

# Generative AI, the law and legal prompt engineering

Published: May 06, 2025



I am quoting a Microsoft publication titled 'Prompt Engineering for Lawyers - Leveraging generative AI in the legal profession'

*<quote>Knowing the basics of prompt engineering – the practice of formulating instructions to achieve specific outcomes from a generative AI tool – will enable lawyers to generate more accurate and pertinent results</quote>*

According to this paper a 'good prompt' will:

*<quote> provide a clear description of the task, explain the role the AI tool needs to play, describe the audience, provide guidance on the tone, style and length of the expected output, and any additional context to be considered. An iterative process will help to refine the result</quote>.*

So, there should be:

- a goal stated,
- context be given,
- expectations be communicated and
- sources be requested and/or cited.

As a lawyer-linguist, at this point, tell me more about your latest findings in regard of any prompt engineering - especially for lawyers (legal prompt engineering).



Recent developments in prompt engineering for lawyers have revealed that the process can be structured into a clear, repeatable methodology—often referred to as:

the four-pillar model:

...

This approach emphasizes a four-step structure:

**Goal:** Explicitly state what you intend to achieve. Whether you're drafting a contract clause, preparing litigation materials, or synthesizing client communications, begin by clearly articulating the desired outcome with precise, action-oriented language.

**Context:** Provide detailed background information that frames the task. This might include the relevant legal field, jurisdiction, targeted audience, and any pertinent case law or client circumstances. By setting the stage properly, you ensure the AI tool understands the environment in which its output will be applied.

**Expectations:** Define the tone, style, format, and depth of the output you expect. For example, in legal drafting, you may require a formal tone with specific formatting rules or a summary that adheres to a particular word limit. Explicit guidance on these points helps generate content that fits seamlessly into your professional workflow.

**Source:** Request that the AI cite sources or reference specific materials when relevant. This adds an extra layer of accountability and compliance, crucial in the legal domain, where accuracy and the ability to back-check information are essential.

These findings stem from practical tests and applications within legal settings—Microsoft's internal experiments, for instance, have demonstrated that adopting such structured prompts can make lawyers up to 32% faster in certain tasks while increasing accuracy by around 20%<sup>2</sup>.

Another important insight is the iterative nature of prompt engineering. Prompts rarely hit the mark on the first try. Lawyers are encouraged to view prompt engineering as a dynamic process where initial drafts are continuously refined. By interacting with the AI in a conversational manner—adjusting your prompt based on the output—you gradually hone results that fully meet your legal standards. This iterative process is key to ensuring that each generation improves in clarity and relevance, aligning closely with established best practices in legal research and document drafting.

Beyond speed and efficiency, prompt engineering also plays a vital role in risk management and compliance. Legal professionals must remain vigilant about safeguarding confidentiality and ensuring that any external or sensitive data is handled appropriately. When a prompt directs the AI to integrate or comment on specific legal data, it should always include safeguards such as referencing only publicly available texts or redacted versions of internal documents.

This careful balance between leveraging advanced technology and protecting client confidentiality is crucial in today's competitive legal environment.

As generative AI continues to evolve, the guidelines for effective prompt engineering are also adapting. Lawyers who integrate these practices not only achieve better initial outputs but also enhance overall workflow efficiency, allowing them to focus more on high-value, strategic tasks. In essence, treating generative AI as a true "copilot" rather than an automatic solution underscores the importance of human oversight in ensuring that technology augments rather than replaces professional judgment.

Looking ahead, beyond the structured four-pillar approach, newer techniques involve incorporating "few-shot" examples within prompts to teach the AI about specific nuances of a legal issue or drafting style. These examples act as templates to further fine-tune responses, making AI outputs more aligned with the lawyer's unique persona and the expectations of the legal audience.

Would you be interested in exploring specific case studies or sample prompts demonstrating these principles in contract drafting or litigation support?



Not for the moment, because in the paper I cited, you can find some sample prompts demonstrating these principles with regard to contract drafting, litigation support (dispute resolution), mergers and acquisitions and regulatory compliance.



That makes sense! Since you have access to those sample prompts, you already have a strong foundation for applying the principles of prompt engineering in legal work.

If you ever want to explore more advanced techniques—like refining prompts for deeper analysis, structuring multi-turn interactions for complex cases, or integrating legal precedent citation into AI-generated drafts—I'd be happy to discuss those. Let me know whenever you're ready to dive in!