

Checklist: AI Use by Lawyers — Prompting Basics & Tips

By TheCorporateCounsel.net

Generative artificial intelligence (“GenAI”) 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.

This Checklist offers practical guidance for prompting GenAI tools, drawing on insights from leading voices on legal applications of GenAI, particularly those shared during the Applied AI Bootcamp for Lawyers, hosted by the UC Berkeley Center for Law and Business in November 2025.¹ For more practical tips and considerations, *see* our [“AI Use by Lawyers” Handbook](#), posted on TheCorporateCounsel.net.

1. 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.

a. 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

¹ 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.

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.

b. 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.

c. 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 _____

2. **Additional Tips for Effective Prompting**

a. **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.

b. 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.

c. 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.

d. **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.

3. 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).