excellent training for developing critical evaluation skills.
- **AI-free rotations.** Designate certain matters or practice periods as AI-free, where junior lawyers do the foundational work unassisted. This ensures they build core skills even as AI becomes standard in other contexts.

The Harvard Center on the Legal Profession study found that firms are not reducing headcount despite massive productivity gains.[4] This suggests the work is being redistributed, not eliminated. The question is whether that redistribution is happening in a way that develops the next generation of skilled lawyers.

Quality Assurance Protocols

Organizations need clear, written protocols for AI-assisted work product. At minimum:

- **Disclosure standards.** When and how to disclose AI use to clients, courts, and counterparties. An increasing number of jurisdictions now require disclosure of AI use in court filings.[13] Several federal courts have adopted standing orders requiring attorneys to certify whether AI was used in the preparation of filings, and state bars are issuing guidance at a rapid pace.
- **Tiered review requirements.** Define who reviews AI-assisted work product and what level of review is expected for different task types. A routine contract summary may require different review than an AI-assisted brief filed with a court.

- **Error tracking.** Systematically monitor AI performance to identify patterns in errors and adjust workflows accordingly. Track false positives, missed issues, and hallucination rates by tool, task type, and practice area.
- **Confidentiality safeguards.** Establish clear rules about what client information can be entered into AI tools. ABA Formal Opinion 512 emphasizes that boilerplate consent in engagement letters is insufficient; lawyers must understand how AI tools process data and obtain informed consent.[12]
- **Accountability.** Clarify that the lawyer, not the AI, is responsible for the final work product. This is not just an ethical principle; it is the legal reality in every jurisdiction. No court has accepted "the AI got it wrong" as a defense.

Culture

The organizations that will use AI most effectively are those that foster a culture of intellectual honesty: where asking questions is valued over appearing certain, where admitting you verified an AI's work is a sign of diligence rather than weakness, and where the emphasis is on getting it right rather than getting it done fast.

This means leadership needs to model the behavior. Partners and general counsel who openly discuss the limitations of AI, who share examples of errors they caught, and who reward thoroughness over speed will create environments where AI amplifies good judgment rather than enabling shortcuts.

Special Considerations for In-House Teams

In-house legal departments face a different set of pressures around AI adoption. The business moves fast, headcount is constrained, and stakeholders often expect legal to keep pace with decisions that are already in motion. AI is attractive in this environment precisely because it promises speed. But that same pressure makes in-house lawyers more vulnerable to the cognitive traps described earlier, especially automation bias and cognitive offloading.

A general counsel overseeing a small team that supports a fast-growing company may be tempted to use AI to handle routine contract review, compliance monitoring, and board preparation all at once. That can work, but only with clear boundaries. The risk is that AI becomes a substitute for headcount rather than a supplement to judgment. When the legal team is stretched thin and AI is doing the first pass on everything, who is doing the second pass? If the answer is "nobody, because we are too busy," then AI is not reducing risk. It is obscuring it.

Practical steps for in-house teams: Identify the three to five highest-risk workflows where AI assistance adds the most value and focus your verification resources there. Build templates and playbooks that encode your organization's risk tolerance so that AI output can be measured against a known standard. Communicate to business stakeholders that AI speeds up legal review but does not eliminate the need for it. When you push back on unrealistic timelines, frame it in terms of risk management, not process for its own sake. And invest in relationships with outside counsel who can serve as a verification backstop for the matters where your team's bandwidth is thinnest.

Conclusion

The story of AI in legal practice is not a story about technology replacing lawyers. It is a story about the nature of legal work itself, and what happens when the information-processing component of that work becomes automated.

What remains, and what becomes more valuable, is everything that requires human cognitive power: the ability to reason about uncertainty, to understand what clients actually need, to navigate ethical complexity, and to make decisions that account for context, relationships, and consequences that no model can anticipate.

The data supports this. Firms are not reducing headcount. Clients are not replacing their lawyers with chatbots. What is changing is the mix of work. The routine, information-processing tasks that used to fill much of a lawyer's day are being absorbed by machines. What is left, and what clients are willing to pay for, is the judgment.

AI is a powerful tool. But tools do not exercise judgment. They do not read the room, weigh competing interests, or take responsibility for the advice they give. Lawyers do.

The path forward is not to resist AI or to embrace it without thinking. It is to use it deliberately: with specific workflows, clear verification habits, and honest self-assessment about where your judgment is adding value and where you are handing it off to the machine.

That is what this guide is about. Not the technology. The thinking that makes you worth hiring.

Endnotes

1. Choi, Monahan, and Schwarcz, "Lawyering in the Age of Artificial Intelligence," Minnesota Law Review, Vol. 109, 2024. Available at https://minnesotalawreview.org/article/lawyering-in-the-age-of-artificial-intelligence/

2. Schwarcz, Choi, and Monahan, "AI-Powered Lawyering," SSRN Working Paper, March 2025. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5162111

3. Kate Barton (CEO, Dentons), interview with Fortune, "Why the World's Largest Law Firm CEO Says Lawyers Must Double Down on EQ," July 2025. Available at https://fortune.com/2025/07/14/lawyers-artificial-intelligence/

4. Harvard Center on the Legal Profession, "The Impact of Artificial Intelligence on Law: Law Firms' Business Models," 2025. Available at https://clp.law.harvard.edu/knowledge-hub/insights/the-impact-of-artificial-intelligence-on-law-law-firms-business-models/

5. Deloitte Legal, "AI and the Legal Profession: Preparing for a 50% Shock," 2025. Available at https://legalbriefs.deloitte.com/post/102kdm3/ai-and-the-legal-profession-preparing-for-a-50-shock

6. Magesh et al., "Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools," Journal of Empirical Legal Studies, 2025. Originally published as Stanford HAI working paper, 2024. Available at https://onlinelibrary.wiley.com/doi/full/10.1111/jels.12413

7. Mata v. Avianca, Inc., No. 22-cv-1461 (PKC) (S.D.N.Y. June 22, 2023). Attorneys sanctioned $5,000 for submitting brief containing six fabricated case citations generated by ChatGPT.

8. As reported by multiple legal trackers by mid-2025, over 600 cases nationwide have involved lawyers citing nonexistent authority generated by AI tools. See, e.g., Cronkite News/Arizona PBS, "As More Lawyers Fall for AI Hallucinations," October 2025.

9. Thomson Reuters, "Future of Professionals Report," 2025.

10. Chiang et al., "Exploring Automation Bias in Human-AI Collaboration," AI & Society, Springer, 2025. Available at https://link.springer.com/article/10.1007/s00146-025-02422-7

11. Vejandla et al., "Impacts of Cognitive Forcing and Need for Cognition on Biased AI-Assisted Decision Making About Mental Health Emergencies," Scientific Reports, Vol. 16, 2026. Available at https://www.nature.com/articles/s41598-025-30506-3

12. ABA Standing Committee on Ethics and Professional Responsibility, Formal Opinion 512, "Generative Artificial Intelligence Tools," July 29, 2024. Addresses competence, confidentiality, communication, candor, supervision, and fees.

13. See, e.g., Standing Order Re: Artificial Intelligence, U.S. District Court for the Northern District of Texas (2024); various federal and state court orders and state bar guidance on AI disclosure obligations.

14. Gerlich, M., "AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking," Societies, Vol. 15, No. 1, 2025. Available at https://www.mdpi.com/2075-4698/15/1/6

15. Clark, A., "Extending Minds with Generative AI," Nature Communications, Vol. 16, Article 4627, 2025. Available at https://www.nature.com/articles/s41467-025-59906-9

16. Barcaui, A., "ChatGPT as a Cognitive Crutch: Evidence from a Randomized Controlled Trial on Knowledge Retention," International Journal of Information Management Data Insights, 2025. Available at https://doi.org/10.1016/j.jjimei.2025.100367

17. Yuvaraj, J., "The Verification-Value Paradox: A Normative Critique of Gen AI in Legal Practice," SSRN Working Paper, October 2025; forthcoming, Monash University Law Review. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5621550