

Prompt Better For Lawyers: Mastering the Art of Asking the Right Questions in the Age of AI



September 2026



"Ask the right questions if you're to find the right answers."

— Vanessa Redgrave —

Foreword

The legal profession has always been about asking the right questions. Whether in a client interview, a cross-examination or a research memo, the quality of the answer depends on the quality of the question. In the age of generative AI, that principle has never been more important.

AI chatbots are extraordinarily capable, but they are not mind-readers. They do not "understand" the law in the way a seasoned solicitor does. They predict the next word based on patterns in their training data. Give them a vague instruction and you will receive a vague, often generic, and sometimes dangerously inaccurate response. Give them a well-crafted prompt and you can obtain targeted, useful and surprisingly sophisticated output that can save hours of work.

This article is a practical guide to prompting better. It weaves together the foundational principles of effective prompting with a collection of ready-to-use prompts that have proven their worth in daily legal practice. It is written for practising lawyers; particularly those in smaller firms or sole practices who do not have access to expensive specialised legal AI platforms. The techniques described here work with the widely available generic chatbots that most of us already use: ChatGPT, Claude, Gemini, Kimi, GLM, DeepSeek and others.

The goal is simple: to help you get better results, faster, while remaining in control of the process and fully aware of the risks. The first half of this article explains the techniques (the “how”). The second half gives you the actual recipes (the “what”) you can copy, paste and adapt tomorrow morning.

Why Prompting Matters More Than the Model

Many lawyers spend too much time worrying about which AI model is “best.” In reality, the difference between a mediocre and an excellent result often lies not in the model but in the prompt.

A well-structured prompt can turn a general-purpose chatbot into a surprisingly effective legal research assistant, drafting partner or document reviewer. A poorly structured prompt will produce the kind of generic, robotic output that makes clients (and judges) roll their eyes.

Think of prompting as the legal equivalent of instructing a very bright but inexperienced trainee. You would never say to a trainee “draft a contract” and expect a usable first draft. You would give background, specify the parties, list the key commercial points, indicate the governing law and describe the tone you want. Exactly the same discipline is required with AI.

The good news is that AI is an infinitely patient trainee. You can iterate, refine and even ask it to improve its own prompts. Once you master this skill, you will find yourself using AI far more often and far more effectively. The lawyers who thrive will not necessarily be those with the most expensive specialised platforms. They will be those who have mastered the art of giving clear, precise, context-rich instructions to the tools that are already freely or cheaply available.

The Core Principles of Effective Prompting

There are several fundamental principles that apply to almost every legal prompt. Master these and everything else becomes easier.

Clarity: Be specific about what you want. Include jurisdiction, time frame, specific legal concepts and the precise outcome you seek. Vague prompts produce vague answers.

Context: Provide relevant background information so the AI understands why you are asking. Context dramatically improves relevance.

Format: Specify how you want the response structured — for example, as a table, bullet points, a formal memorandum or slides.

Examples and Constraints: Give examples of the style or content you want, or set limitations such as word count, excluded topics or required elements.

Iterations: Start broad, then refine with follow-up prompts. Unlike a web search, an AI conversation retains context within the same thread.

Industry-specific terminology: Use accurate legal language. Models trained on legal texts respond far better to proper terms than to colloquial ones.

These principles sound obvious, yet they are routinely ignored. The internet is full of generic “prompt libraries” that produce equally generic results. Lawyers need something more precise.

From Vague to Useful: Three Illustrations

Let us look at how a modest amount of extra information transforms a prompt. These examples come from real legal work.

Example 1: Negligence Research

Vague: "Find cases about negligence."

This will produce an unmanageable flood of results from every common-law jurisdiction.

Improved: "Find cases in Hong Kong regarding negligence in slip-and-fall accidents involving wet markets in the past 5 years."

The improved version specifies jurisdiction, type of accident, factual context, and time frame. The AI now has a fighting chance of giving you something usable.

Example 2: Statutes and Ordinances

Vague: "Tell me about privacy laws."

Improved: "Summarise the key provisions of the Hong Kong Personal Data (Privacy) Ordinance regarding data breaches, including notification requirements and potential penalties."

Again, the extra specificity (the exact ordinance, the particular topic, and the elements you care about) makes all the difference.


Example 3: Finding Precedents

Vague: "Does someone who lends money to a corporation in Singapore and takes a guarantee fall foul of the Moneylenders Act 2008?"

Improved: "Find cases in Singapore that address the meaning of 'excluded moneylender' under Section 2(e)(iii) of the Moneylenders Act 2008. Specifically, if someone lends to a corporation and takes a personal guarantee as security, is that person still an 'excluded moneylender'? Highlight any established tests or criteria used by the courts."

The improved prompt tells the AI exactly which section to focus on and what factual scenario you are interested in.

Notice that none of these improved prompts is particularly long. You do not need to write an essay. You simply need to think like a lawyer instructing a junior.



A Prompting Strategy That Works

Unlike a traditional Google search, an AI conversation has memory within the same thread. You can therefore start broad and then refine.

A good sequence is:

1. Ask for an overview or list of key issues.
2. Follow up with more specific questions based on the first response.
3. Ask the AI to cite sources or give examples.
4. Ask it to identify weaknesses or counter-arguments.

This iterative approach is far more effective than trying to write the perfect prompt on the first attempt.

One of the most under-used techniques is simply asking the AI to help you write a better prompt. Try this:

"I am researching [issue]. I wish to use [name of chatbot] to find information concerning [subject]. Please give examples of prompts I should use and explain what additional information or context I should include."

The AI is often surprisingly good at suggesting improvements to its own instructions. This is the starting point for everything that follows.

Advanced Prompt Engineering Techniques for Lawyers

Once you have mastered the basics, a range of more sophisticated techniques can take your results to another level. These methods allow you to handle complexity, reduce errors, and tailor output with precision. They are drawn from practical experience with legal work and from systematic testing with the major chatbots.

The Iterative Technique

This is the fundamental approach for any complex legal task. You begin with a general prompt that defines the problem, evaluate the response for clarity and accuracy, refine the prompt to address gaps or inaccuracies, and re-run it. Repeat until the output meets your professional standard.

This mirrors how we already work: analyse, test, and optimise. For example, when drafting a complex clause, start with a high-level request, review the first version, then add constraints such as "make it more concise while preserving the indemnity protection" or "align the language with Hong Kong law". The AI improves with each cycle. In practice, three to five iterations usually produce a usable first draft that would have taken a junior lawyer half a day.

Chain-of-Thought Technique

For abstract or multi-layered legal questions, instruct the AI to break the problem into a series of logical steps, with each step building on the previous one. This improves accuracy, transparency, and explainability.

Example: "Analyse whether the following facts give rise to a claim in unjust enrichment under English law. First, identify the enrichment. Second, determine whether it was at the claimant's expense. Third, consider whether it was unjust. Fourth, examine any defences. Finally, synthesise your conclusions into a single memorandum."

The AI is forced to show its working, which makes it easier for you to spot errors. This technique is especially powerful for opinion work, where you need to see the reasoning rather than just the conclusion.

Self-Criticism Technique

Ask the AI to evaluate and critique its own output. Add phrases such as "Please critique your own analysis", "Evaluate the potential limitations of your proposal" or "Identify any assumptions or biases that may have influenced this advice".

This is particularly useful for legal research or opinion drafting. The model is forced to consider its own limitations, resulting in more refined and cautious responses. You can then decide which criticisms to accept. In my experience, this single extra sentence often surfaces the very weaknesses a supervising partner would have spotted.

Positive and Negative Technique

Give the AI both positive examples (what you want) and negative examples (what you do not want). This helps it understand the boundaries of the task.

Example: "Draft a non-compete clause. Positive example: a clause that is reasonable in duration (12 months) and geography (Hong Kong only). Negative example: a clause that is overly broad and likely unenforceable. Use clear, commercially sensible language."

The contrast sharpens the output considerably. You can apply the same method to almost any drafting task: "This is a good example of a clear force-majeure clause. This is a bad example because it is too vague."

Task Decomposition Technique


When a task is too large or ambiguous, break it into smaller sub-tasks. Prompt the AI on each sub-task separately, then ask it to consolidate the results.

This is ideal for due diligence, complex contract reviews, or multi-jurisdictional research. Write down the overall prompt, split it into focused sub-prompts (the AI can even help you do the splitting), run each one and finally instruct: "Now combine all the above into a single coherent report." The result is far more reliable than asking for everything at once.

Temperature and Determinism Control

Temperature is a parameter that controls how random or conservative the AI's output will be. Low temperature (close to 0) produces focused, predictable, "safe" responses — perfect for factual legal research or precise drafting. High temperature (1.0 or above) produces more creative, varied output — useful for brainstorming arguments or generating marketing copy.

Most chatbots allow you to set this, either explicitly or through phrasing such as "be conservative and stick closely to established law" versus "be creative and suggest novel approaches". For most legal work, keep the temperature low.

A decorative graphic in the bottom right corner consisting of several overlapping, light-colored circular arcs and a trail of small dots, creating a sense of motion or a stylized path.

Persona Layering

A persona is a fictional character with defined expertise, tone, and audience. Embedding a persona in the prompt aligns the AI's responses with the characteristics you need.

Example: "Voon Keat LAI is a senior corporate lawyer in Hong Kong with expertise in mergers and acquisitions, capital markets and funds. His language is specific and concise. Use this persona when explaining the following assignment clause..."

You can create different personas for different audiences — a general counsel, a board of directors or even a 12-year-old student whose second language is English. Some platforms do not remember personas between conversations, so restate them at the start of each thread. Some aggregators allow you to save custom bots with a permanent persona.

Other useful techniques include specifying output formats (tables, slides, or numbered lists), setting strict word or character limits, requesting particular writing styles, defining how the AI should handle uncertainty ("if you are less than 80% confident, say so"), instructing it to exclude certain topics or approaches and using prompts to construct better prompts. You can also ask the AI to verify its own understanding before tackling a complex task or to perform progressive knowledge checks.

These techniques are not theoretical. They are the difference between an AI that produces something you can use and an AI that produces something you have to throw away.

Great Prompts for Everyday Legal Work

Here is a collection of prompts that have proven useful in practice. Each starts with a basic version and then shows how a few extra details — often using the techniques above — produce a much better result. Copy them, adapt them, and make them your own.

Case Review

Basic: "Please summarise the attached case. What are the legal principles established in the case?"

Improved: Add "What were the reasons given by the judge in arriving at the decision? What case law did the judge rely on? Distinguish between ratio and obiter if possible."

AI still struggles to distinguish ratio from obiter with complete reliability, but a well-directed prompt can dramatically reduce the time a knowledgeable lawyer needs to extract the key points.

Client Communications

Basic: "Please draft a [document] to a client dealing with [purpose]."

Improved: Add background, specific issues to cover (bullet points work well), desired tone, audience, and length.

Example: "Please draft a legal memorandum to a client explaining the fiduciary duties of directors under the English Companies Act and illustrate the point by way of examples and case law. The tone of the letter should be friendly. The style should be professional and legalistic. The client is involved in motor trading. The audience is the general counsel of the company. Keep it under 1,200 words."

Document Drafting

Generic AI still cannot produce a sophisticated 80-page SPA from a single prompt. It can, however, produce a very usable first draft of a simpler document, or help you build a complex document block by block using task decomposition.

Example: "Please draft a sale and purchase agreement for the purchase of a second-hand non-motorised sailing boat. The boat is moored in Hong Kong. The price is HK\$1,000,000. The Seller is an individual and the Buyer is a corporation. Both are based in Hong Kong. Specifically, the agreement should cover all the standard provisions of a SPA relating to the sale of a boat including the relevant warranties, closing adjustments, conditions precedent, how delivery should be made, post-completion obligations, breaches and remedies, when the risk of loss or damage passes, any registration requirement and the standard boilerplate provisions. The SPA is governed by Hong Kong law. Use clear, commercially sensible language and keep each clause as concise as possible."

Always treat the output as a starting point that requires thorough review by an experienced lawyer.

Document Review

When reviewing a lengthy agreement, give the AI a clear list of issues you care about. Use task decomposition if the document is very long.

Example (private equity fund): "The attached is a Limited Partnership Agreement for a USD [amount] fund. Please evaluate if its key terms and provisions conform to typical market standards and practices for Private Equity funds of this size. Our proposed investment is USD [amount]. What other rights should we seek given our investment size? Specifically analyse: management fees, carried interest, key-man provisions, no-fault divorce, reporting requirements, transfer rights and any other noteworthy terms. Flag any provisions that are unusually one-sided."

Legal Research

Always ask for illustrations and citations, then verify those citations independently. Combine with chain-of-thought for complex issues.

Example: "Please give me the principles of estoppel under English law. Identify the tests required to maintain a defence of estoppel. Please cite relevant cases and give me a brief summary of each of these cases. Then critique your own answer for any jurisdictional confusion or outdated authorities."

A second research prompt that works well: "In the context of setting up a Type 1 licensed corporation in Hong Kong, please give the detailed qualification requirements. Further, please estimate how long this might take and what additional information is needed. Then list any common pitfalls."

Due Diligence and Financial Analysis

AI is surprisingly good at spotting patterns in financial statements once you tell it what to look for. Follow up with more targeted questions.

Example: "Based on the attached financial statements of [Company] for FY 2021, 2022 and 2023, conduct a financial due diligence analysis. Identify any red flags, assess the company's financial health and highlight areas that may pose legal or financial risks. [Company] is involved in [industry] and listed on [exchange]. Then critique your own analysis for any assumptions that may not hold."

You can then follow up: "Did you find any patterns or transactions that may indicate financial fraud, such as revenue recognition issues, expense misclassification or off-balance-sheet liabilities?"

Paraphrasing and Rewriting

Lawyers are often terrible at writing clearly. AI can help.

"Rewrite the following clause/article into [XXX] words but keep its legal effect. Please also illustrate with examples how the clause works. Use a professional but accessible tone."

Domain Expertise / Industry Knowledge

"Please explain the difference between European waterfall and American waterfall in the context of funds. Then give a simple numerical example that a non-finance lawyer can understand."

Jurisdictional Confusion

AI frequently mixes up the law of different common-law jurisdictions. A prompt about Hong Kong law may produce a perfectly accurate statement of English or Singaporean law, complete with the wrong section numbers.

Always specify the jurisdiction clearly and then check the citations against the actual ordinance or statute.

Help with Prompts

When you are stuck, ask the AI itself:

"I am researching [issue]. I wish to use [AI chatbot] to find information for me concerning [subject]. Please give examples of prompts to be used and explain what additional information or context I should include in those prompts."

These prompts are not magic. They work because they incorporate the techniques described earlier: specificity, context, format, constraints, iteration, and self-criticism.

Managing the Risks: Hallucination, Bias and Over-Reliance

No discussion of prompting would be complete without a frank acknowledgement of the risks. The current generation of AI models, while powerful, have important limitations.

Hallucination

AI will sometimes invent cases, statutes or quotations that sound completely plausible. This is not malice; it is a statistical artefact of how the models work. They prioritise fluency over accuracy.

Hallucination can be easily managed and mitigated.

Be Specific: Be as specific as possible in your prompts. Always ask for citations and then verify them independently (HKLII, Westlaw, Lexis, official government websites or at least a Google search). Treat any case that "perfectly" fits your facts with particular suspicion until you have confirmed it exists. Use the self-criticism technique: "Please critique your own citations and flag any that you are not 100% certain exist."

Web Searches: Most modern AI platforms now support real-time web search. This capability is especially invaluable for legal work. Almost every meaningful point of law has been written about somewhere online. Whenever possible, instruct your AI to perform a web search to verify the information it provides.

Multi-LLM: A particularly powerful anti-hallucination method is the multi-LLM technique. Even the best single model still has its own blind spots and can produce fabricated citations or incorrect statements of law, especially on complex, multi-turn legal tasks. Legal hallucination rates remain unacceptably high for professional work if you rely on one system alone. The most reliable solution is to use multiple models. Run the same prompt through different LLMs, or (my preferred method) ask one bot to critique the output of another. When the models converge, you gain genuine confidence. When they diverge, they immediately surface the disagreement for you to investigate. This is something no single-model system can offer on its own. In daily legal practice, this multi-model, multi-turn checking drives citation hallucinations down to a very low level. LLM-agnostic platforms and aggregators such as Poe and my own AskKEMP.ai make the workflow straightforward. For high-stakes work, it is the safest way to keep the AI honest.

Bias

Training data contains the biases of the internet. US-trained models and Chinese-trained models can give noticeably different answers on politically sensitive topics. Using several different models can help surface these differences. Be vigilant about biases.

Over-reliance

AI is not a substitute for professional judgment. Cross-check using different bots, apply common sense and never accept everything as gospel. There is no substitute for a lawyer's critical eye.

Other Risks

Plagiarism can be an issue if you present AI-generated text as your own without review.

Claude now reportedly marks its AI-generated content so that someone with the right App can find out if the content is generated by Claude. I don't think there is anything wrong with content generated with the help of AI, but be honest with yourself and your client whether AI assistance is used.

The golden rule remains: AI is a tool to augment your professional judgment, not replace it. You remain fully responsible for the advice you give.

Specialist Legal Bots

If you have access to specialist legal bots, use them.

Specialist legal bots can take the benefits of AI even further than generic models, however well prompted. As they are purpose-built or fine-tuned for legal practice, they bring deeper domain knowledge of particular practice areas, document types and jurisdictions, together with workflows designed around how lawyers actually work. The result is faster, more reliable output on research, drafting, review and analysis, often with built-in awareness of local law and professional standards.

They typically offer enterprise-grade security, data protection and compliance and the ability to handle multi-jurisdictional work.

My own platform, AskKEMP.ai, provides exactly this: a team of specialised AI legal companions providing common law support particularly for lawyers in Hong Kong, the UK, the UAE (DIFC and ADGM), Malaysia and Singapore.

Specialist legal bots are designed not to replace the lawyer but to supercharge the practice — reducing research time dramatically, accelerating document production and firms and practitioners to scale sustainably while remaining fully in control of the advice given.

Conclusion: Prompting Is the New Legal Skill

The lawyers who thrive in the coming years will not necessarily be those who have access to the most expensive specialised AI platforms. They will be those who have mastered the art of prompting the tools that are already freely (or cheaply) available.

Prompting is not a technical skill. It is a legal skill — the skill of giving clear, precise, context-rich instructions. It is the same skill we have always used when instructing juniors, counsel or expert witnesses. We are simply applying it to a new kind of assistant.

The future of law is not AI versus lawyers. It is AI-empowered lawyers. And the lawyers who prompt better will be the ones who stay ahead.

About the author: Voon Keat Lai is the author of "The AI-Empowered Lawyers' Cookbook" and the founder of the AI Legal Platform AskKEMP.ai.

Should you require any further information, please contact:



Voon Keat Lai
Group Managing Partner
M: +852 9300 7047
E: vk lai@kempllp.com

88.007.1-2609040719.1