

AI Use by Lawyers Handbook

Editors:

Meredith Ervine

Senior Editor, TheCorporateCounsel.net

Meaghan Nelson

Partner, Stoel Rives LLP, and
Contributing Editor, TheCorporateCounsel.net

February 2026

Chapter 29

AI Use by Lawyers Handbook

Editors:

Meredith Ervine

Senior Editor, TheCorporateCounsel.net

Meaghan Nelson

Partner, Stoel Rives LLP, and
Contributing Editor, TheCorporateCounsel.net

Table of Contents

1. Core Competencies 29-2

2. Understanding AI Jargon 29-3

- Core Terms You May Already Know29-3
- Terms That May Be Unfamiliar 29-4
- Technical Term You May Hear from IT29-7

3. Legal Ethics 29-7

4. AI Tools & Use Cases 29-9

- Current Legal Applications of AI..... 29-9
- Practice-Specific Examples29-10
- Choose Your Tool Strategically29-11
- Practical Considerations 29-11

5. AI Use Policies..... 29-12

6. Contractual Restrictions on Use..... 29-13

7. Request Process Workflow29-14

8. The Basics of Prompting 29-15

- Structure29-15
- Natural Language29-15
- Markdown..... 29-16

9. Additional Tips for Effective Prompting 29-18

- General Best Practices..... 29-18
- Prompting Options29-19
- Comparison & Redlining Guidance 29-20
- Methods of Testing & Validation..... 29-20

10. AI-Generated Content Checklist 29-21

11. Keep Learning 29-22

1. Core Competencies

Artificial intelligence (“AI”) is rapidly transforming the practice of law. You have likely heard the quote from Ginni Rometty, Former CEO of IBM:

AI will not replace humans, but those who use AI will replace those who don’t.

Recent data underscores this trend in the legal field. According to surveys presented by Anthropic Associate General Counsel Mark Pike at the November 2025 Applied AI Bootcamp for Lawyers, 77% of lawyers reported using generative AI (“GenAI”) in the last 12 months, with the highest adoption rates among mid-career practitioners.

For those concerned about AI displacing lawyers in the near term, note that leading providers of general-purpose AI models (often used by non-lawyers) have recently updated their usage policies on professional advice. OpenAI recently prohibited using ChatGPT models for “the provision of tailored advice that requires a license, such as legal or medical advice, without appropriate involvement by a licensed professional.” Anthropic made similar changes to its Usage Policy in August 2025, which now requires that users implement additional safety measures (including human-in-the-loop) for “use cases related to legal interpretation, legal guidance, or decisions with legal implications.” See Alexandra P. Moylan and Alisa L. Chestler, *OpenAI Updates Usage Policies: Key Considerations and Next Steps for Organizations Deploying AI*, Baker Donelson Publication (November 18, 2025).

For lawyers looking to get started, the learning curve remains steep. Hallucinations continue to create risk, and issues related to bias, confidentiality, privacy, intellectual property and privilege persist. Not to mention that AI’s capabilities are constantly and rapidly evolving. Without training, early attempts to use AI may reinforce your skepticism about its value. As our colleague Meaghan Nelson observed on [“The Mentor Blog”](#) on TheCorporateCounsel.net:

[A] barrier to entry on the use of AI could be trepidation about “garbage in, garbage out” and not feeling confident that your prompts will produce meaningful results.

And, in fact, all the risks associated with AI use are exacerbated when it is used by lawyers who lack AI literacy. A Fisher Phillips alert identifies four core competencies that define AI literacy:

- **Awareness:** Understanding AI’s capabilities and limitations, including risks such as bias, hallucination and privacy concerns.
- **Application:** Leveraging AI tools to draft content, analyze data and streamline workflows—while maintaining appropriate human oversight.
- **Adaptability:** Remaining curious, experimenting responsibly and continuing to learn as technology evolves.
- **Accountability:** Recognizing when to question, verify or override AI-generated output.

See Fisher Phillips Insight, *What AI Skills Should Hiring Employers Look For? How to Define and Seek Out Workers with “AI Literacy”* (November 17, 2025).

This *Handbook* offers practical guidance drawn from leading voices in the field, particularly insights shared during the panel “AI & The In-House Legal Department” at Northwestern Pritzker Law’s Corporate Counsel Institute in June 2025¹ and the Applied AI Bootcamp for Lawyers hosted by UC Berkeley Center for Law and Business in November 2025.² We hope this *Handbook* will strengthen your awareness and accountability, encourage you to find practical applications for AI in your practice and inspire adaptability so your AI use evolves as these tools evolve.

2. Understanding AI Jargon

One of the initial hurdles in adopting AI is the unfamiliar terminology and technical jargon you may hear in training sessions, from early AI adopting colleagues or from your IT department. Below is a curated list of commonly used terms, explained in plain language for legal professionals.

Core Terms You May Already Know

Perkins Coie has developed a reference guide defining AI concepts in plain language and linking each concept to legal workflows, risks and best practices. See Perkins Coie AI Glossary, available at <https://perkinscoie.com/ai-week-glossary>. Here are some terms from that guide that are probably already familiar:

- **Generative AI (“GenAI”)**: “Generative AI is a class of AI models that can create new, original content rather than just analyzing or classifying existing data. This content can include text, images, code, and more.”
- **Large Language Models (“LLMs”)**: “A Large Language Model is an expansive NLP [Natural Language Processing] model built using deep learning techniques that is trained on vast amounts of text data to understand and generate human-like language. LLMs are the foundation for most modern generative AI tools, enabling them to perform a wide range of tasks like drafting, summarization, and question-answering.”

¹ This panel featured Christina Roupas, Cooley LLP, Chicago; Aaron Gin, McDonnell Boehnen Hulbert & Berghoff LLP, Chicago; Amanda Stuart, Assistant Director of Practice Innovation – Americas, Mayer Brown LLP, New York City; and Jonathan Yellin, Executive Vice President & General Counsel, Charles River Associates, Boston.

² This program featured Jasmine Singh, General Counsel, Ironclad; Cecilia Ziniti, CEO & Co-Founder, GC AI; Loni Mahanta, Chief Legal & Corporate Affairs Officer, HopSkipDrive; Lindsay Smith, Solutions & Innovation Attorney, GC AI; Gloria Lee, Chief Legal Officer, Everlaw; and Javed Qadrud-Din, Director Machine Learning & Applied AI, Thompson Reuters, and Mark Pike, Associate General Counsel, Anthropic.

- **Hallucination:** “A hallucination is an AI-generated output that is factually incorrect, nonsensical, or entirely fabricated, yet presented with the same confidence as a factual statement.”
 - Why Hallucinations Occur: LLMs predict the next most likely word based on patterns in text, not certainties. This probabilistic approach can produce accurate answers or misleading ones. At the Applied AI Bootcamp for Lawyers, Javed Qadnud-Din, Director Machine Learning & Applied AI, Thompson Reuters, illustrated this with two examples:
 - *Flightless birds living in the Antarctic are _____.*
 - Easy for GenAI to answer correctly
 - *Sue went to the store to buy some _____.*
 - Highly ambiguous; GenAI is unlikely to guess correctly and may guess differently each time
 - GenAI models rarely admit uncertainty unless you specifically tell it not to speculate—they will guess rather than decline to answer. Prompts that imply a desired outcome increase the risk of hallucination because the model tries to satisfy the request. Neutral, factual prompts are more likely to yield neutral, factual results.
- **Prompt Engineering:** “Prompt engineering is the art and science of designing effective inputs (prompts) to guide a generative AI model toward a desired output. A well-crafted prompt is clear, specific and provides sufficient context, constraints and format parameters. It is the primary method users have at their disposal to control the model’s behavior, and it significantly impacts the quality of the response.”

For more on basic terms, *see* Charu A. Chandrasekhar, Avi Gesser, Matt Kelly, Diane Bernabei, Carl Lasker, Achutha Raman, Adam Shankman and Nicholas T. Ziebell, *What Exactly is an “AI System” or an “AI Solution?” Debevoise’s Practical Definitions for Common AI Terms*, Debevoise Data Blog (November 23, 2025).

Terms That May Be Unfamiliar

Here are several important terms from the Perkins Coie guide that may be less familiar:

- **Fine-Tuning:** “Fine-tuning is the process of taking a pretrained AI model and training it further on a smaller, specialized dataset. This adapts the model to excel at a specific task or to adopt a particular style or domain knowledge, such as aligning an AI model so that it reasons and writes more like a lawyer. It is a much more complex and resource-intensive process than simply connecting the model to a document database (*see*

Grounding) or providing a system prompt that always runs in the background. Fine-tuning actually modifies the core model itself.”

- **Grounding**: “Grounding is the concept of anchoring an AI model’s responses in specific, verifiable sources of truth to avoid hallucinations and improve the quality of responses. The goal is to prevent the model from relying solely on the generalized, and sometimes incorrect or incomplete, information it learned during training, and instead force it to base its answers on a provided, authoritative context. Grounding can happen in a variety of ways. The simplest is by uploading a document during your chat session and asking questions about the specific file. Alternatively, specialized tools are often connected to databases of information that can be searched before responding (*see* Retrieval-Augmented Generation). Still other tools have ‘web search’ modes, which allow them to search the internet and parse the results before generating a response.”
- **Inference**: “Inference is the process of using a trained AI model to make a prediction or generate an output based on new, previously unseen input. It’s the ‘live’ or ‘runtime’ phase, as opposed to the ‘training’ phase. Every time you submit a prompt to a generative AI tool, you are running an inference task. The speed and efficiency of inference (*see* Latency & Throughput) are critical for a good user experience.”
- **Latency & Throughput**: “These two terms describe the performance of an AI system. Latency is the time delay between when you send a prompt and when you start receiving a response (user-visible speed). Throughput is the total number of requests the system can handle simultaneously (system capacity). Balancing low latency for a responsive user experience with high throughput to accommodate many users is a key challenge in deploying AI.”
- **Multimodality**: “Multimodality refers to the ability of an AI model to process and understand information from multiple types, or ‘modes,’ of data simultaneously. This could include text, images, audio, and video. A multimodal AI model can, for example, look at a diagram (image) and answer a question about it (text), providing a more holistic understanding of information.”
- **Retrieval-Augmented Generation (“RAG”)**: “RAG is a technique to improve LLM accuracy by connecting the LLM to external knowledge sources. When you ask a question, the RAG system first searches the database for relevant information. It then provides this retrieved information to the LLM as part of the prompt, instructing the LLM to use these sources to formulate its answer. This grounds the response in verifiable facts and significantly reduces (but does not eliminate) hallucinations.”

Here are a few additional terms that you may or may not already know from the blog "30 Essential AI Terms Every Attorney Should Know in 2025" by Jeff Kerr, of Casefleet.com (August 21, 2025):

- **Bias/Algorithmic Fairness:** "The tendency of AI systems to produce discriminatory results based on skewed training data. Critical for compliance with anti-discrimination laws, especially in hiring, criminal justice applications, or client screening."
- **Chain of Thought/Reasoning:** "The AI's step-by-step explanation of how it reached a conclusion, not just the final answer. Essential when you need to defend the AI's analytical process in court or to clients. Always ask vendors if their tools provide this."
- **Chunking:** "Breaking a long document (e-mail thread, contract, transcript) into smaller segments so the model can process it without hitting its size limits, and so that each chunk stays on point for better answers. Like organizing exhibits by topic rather than presenting one massive binder."
- **Citations:** "Metadata or inline links that show *where* the model found each fact. Non-negotiable for legal work—always insist on citation capabilities when evaluating AI tools. Without citations, you're flying blind on verification."
- **Context Window:** "The maximum 'working memory' of a model—the total number of tokens it can consider at once. Affects how much of a case file you can analyze in one session. If you exceed the window, older text falls out of scope."
- **Embedding:** "A way AI represents text as numbers to understand meaning and similarity. Embeddings power features like 'search across all case files.'"
- **Guardrails:** "Policy or technical limits placed on a model—like blocking certain topics, requiring citations, or preventing privilege waiver—to reduce risk. Always ask vendors about their guardrails and testing procedures. Think firm policies that keep junior associates from sending unvetted client emails."
- **Model:** "The underlying AI program (with its parameters or 'weights') that turns inputs into outputs—the 'brain' behind any AI tool you're evaluating. Different models have different strengths, biases, and costs."
- **Model Cards:** "Documentation describing what a model was trained on, intended uses, limitations, and ethical considerations. Essential for compliance checklists when evaluating AI vendors. If a vendor can't provide this, proceed with caution."
- **Rate Limiting:** "Caps on how many requests you can send per minute or day—affects both pricing and workflow planning."
- **Temperature:** "Controls how 'creative' vs. consistent the model's outputs are (typically on a 0-1 scale, where 0 = most consistent). Low temperature for contract clauses that

need uniformity; higher for brainstorming litigation strategies. For most legal work, keep it low.”

- **Tool Use:** “Letting the model interact with external systems—databases, calendars, e-discovery platforms—to fetch data or perform actions. Critical for AI tools that need to work with your existing tech stack.”
- **Vectorization:** “Turning text into numerical embeddings so computers can measure semantic similarity. Enables searching for all documents about ‘breach’ even if they use terms like ‘default’ or ‘violation.’ Foundation of modern legal search.”

Technical Term You May Hear from IT

- **Model Context Protocol (“MCP”):** Anthropic announced in November 2024 that it was open-sourcing the “Model Context Protocol,” a “standard for connecting AI assistants to the systems where data lives, including content repositories, business tools, and development environments.” MCP could allow, for example, Claude to communicate with Slack.

3. Legal Ethics

All lawyers employing AI in the practice of law should be familiar with the American Bar Association Formal Opinion 512 (2024), which provides guidance on the ethical use of generative AI tools by legal professionals. Referencing the ABA Model Rules of Professional Conduct, the opinion addresses the duties of lawyers to “provide competent legal representation, to protect client information, to communicate with clients, to supervise their employees and agents, to advance only meritorious claims and contentions, to ensure candor toward the tribunal, and to charge reasonable fees” and discusses those responsibilities in the context of using AI in the practice of law.

- **Competence:** Lawyers must have a “reasonable understanding of the capabilities and limitations” of any AI technology the lawyer uses in the practice of law.
- **Confidentiality:** Lawyers must make “reasonable efforts” to prevent inadvertent disclosure of information relating to the representation of clients and must evaluate the risks of inadvertent disclosure before inputting information into an AI tool.
- **Communications:** A lawyer must “reasonably consult” with a client about the means by which the client’s objectives are to be accomplished, which may require lawyers to disclose their AI use to clients or obtain their informed consent.
- **Meritorious Claims and Candor:** Lawyers must carefully review any AI output to ensure that any assertions made to a court do not cite nonexistent opinions, inaccurately analyze authority or use misleading arguments.

- **Supervision:** Lawyers charged with managerial and supervisory responsibilities must establish policies regarding permissible use of AI and ensure subordinate lawyers and nonlawyers are trained on ethical and practical use of AI and risks associated with AI use.
- **Fees:** If a lawyer uses AI to draft a pleading, the lawyer may charge for the time spent to input relevant information and for the time spent reviewing and revising the draft. Lawyers generally cannot charge a client for learning how to use AI tools.

These examples of how AI use implicates the ethical duties of lawyers are not exhaustive. For example:

- The duty to provide competent representation also means that a lawyer should not use AI for a task if the lawyer lacks the skills to validate the AI's output.
- The use of AI tools can also raise conflicts of interest. As noted in footnote 4 of the opinion, the use of information developed in the representation of one client to inform the representation of another client could run afoul of Model Rules 1.7 and 1.9. Safeguards may be necessary to prevent the sharing of information, even among attorneys within the same firm (e.g., to ensure ethical walls are observed). For this reason, some firms require an individual AI account for each attorney rather than group access to an AI tool.

Attorneys should also be familiar with any orders and guidance issued by bar ethics committees and applicable to the jurisdiction(s) in which they are licensed to practice law or by federal and state courts in which they appear. For example:

- The California State Bar Board of Trustees approved the publication of *Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law* in November 2023. The guidance was developed by the Committee on Professional Responsibility and Conduct and focuses on the ethical obligations of California attorneys when using AI. The "Ethics and Technology Resources" page of the State Bar of California's website notes that this is intended to be a "living document that is periodically updated as the technology evolves and matures, and as new issues are presented," so California licensed attorneys should check that page for updates. In addition to the topics addressed in ABA Formal Opinion 512, this guidance addresses:
 - **Duty to Comply with the Law (Bus. & Prof. Code, § 6068(a); Rule 8.4; Rule 1.2.1):** "A lawyer must comply with the law and cannot counsel a client to engage, or assist a client in conduct that the lawyer knows is a violation of any law, rule, or ruling of a tribunal when using generative AI tools.

"There are many relevant and applicable legal issues surrounding generative AI, including but not limited to compliance with AI-specific laws, privacy laws, cross-border data transfer laws, intellectual property laws, and cybersecurity concerns. A lawyer should analyze the relevant laws and regulations applicable to the attorney or the client."

- **Prohibition on Discrimination, Harassment and Retaliation (Rule 8.4.1):**
“Some generative AI is trained on biased information, and a lawyer should be aware of possible biases and the risks they may create when using generative AI (e.g., to screen potential clients or employees).

“Lawyers should engage in continuous learning about AI biases and their implications in legal practice, and firms should establish policies and mechanisms to identify, report, and address potential AI biases.”

Lawyers (or their clients) employing AI in the practice of law also need to understand how the use of GenAI tools impacts attorney-client privilege and work product doctrine. For example, the U.S. District Court for the Southern District of New York determined in a February 2026 bench ruling that documents outlining a defense strategy and possible legal arguments created by a client using GenAI (on his own accord, and not directed by counsel) then sent to his lawyer for review were not protected by privilege or work product doctrine. The written opinion has yet to be issued, but the court’s decision turned on the fact that the GenAI tool was not confidential and the documents were not prepared by or at the direction of counsel. *See* Charu A. Chandrasekhar, Avi Gesser and Caroline H. Wallace, *SDNY Rules AI-Generated Documents Are Not Protected by Privilege*, Debevoise Publication (February 12, 2026) (describing the bench ruling in *United States v. Heppner*, No. 25-cr-00503-JSR (S.D.N.Y. Feb. 6, 2026), Dkt. No. 22).

4. AI Tools & Use Cases

Current Legal Applications of AI

According to surveys presented by Anthropic Associate General Counsel Mark Pike at the November 2025 Applied AI Bootcamp for Lawyers, the most common legal use cases for AI include:

- Legal research (case law searches, jurisdiction comparisons, precedent hunting)
- Document Work (drafting memos, email templates, first-draft contracts)
- Summarization (case law digests, deposition summaries)
- Contract review
- eDiscovery (document review, pattern detection, privilege flagging)

Conversely, lawyers **are NOT** (yet) using AI for:

- Advanced negotiation
- Securities filings
- High-stakes litigation

Practice-Specific Examples

Mark Pike also shared some examples by practice area, which we have supplemented with examples from other speakers at the Berkeley Applied AI Bootcamp for Lawyers:

- **Corporate**
 - Finding precedent, understanding market terms and benchmarking
 - Data extraction across large data sets (e.g., due diligence)
 - Redline analysis and comparison
 - Clause summarizing and explaining
 - First draft generation
- **Legal Research**
 - Summarizing laws and regulations
- **Intellectual Property**
 - Patent application summaries
 - Identifying patentable innovations
- **Litigation, Investigations & Support**
 - Employment legal needs
 - Analyzing large sets of data (e.g., eDiscovery)
 - Deposition analysis and expert preparation
 - Regulatory investigation responses
 - Customer support improvements
- **Privacy**
 - Privacy Impact Assessments ("PIAs")
 - Data Protection Impact Assessments ("DPIAs")
- **Regulatory, Compliance**
 - RFPs for compliance projects
 - Governance documentation
 - Marketing reviews

On LinkedIn, Todd Carpenter, Chief Legal Officer of Canva, announced in late 2025 that their legal team has gone “AI native” and shared their AI-Native Legal Team Charter, their “blueprint for how humans and AI can work together.” Mary Anne Becking, Head of Corporate Legal at Canva, also highlighted a recent use case in a follow-up post:

Canva’s corporate legal, equity ops and legal ops teams partnered to build and train a Tender-bot that helped answer hundreds of questions from employees and investors during our recent tender offer.

Choose Your Tool Strategically

Different AI tools have distinct strengths and limitations. Selecting the right tool for your task is critical. You might use or encounter any of the following broad categories of AI models used for legal work:

- General-purpose AI models (*e.g.*, Anthropic’s Claude, OpenAI’s ChatGPT, Microsoft’s CoPilot and Google’s Gemini);
- AI models built for legal research (*e.g.*, Thompson Reuters’ CoCounsel, Harvey AI, Lexis+ AI, GC AI); and
- Task-specific AI tools (*e.g.*, Ironclad, Juro, Diligen, Clio, Everlaw, DragonGC, legal plugins for Claude Cowork).

Within each of the above examples, there may be multiple models. For example, as of the time of publishing, within GPT Enterprise, there are three GPT-5 models (Flagship/Router, Thinking and Pro), not to mention several legacy models. Some legal-focused platforms integrate multiple models and may allow you to select your preferred option, while others automatically choose the best fit for your prompt.

Practical Considerations

- Task-specific tools or custom GPTs generally outperform general-purpose models for their intended function but may be less effective for other tasks.
- Because task-specific tools are trained for specific workflows, prompts can often be shorter and more efficient.
- Certain general-purpose models can excel at specific tasks—for example, we have heard that Claude is great for research and ChatGPT is great for organizing information and creating tables. The capabilities of these AI tools are ever-changing, so we will not attempt to suggest any particular tool for a specific use case. For more, *see* Varun Magesh, Faiz Surani, Matthew Dahl, Mirac Suzgun, Christopher D. Manning and Daniel E. Ho, *Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools*, *Journal of Empirical Legal Studies* (April 23, 2025).

- Investing time upfront to identify the best tool for your objective may ultimately save you time iterating on prompts and editing output.
- Always consider your organization's policies, approved tools, data sensitivity and confidentiality requirements, as these considerations may ultimately dictate the tool you can use.

5. AI Use Policies

Corporate AI use policies may address a number of AI-related risks by:

- Requiring human validation of AI output;
- Addressing the protection of confidential client and organizational information; and
- Requiring employees to use company-issued devices and specific, enterprise versions of AI tools that have been formally approved and procured by the organization and are covered by vetted terms and conditions.

Organizations may impose limitations on how information is processed, stored and shared within enterprise AI platforms and can require that employees acknowledge these restrictions each time they access the tool. Those barriers cannot be imposed when employees use personal accounts or access an AI tool on their personal devices.

In addition to addressing privacy, IP, consumer protection and legal compliance matters, a Debevoise & Plimpton update, *Does Your Company Need a ChatGPT Policy? Probably* (February 8, 2023), suggests that a corporate AI use policy may:

- Require that all AI uses be reported to a team to track, risk assess and/or approve;
- Require users to label AI-generated content;
- Require that certain external content be labeled as AI-generated; and
- Require maintaining a record of when certain content was generated and the prompt used.

See the section ["Request Process Workflow"](#) below.

A Sheppard Mullin AI and Privacy Blog says if your company uses AI notetaking or recording tools, your AI policy should also cover: "i) managing and documenting notice and consent; ii) dealing with nonconsenting parties participating in a call being recorded; iii) inaccuracies of AI-generated transcripts and summaries; iv) AI-generated sentiment analysis/emotion detection; v) confidentiality and privilege issues; vi) retention and/or deletion of recordings; vii) vendor diligence on these tools and approval process for specific tools; and viii) knowing the technical features of some tools that can help mitigate risk and others that can create more risk." Sheppard Mullin AI and Privacy Blog, *"Listen Up" if Your AI Policy Does Not Cover AI*

Recording Issues—Another Class Action Lawsuit Filed Over Third-Party AI Recording Service (July 9, 2025).

Finally, to avoid the use of “shadow AI,” organizations should educate their employees on the risks of using an unauthorized AI tool and provide training on how to comply with their AI use policies, in addition to training on how to effectively use the authorized tools.

In developing a policy, start with a discussion with your operations and security colleagues, and specifically ask them about frameworks developed by the National Institute of Standards and Technology (“NIST”) related to AI risk management and cybersecurity and how your organization aligns its practices to these frameworks.

Some law firms have adopted policies requiring attorneys to offer firm clients an opportunity to “opt out” of the firm’s use of AI in matters related to their representation, despite the reluctance of many lawyers to initiate these conversations. And, in fact, contacting clients about AI use might be necessary due to the attorneys’ ethical communication obligations and depending on the terms of existing engagement letters and outside counsel guidelines, which may prohibit using AI with client data without the clients’ consent. *See* the section [“Contractual Restrictions on Use”](#) below.

During a panel discussion on “AI & The In-House Legal Department” at Northwestern Pritzker Law’s Corporate Counsel Institute in June 2025, one panelist from an AmLaw 200 firm shared that most clients request more information regarding how AI is used in their representation and do not ultimately prohibit its use altogether. Generally, clients accept the use of AI to enhance efficiency, provided that their data is not repurposed for other clients. Some clients have even started insisting that AI is used in their representation to save on costs. That preference may even be reflected in outside counsel guidelines, which may encourage or require law firms to “use AI whenever appropriate.” (Be sure guidelines don’t go so far as to require inappropriate use.) However, certain clients have imposed limitations on AI use in specific contexts, and a few have issued blanket directives prohibiting its use entirely. As AI use becomes ubiquitous, we will continue to see policies adapt.

Data retention is also a policy consideration. A Debevoise Data Blog post on “Top 10 Predictions for Law Firm AI Use in 2026” by Charu A. Chandrasekhar, Avi Gesser and Karen Levy (January 13, 2026) suggests that “law firms will struggle with complex data retention choices for chat histories, custom workflows, and AI project artifacts. For joint client-firm AI projects, disputes will arise as to who owns the data and the workflows, as well as any model customization, especially in circumstances when clients want to switch firms.”

6. Contractual Restrictions on Use

A Debevoise Data Blog post on “AI’s Biggest Enterprise Challenge in 2026: Contractual Use Limitations on Data” by Charu A. Chandrasekhar, Avi Gesser, Adam Shankman and Patty (Virtual AI Specialist) (November 17, 2025) predicts that NDAs and other contractual

limitations on content and data use may be the biggest challenge to expanding enterprise AI adoption in 2026.

[M]any of the new features recently announced by frontier LLM providers like OpenAI and Anthropic are designed to provide the models with access to high-quality, non-public internal firm context from work emails, SharePoint sites, databases, customer service calls, and so on. But often, the firms that want to use these materials do not clearly own them or the right to use them in this way. Many of these materials were provided by clients, customers, or other third parties, with use conditions attached to them. Specifically, NDAs, engagement letters, and contractual terms and conditions may place significant limitations on how those documents can be used.

This is yet another issue when considering additional applications of AI tools. The blog notes:

Unlocking the value of AI will increasingly involve the messy exercise of assessing and addressing contractual restrictions on high-quality internal non-public data.

7. Request Process Workflow

Many organizations are implementing structured workflows to request and approve new AI use cases, along with establishing best practices for each approved application. Organizational culture plays a critical role in this process. Employees must feel comfortable proposing the use of a new tool or an existing tool for a different purpose, and they should understand the rationale behind any restrictions or guidelines. Accordingly, the approval process should include clear communication about why certain tools and use cases are permitted while others are not.

A senior public company legal executive described an example of such a process at a high level during the panel "AI & The In-House Legal Department" at Northwestern Pritzker Law's Corporate Counsel Institute in June 2025. It might involve:

- Submitting a formal request and completing a survey detailing the proposed tool or use case;
- Initial automated review by an AI system;
- Subsequent manual review by a designated team;
- Development of any necessary restrictions or best practices; and
- Evaluation of whether a new tool requiring purchase could be leveraged across other parts of the organization.

Both automated and manual assessments should address factors such as security, confidentiality, data privacy, contractual restrictions, regulatory compliance, ethical considerations and bias. The

goal is to ensure that any new AI product does not expose the organization to undue reputational risk and that the risk is proportionate to the anticipated benefits. Importantly, risk levels may vary by use case. For example, applications in recruiting or other HR functions may present heightened concerns.

8. The Basics of Prompting

In general, prompting functions similarly across AI tools and platforms, allowing users to issue requests in natural language. The keys to effective prompting are clarity and specificity, much like delegating tasks to an associate. (Indeed, we've heard that many senior partners often excel at prompting because they are experienced in giving precise instructions to more junior attorneys.) One important distinction between prompting AI and delegating to a junior attorney is that AI systems are highly literal. AI systems are less likely to seek clarification or pose follow-up questions and rarely respond with "I don't know" or "I can't find anything," even when such an answer would be appropriate.

Structure

A well-constructed prompt should include relevant context, clear instructions, the desired output and, where appropriate, the methodology to be applied. Most prompts should include the following components, which have been adapted from the prompting recommendations of Jasmine Singh, General Counsel of Ironclad:

- **Action Item:** Specify the task you want the AI to perform (*e.g.*, summarize, research, analyze, edit/redline, make, change, draft memo, translate to plain English).
- **Context:** Explain who you are, what you aim to accomplish and why. Context may also include the intended audience (*e.g.*, other attorneys, contract managers, sales team); the communication channel (*e.g.*, email, Slack); and the desired tone (formal, conversational, etc.).
- **Detail:** Supply any additional information necessary for the AI to perform effectively. This may include areas of emphasis or de-emphasis, desired length and output format, and limitations on sources (*e.g.*, require primary sources or define acceptable secondary sources). You may also provide precedents or a playbook for reference. If you are providing data, attaching files or referencing a website, clearly identify in your prompt what you are sharing.

Natural Language

Simple, natural language prompts are often sufficient. Here is an example we adapted from remarks by Jasmine Singh, General Counsel of Ironclad, at the Applied AI for Lawyers Bootcamp:

Please redline this agreement to make it more favorable to my company, the purchaser, a [_____] manufacturer. We are procuring [_____] for [_____] from a third-party vendor. The edits will be shared internally with non-lawyers and externally with the counterparty. Use the [_____] language from the attached template as a reference to suggest edits and provide context and reasons for each change. Keep the redlining surgical, avoid wholesale changes where possible while ensuring we achieve the same substantive outcome.

Markdown

In some cases, lawyers may find it helpful to use Markdown to structure more complex prompts. Because AI models interpret text rather than visual formatting, elements such as headings, subheadings, bold or underlining—though useful for human readers—are not recognized by the model. Markdown provides a way to convey structure using plain text symbols. For example:

- Use the # symbol for headings and ## for subheadings
- Use a dash (-) for bullet points

When applying Markdown, you can still include line breaks for readability, but the symbols are essential for the model to understand the intended hierarchy and organization. Markdown does not need to be used to dictate the format of the output (although it can be); it can also serve as a tool to organize complex prompts into a clear outline. Consistency is key. Maintain uniform formatting patterns throughout your prompt.

For a quick reference, consult the “cheat sheet” available at <https://www.markdownguide.org/cheat-sheet/>.

Here are two prompts from the extensive sample prompt library available on GC AI that use markdown formatting to organize these complex prompts. Keep in mind that these prompts may lack some contextual information you would otherwise provide in a prompt because GC AI allows users to save a company profile to avoid the need for repetition in every prompt.

- **Sample Prompt No. 1:**

Legal Topic: [INSERT YOUR TOPIC HERE]

Task

For the above a legal topic, I’d like you to educate me about it with explanations for three different audiences, in decreasing levels of complexity. I’d like you to look up articles and sources on the topic when needed. Below are the audiences. For each audience, write 4-6 sentences in prose.

Three Audiences

1. Start off explaining like I am a senior partner at a law firm who practices in the topic. Use and link to case and/or statutory citations that you find, and in general assume a very high level of sophistication.
2. Then I want you to explain assuming that you are writing for a smart group of in-house lawyers who are great and sophisticated lawyers. They want the legal framework with some legal nuance, plus the practical effect of the law.
3. Then assume you're explaining it for business people at our company who are not lawyers. They care about the law as applied to us. Use no more than a 12th grade reading level.

Follow Up

After you provide the three explanations, ask me if I'd like you follow up or learn more, or perhaps if I'd like you to provide more analysis or help me work through the implications.

- **Sample Prompt No. 2:**

Please review this client alert or article for me.

Client Alert

CLIENT_ALERT_LINK

Task

Please give me:

1. The key takeaways from the article and law or other development it describes.
2. Your assessment of whether this applies to my company.
3. Action items for me or my legal team from it.
4. Ask me for any documents or policies for my company that might need to be updated based on your answer to 3.

Speakers at the Applied AI Bootcamp for Lawyers gave a few simple prompt outlines:

- **Example 1:**

You're acting as _____.

Your task is _____.

Compare against _____.

Your response / output must be _____.

- **Example 2:**

context _____

task _____

compare _____

output _____

- **Example 3:**

##Role _____

##Input _____

##Requirements _____

##Examples _____

9. Additional Tips for Effective Prompting

General Best Practices

- *Avoid overly broad or open-ended tasks.* They often produce vague or incomplete results. Consider whether your request will produce a result that you are able to validate. If not, break down the task into smaller, concrete questions.
- *Avoid suggesting outcomes or showing preference.* If you provide a desired outcome, the model will attempt to deliver it—even if it does not exist—increasing the risk of hallucination. For instance, do not prompt AI to find a case with a specific holding unless you know that case exists. As noted above, neutral, factual prompts are more likely to yield neutral, factual results.
- *Be wary of the context window.* AI models have a maximum token limit that constrains the amount of information they can process in a single interaction. If your prompt exceeds the tool’s limit, the AI may ignore parts of your input or produce incomplete answers. You can ask AI for a lot of information about a discrete thing or a little information about a lot of things. Avoid overload by chunking lengthy documents into smaller segments or lengthy projects into discrete tasks. For example, Lindsay Smith, Solutions Attorney at GC AI, notes that asking AI to conduct a 50-state survey in one prompt is unlikely to produce reliable results; start with one state and build from there.
- *Know whether your AI tool retains memory between sessions.* If it does, understand what information persists. Some tools allow preloading company details to avoid repeating context or allow you to ask the tool to “remember” certain facts or outputs for future reference. If the tool you are using does not retain memory between sessions, you must

provide relevant details—such as your role, company name, industry and examples—each time.

- *Consider output restrictions carefully.* AI is highly literal. If you require a specific length or number of bullet points, the model may omit key information to meet those requirements. If you first request an unrestricted summary, then a condensed version, you can compare the two.
- *Mind your prompt's tone.* Including “please” and “thank you” can improve output quality. Studies suggest polite prompts yield more detailed, better-formatted responses, likely because they provide additional context and mirror human communication patterns. AI may mimic your tone. Professional, courteous prompts often result in professional, courteous output. *See Jodie Cook, Should You Say Please and Thank You to ChatGPT?, Forbes (May 2025).*
- *Expect iteration.* Prompting is often iterative, whether to add details you missed in the first prompt or otherwise refine the output. But repeated iterations can degrade results if the model revisits tasks it previously completed well. In those cases, start a new chat and include the initial response as input.
- *Start a new chat for new tasks or topics.* This resets the context window and avoids confusion from prior prompts.

Prompting Options

- *Assign the AI a role or persona.* For example, instruct it to act as a junior associate at a law firm. GC AI suggests you try instructing AI to adopt a certain role or persona and using it as a sparring partner to anticipate questions (*e.g.*, from an internal audience of non-lawyers) or counterarguments (*e.g.*, from opposing counsel). You may need to clarify the role's scope and limitations and be careful of unrealistic assumptions.
- *Ask AI to draft a prompt or explain your objective and invite it to ask questions.* For example, Cecilia Ziniti, CEO & Co-Founder of GC AI, suggested something like the following:

Please draft a prompt with formatted ## breakers to be used with Gemini for this use case: An in-house lawyer needs to identify deviations in this agreement from my company's standard playbook, assess risk for each deviation (flagged red, yellow, or green), and provide a summary of implications and remediation options.

Or:

I'm an in-house lawyer who needs to identify deviations in this agreement from my company's standard playbook, assess risk for each deviation

(flagged red, yellow, or green), and provide a summary of implications and remediation options. Ask me any questions you need to complete this task.

- *Ask clarifying questions.* For instance, "Do you understand this term?"
- *Leverage sample prompts when available.* Many AI tools provide sample prompts, and some law firms and in-house teams have developed prompts for common tasks. If you create your own prompt for a recurring task and find it effective, save it in your personal prompt library.

Comparison & Redlining Guidance

- *Provide clear instructions for redlining when comparing contracts against precedent or a playbook.* When using an AI tool to redline a contract using a precedent or playbook you provide, without specific instructions, it may make unnecessary grammatical corrections or fully replace each provision with the comparable provision from the precedent. You may want to prevent excessive redlining by:
 - Instructing AI not to edit boilerplate language (if you are using an AI model that will understand this request and what it means to you);
 - Instructing AI not to edit specific sections that you identify;
 - Requesting a surgical redline limited to substantive changes aligned with the precedent or your playbook; or
 - Instructing AI to ignore grammatical issues unless they have substantive implications.
- *Handle defined terms carefully.* General-purpose AI models may not understand the phrase "defined terms" if included in a prompt.
- *Specify if you require a point-by-point comparison.* Without this, AI tools may only compare key terms. For example, you might prompt the AI to compare two documents concept by concept and present the results in a table with a column for each document's approach and a third column explaining the differences.

Methods of Testing & Validation

- *Request citations for legal work.* Prompt the AI to provide references to all cases, statutes, rules, documents, agreement sections or subsections, or webpages, to aid in your review of the AI output. But do not rely solely on these citations; they may be incomplete. Verify the validity, content and accuracy of all citations independently, because AI-generated references may be fabricated or outdated. This includes validating that the referenced source exists, is cited correctly, is appropriately used for the purported purpose and is accurately summarized.

- *Do not ask the AI tool to assess its own confidence or validate its prior output.* It will likely reinforce its previous answer rather than provide an objective evaluation. If you want to use an AI tool as a first level of validation, you would be better served by running a fresh prompt or running the same or similar prompt in another AI tool (or several other AI tools). No matter how many AI tools you use, checking with AI is never a replacement for human review.
- *Test your prompt and the AI tool.* For example, when assessing a policy for compliance with a particular law, add a term you know violates that law to confirm the AI identifies it. But be aware that this only reveals errors, not omissions. It may help you improve your prompt, but it cannot ensure the AI's output is complete.
- *Lawyer-in-the-loop.* Most importantly, AI output should never be treated as final. Manual review, validation and potential correction of any AI output is always necessary before sharing externally.

10. AI-Generated Content Checklist

The ABA Law Practice Division released a “ten-step checklist for reviewing AI-generated legal content before it’s shared with clients, courts, or opposing counsel” in January 2026:

1. Use firm-approved tools

Avoid public AI models for client work. Use secure, legal-specific tools designed for law firms, such as Clio Work, that keep work inside your firm’s existing systems.

2. Confirm security and confidentiality

Verify that the tool guarantees Zero Data Retention and meets recognized security standards like SOC 2 Type II. Never assume privacy. Confirm it.

3. Check for factual accuracy

Manually verify all facts, citations, and assertions. Treat AI output the same way you would an associate’s draft.

4. Cross-check sources

Rely only on verified legal databases such as Clio Library. Confirm that cited cases, statutes, and regulations actually exist and are current.

5. Analyze reasoning quality

Review the logic, not just the conclusion. Confirm proper use of IRAC/CRAC, correct doctrinal tests, and coherent legal reasoning.

6. Confirm the correct jurisdiction

General AI frequently blends jurisdictions. Ensure the content reflects the correct federal, state, provincial, or local law, including terminology and standards.

7. Look for bias or mischaracterization

Check that cases are accurately described and not selectively framed. Generative AI can reflect bias from training data if left unchecked.

8. Verify formatting and procedural rules

If the content is headed to court, confirm formatting requirements, captions, headings, signatures, and local procedural rules.

9. Ensure ethical compliance

Follow ABA Model Rules, state rules of professional responsibility, or Law Society guidance on AI-assisted work.

10. Require final human sign-off

Every piece of AI-generated legal work must be reviewed and approved by a lawyer before release. No exceptions.

Clio, ABA Law Technology Today, *A Practical Checklist for Using AI Responsibly in Your Law Firm* (January 5, 2026).

11. Keep Learning

The considerations outlined in this *Handbook* are not exhaustive. As noted in the Fisher Phillips alert, AI literacy has become a highly valued skill for employers—even in roles that are not technology-focused—and many organizations are actively encouraging adoption among existing employees:

Employers are discovering that “AI-fluent” workers often outperform their peers by leveraging automation to handle routine work, which frees time for strategy, creativity, and relationship-building.

However, both AI capabilities and employer expectations are evolving at an unprecedented pace. Guidance based on today’s tools may be outdated tomorrow. Treat AI literacy like any other core legal competency by committing to continuous learning and professional development. Take advantage of opportunities for training in AI and prompt engineering, particularly programs offered by your firm or organization that are tailored to the tools you use daily or to your specific practice area. Staying informed and adaptable will position you to leverage AI responsibly and effectively in your work.

If AI follows the trajectory of prior technological advancements in the legal profession, adoption will be essential to keep pace with the modern practice of law. As more lawyers integrate AI into their workflows, those who don't risk falling behind.

Good luck!