

LAST PRINTED ISSUE OF *PROBATE & PROPERTY*. WE ARE GOING FULLY DIGITAL. SEE PAGE 5 FOR DETAILS.



LITIGATION
AND THE
JACKSON ESTATE


INSURANCE AND
ANNUITIES IN MODERN
ESTATE PLANNING

INITIAL TAXPAYER EFFECT
OF THE GST FINAL
REGULATIONS

VOL 40, NO 4
JUL/AUG 2026

PROBATE & PROPERTY

A PUBLICATION OF THE AMERICAN BAR ASSOCIATION | REAL PROPERTY, TRUST AND ESTATE LAW SECTION

A large, abstract digital graphic composed of many small dots in red and blue, forming a complex, swirling shape that resembles a stylized 'P' or a digital wave. It occupies the lower half of the cover.


Becoming the
Director of Artificial
Intelligence
Through Prompt
Engineering

All Right, Mr. AI, I'm Ready for My Close-Up

Becoming the Director of
Artificial Intelligence Through
Prompt Engineering

By J. Cliff McKinney II

Getty Images



Every director knows the performance depends on the direction. Artificial intelligence (AI) is no different. It is a remarkably capable performer that can draft, edit, analyze, and organize information with impressive speed. But without direction, it produces generic, inconsistent, and sometimes unreliable work, the equivalent of a talented actor left to improvise without a script or a vision. Lawyers are now expected to take on the role of director when using AI. Becoming that director requires more than access to the technology or a willingness to experiment. It requires a shift in how lawyers think about the work itself, and prompt engineering is the language of that shift.

J. Cliff McKinney II is a managing member of Quattlebaum, Grooms & Tull PLLC in Little Rock, Arkansas. For a much more detailed discussion of AI ethics and prompt engineering, the author wrote a multipart series titled *Ethics of Artificial Intelligence for Lawyers*, published by the University of Arkansas Law Review's Law Notes Series, which can be found at <https://scholarworks.uark.edu/arlnlaw/>.

The way to combat hallucinations is to treat AI output the way you would treat a first draft from a junior associate: Review it critically, verify anything that matters before it goes out, and never cite a case, statute, or fact that you have not independently confirmed.

The potential ramifications of AI are hard to overstate. AI may prove to be the most transformative technology to touch the legal profession: more than word processing, personal computers, electronic research, or e-filing. This transformation could be profoundly positive. Lawyers who use AI well will do better work in less time. But using AI incorrectly will lead to errors and erode the judgment clients are paying for.

The transformation is already underway in everyday practice. Lawyers are experimenting with AI to draft lease provisions, summarize title commitments, analyze case strategies, and help prepare estate planning instruments. Sometimes the output is impressive, producing organized, coherent work that closely meets the needs. Other times, the result is useless or just plain wrong. The same tool can save an hour on one matter and create extra work on the next.

Understanding why that inconsistency happens is the first step toward fixing it. Lawyers are accustomed to precision in drafting and analysis, but many have not applied that same precision when communicating with AI. Two problems arise most often. First, AI does not “know” whether its answer is correct. It generates responses that are statistically likely to be coherent and relevant, which is very different from being right. Second, and worse,

AI delivers the wrong answer with the same confidence as right answers. These types of wrong answers are called “hallucinations”: output that sounds authoritative but is incomplete, imprecise, or simply false. The way to combat hallucinations is to treat AI output the way you would treat a first draft from a junior associate: Review it critically, verify anything that matters before it goes out, and never cite a case, statute, or fact that you have not independently confirmed. Also, the more specific and structured the prompt, the less room there is for hallucination, but verification discipline should be constant regardless of how good the output looks.

Another problem is that lawyers tend to treat AI like a search engine. A short, context-free prompt such as *draft a right of first refusal* or *summarize this lease* leaves too much for the system to guess at. Without knowing which party the lawyer represents, how risk should be allocated, what jurisdiction applies, or what style the lawyer prefers, AI defaults to a generic response that rarely serves the lawyer’s actual objectives.

The key to unlocking AI’s value lies in how instructions are given, in learning to direct the performance rather than simply starting it. That skill has a name: prompt engineering. Prompt engineering is the process of

communicating with AI in a deliberate, structured way to produce useful and reliable results. For transactional lawyers, it improves drafting and consistency. For litigators, it sharpens issue spotting and argument development. For both, it supports rather than replaces professional judgment. This article focuses on how to use prompt engineering in practice to improve AI use.

What Generative AI Is and Why It Works Differently

Before discussing how to use AI effectively, it is important to understand what AI is and what it is not. Lawyers are familiar with traditional software tools like word processors, document management systems, and legal research platforms that operate by retrieving stored information or applying predefined rules. Sophisticated research tools like Westlaw and Lexis are designed to locate existing authority or data in response to a query. Generative AI operates differently on a fundamental level.

Rather than retrieving information, AI generates new content. When a lawyer asks AI to draft a lease provision, summarize a title commitment, or analyze a legal issue, the system is not pulling a prewritten answer from a database. Instead, it predicts and constructs a response based on patterns it has learned from vast amounts of text. Those patterns include language, structure, and reasoning styles commonly found in legal and business writing, and this pattern recognition comes from training on literally trillions of pieces of information. But this same reliance on pattern recognition is the major flaw of AI because it will fail whenever it guesses the pattern incorrectly.

Five Techniques That Improve Results

Prompt engineering comes from computer science, where it refers to structuring inputs to produce better outputs. For lawyers, prompt engineering requires giving clear, deliberate instructions by telling the system



exactly what you need and why. Several techniques are particularly useful in legal practice.

1. Personas: Giving AI a Role

One of the most practical ways to improve AI output is to assign it a defined role. Without a role, the system has no perspective, no client, and no objective, so it defaults to a neutral, generalist output that produces generic results. Assigning a “persona” gives AI the focus it needs to produce useful results. It reorients the system from merely answering a question to advocating for a position, which is precisely what legal work requires.

Lawyers know their clients and understand their goals and objectives. Lawyers do not draft generic agreements or generic pleadings. Instead, they approach every task from a particular vantage point, advancing a specific client’s interests with purpose and deliberation. A persona directs AI to do the same thing.

Assigning a persona means expressly instructing the AI as to what role to play, whose interests to advance, and what the client is trying to achieve. A well-constructed persona instruction identifies three things: the professional role (e.g., experienced commercial real estate attorney), the client relationship (e.g., representing a landlord in a long-term net lease), and the objective (e.g., maximize the landlord’s protection and minimize exposure to tenant default risk). That instruction reframes how the system tackles the task from the ground up.

Consider the difference between these two prompts.

- The first: *Draft a trustee succession provision for a revocable trust.*
- The second: *Act as an experienced estate planning attorney representing an elderly client establishing a revocable trust. Your client has a blended family with children from two marriages, significant real estate holdings, and a primary concern*

about ensuring a smooth and conflict-free transition of trust management in the event of incapacity or death. Draft a trustee succession and removal provision that establishes clear, objective standards for determining incapacity, provides a smooth mechanism for successor trustee appointment, grants the successor trustee broad administrative authority, and minimizes the risk of family conflict or court intervention.

The first prompt produces a serviceable, generic provision. The second produces a clause that reflects the specific client’s family dynamics, planning objectives, and risk tolerance.

2. Chunking: Breaking Complex Tasks into Manageable Parts

One of the most effective ways to improve AI output is also familiar to lawyers: breaking a complex task into smaller, more manageable pieces. In prompt engineering, this is called “chunking.” Lawyers already read, draft,



edit, and analyze documents one section at a time. AI performs best when handling tasks the same way.

When asked to examine or draft an entire document in a single prompt, AI tends to produce high-level, generalized output. It may identify obvious issues but often misses nuance, and the response can feel more like a summary than a meaningful analysis. A related constraint is the “token” limit built into most AI systems. Tokens are the units of text that AI processes at one time, roughly corresponding to words or parts of words, and every system has a maximum per session. These limits usually depend on how much the system costs to use, with free or low-cost versions having lower token limits. When a lengthy document exceeds the token limit, the system may truncate content or produce degraded output. Chunking addresses both problems: It focuses AI’s attention on a discrete task

to produce deeper, more nuanced analysis and it keeps each prompt within token limits to avoid truncation or degraded output.

Breaking the task into components produces a very different result. Instead of asking AI to *review this lease*, a lawyer should direct it to analyze specific sections in sequence: first, rent and escalation provisions; then, assignment and subletting; then, indemnity and insurance; and finally, default and remedies. The same approach applies to estate planning. Rather than asking AI to *review this trust agreement*, a lawyer can direct it to analyze the trustee powers, then the distribution provisions, then the spendthrift and creditor protection language, and finally the amendment and revocation provisions, each in a separate, focused prompt. The same approach applies to drafting: A lawyer can build a document piece by piece, refining each provision

before moving to the next. Each prompt focuses the system’s full attention on a single issue, producing sharper, better results.

3. Few-Shot Prompting: Teaching Through Examples

Breaking a task into parts improves focus, but it does not by itself guarantee quality. Once each piece is isolated, the next challenge is ensuring AI follows the right approach to the problem. Lawyers are rarely looking for just any acceptable clause or analysis. They are looking for something that expresses a particular style, one shaped by the lawyer’s experience and judgment. That is where examples become especially useful. In prompt engineering, this is called “few-shot prompting.”

Few-shot prompting is the practice of supplying AI with examples before asking it to produce output, teaching it to model the examples. Experienced

lawyers already rely on prior work in exactly this way, including forms, model provisions, and language that the lawyer has refined over time. When assigning work to a junior attorney, senior lawyers often provide examples of what they want, knowing that those examples will shape the result. When given the same sample language or a model structure, AI can match the tone and approach with surprising accuracy.

That accuracy becomes especially valuable in transactional drafting. Consider a lawyer drafting an indemnity provision for a commercial lease who has a clear view of how risk should be allocated with broad protection for the landlord, specific carve-outs, and a defined structure for defense and reimbursement obligations. If that lawyer asks AI to *draft an indemnity clause*, the result will be generic. It may be usable, but it will almost certainly require revision. By contrast, when the lawyer supplies two or three samples of indemnity clauses that demonstrate the preferred style and directs AI to compose a similar clause, the outcome changes significantly. The result better matches expectations, the language reflects the intended level of protection, and the amount of revision required is reduced.

4. Iterative Refinement: Improving Through Successive Drafts

Even with clear structure and strong examples, AI's first output is rarely the final product. No experienced attorney expects a first draft to be ready for immediate use, whether it comes from a junior associate, a form file, or AI. Drafting is iterative, and the "iterative refinement" technique mirrors that reality. Iterative refinement uses successive revisions to sharpen the language and close the gaps.

One of the most common mistakes lawyers make is treating the first AI output as a finished product. When a response falls short, the natural instinct is to discard it and start over, but a more effective approach is to treat that initial output as a working draft. Targeted follow-up instructions can strengthen a provision, clarify the language, adjust

the tone, or refine a specific issue, with each successive prompt building on the version before it. The more specific the instruction, the better the revision, and multiple rounds are often necessary before the output is truly useful.

For example, a lawyer drafting a no-contest clause for a client with a blended family and a high risk of post-death disputes might begin by requesting a clause. The initial version may capture the basic structure but lack the precision the client's situation requires. The lawyer can then direct AI to broaden the clause to cover trust contests as well as will contests, clarify the probable cause exception, and address challenges brought by a guardian or conservator acting on behalf of a beneficiary. A subsequent prompt might address coordination with a specific bequest structured to give potential challengers a financial incentive to accept the plan rather than litigate. With each iteration, the clause becomes progressively tailored to the client's specific family dynamics and more useful as a planning tool.

Iterative refinement also improves efficiency. Rather than starting from a blank page for each revision, the lawyer is building on an existing draft. The system retains context, allowing it to incorporate prior changes and preserve consistency across revisions. At the same time, refinement reinforces the lawyer's role. AI can generate language, but it does not determine whether that language is appropriate. The lawyer remains responsible for evaluating the output and directing what refinement comes next.

5. Flipped Interaction and Perspective Switching: Stress Testing Ideas and Work Product

Even a well-organized workflow can miss issues. The quality of an analysis depends as much on the questions asked as on the answers given. Two prompt engineering techniques help with this. "Flipped interaction" reverses the conversation: Instead of asking AI to immediately produce an answer or tackle a project, the lawyer instructs it to ask clarifying questions

first, surfacing gaps and assumptions before any drafting begins. "Perspective switching" later directs AI to evaluate the work product from multiple viewpoints, such as those of the client, opposing counsel, or a court. Used in sequence, flipped interaction broadens the scope of inquiry before work begins, and perspective switching stress-tests the resulting work once it is underway.

Flipped interaction is especially effective at the outset of a task. Rather than immediately asking AI to draft or analyze, the lawyer instructs it to ask clarifying questions first. Before drafting a complaint, a litigator might prompt: *Before you begin, ask me every question you need to develop the strongest possible argument for my client.* In transactional work, a lawyer might instruct: *Before drafting any provisions, ask me the questions necessary to understand the deal structure, my client's priorities, and the key risks I want to address.* In probate practice, the same approach surfaces issues involving fiduciary duties, tax exposure, or family dynamics before any drafting begins. In each case, AI drives the inquiry, helping surface key issues before work begins.

Once those issues have been identified and first drafts developed, perspective switching puts the results to the test. The lawyer directs AI to evaluate a clause, argument, or strategy from multiple viewpoints in sequence: first from the client's perspective, then from opposing counsel's, and finally from the judge's. Each viewpoint exposes different concerns and demonstrates AI's ability to switch roles midstream like an improv actor. The client's perspective might reveal issues that still need to be addressed. Opposing counsel may identify weaknesses or likely negotiation points. A judicial perspective may concentrate on clarity, enforceability, and fairness. Together, they reveal a more complete picture than any single vantage point can provide. A lawyer also can ask AI to prepare counterarguments and potential compromises to the issues surfaced during these sessions.

Consider a landlord-protective lease provision. Viewed from the landlord's perspective, it may appear strong.

Analyzed from the tenant's perspective, it may reveal points of likely resistance or unintended ambiguity. A judicial perspective may show that the wording is overly broad or difficult to enforce. The same dynamic applies in litigation. An argument that appears persuasive at first may look far less secure when tested from the other side's perspective. Identifying those vulnerabilities early gives the lawyer the opportunity to refine the work before it becomes a real problem. Then, working with AI to develop counterarguments and compromises permits more effective negotiation.

These two techniques create a more disciplined analytical process. Flipped interaction expands the scope of inquiry by surfacing questions the lawyer may not have considered. Perspective switching then stress-tests the resulting work. Together, they mirror how experienced lawyers already think by finding the weaknesses in their own analysis before someone else does.

A Simple Workflow You Can Use Today

Prompt engineering techniques are most effective when used together. Consider a commercial lease negotiation. The lawyer breaks the lease into sections, works through each piece with targeted examples, and refines the drafts iteratively rather than accepting the first version. The lawyer can then ask AI to identify missing risks and test the work from the perspective of opposing counsel and a court. Each step reflects deliberate direction, not a single open-ended request.

To demonstrate how these techniques work together, consider an example: reviewing a commercial lease for a client acquiring a property or business. Instead of asking AI to *review the lease*, direct the process in steps.

Start like an improv director by assigning a role with a clear objective. For example: *Act as an experienced commercial real estate attorney representing a purchaser of a shopping center. Your objective is to identify provisions in this lease that could negatively affect value, impair future leasing flexibility, or create*

unexpected financial exposure for the landlord. Focus on practical risks, not just legal summaries.

That single prompt dramatically changes the output because the system is no longer summarizing but is instead analyzing from the client's perspective. Then, ask AI to help you plan the task: *Ask me five key questions that will identify the most important points to protect the client.* Next, break the lease into sections. Rather than uploading the entire document and asking for a general review, direct AI to focus on one issue at a time. For example: *Review only the rent and escalation provisions. Identify any unusual escalation formulas, pass-through expenses, or provisions that could materially increase operating costs or reduce predictability of income.*

Then move to the next issue: *Now review the assignment and subletting provisions. Identify any restrictions that could impair a future sale, refinancing, or re-tenanting of the property. Be specific about consent standards, recapture rights, and any landlord obligations.*

Where style and drafting matter, provide examples. For instance: *Using the style and structure of the following indemnity provisions, draft a revised clause that maximizes protection for the landlord while remaining commercially reasonable.*

From there, refine the output using targeted follow-ups applicable to the specific situation: *Strengthen this provision further in favor of the landlord. Add clear defense obligations, expand indemnified parties, and eliminate any ambiguity in scope.* Then: *Now revise this clause to address enforceability concerns under Arkansas law and avoid overbreadth that could create judicial resistance.*

Once the core provisions are developed, shift back to issue spotting using flipped interaction: *Before we finalize this review, ask me every question you need to identify risks or issues in this lease that I may have missed as a buyer acquiring the property. Do not begin your analysis until you have asked me those questions.* That prompt reverses the conversation by allowing AI to drive the inquiry and surface gaps the lawyer may not have considered. Finally, stress-test the result using perspective switching: *Now*

analyze this lease from the tenant's perspective. What provisions are most likely to be challenged or renegotiated? Then: Now switch roles again. You are a judge reviewing a dispute over this lease. Identify provisions that are ambiguous, overly broad, or likely to create enforcement problems.

The same workflow applies directly to litigation. Begin by assigning a role with a focused objective: *Act as an experienced trust and estate litigation attorney. Your client is a beneficiary of a revocable trust who believes the trustee has been self-dealing and mismanaging trust assets for personal benefit. Your objective is to develop the strongest possible breach of fiduciary duty claim, identify every viable theory of recovery, and anticipate the trustee's defenses.* Then use flipped interaction before any drafting begins: *Before you outline any arguments, ask me every question you need to fully understand the trust instrument, the trustee's conduct, the financial records, and any relevant communications. Do not proceed until you have asked those questions.* Once the facts are developed, chunk the analysis by issue: first the elements of breach of fiduciary duty and self-dealing, then damages and surcharge remedies, then removal of the trustee and appointment of a successor, and, finally, expected defenses such as exculpatory clauses or beneficiary consent. Use few-shot prompting where style matters: supply a sample demand letter or motion section and direct the system to match that structure and tone. Refine iteratively, strengthening arguments and tightening language with each pass. Finally, stress-test using perspective switching: *Now evaluate this case from the trustee's perspective. Identify the three strongest defenses and the weakest points in the beneficiary's theory.* Then: *You are the probate court judge assigned to this case on a motion for summary judgment. Identify the elements of the beneficiary's claim that are most vulnerable, any evidentiary gaps that could be fatal, and arguments that are likely to fail on the current record.* The result is a structured, thoroughly stress-tested litigation strategy developed through deliberate direction rather than a single open-ended request.

In each case, the lawyer is directing the development of the final output. This process takes only a few additional prompts, but it transforms the result. Instead of a generic summary, the lawyer receives a comprehensive product that can be reviewed, refined, verified, and perfected. This approach can be adapted to nearly any legal task.

Conclusion

AI does not eliminate the need for legal judgment; instead, it increases the need for it. The lawyer still decides which risks matter, which provisions need work, and which outputs can be trusted. AI that produces polished, confident writing does not reduce professional responsibility. Speed means nothing if the lawyer has not verified the results. Prompt engineering is the technique of intentionally directing AI, checking and refining the output, and owning the result.

Give AI a clear role and a focused task, and it becomes a useful assistant, one that can sharpen a draft, surface a risk you missed, or stress-test an argument. Leave it undirected, and it will give you something that looks right but may not be. The lawyers who get the most out of AI will not be the ones who simply adopt it earliest. They will be the ones who direct it most deliberately and take full ownership of the process and the end result. The tools are changing quickly, but the standard for good lawyering is not.

Of course, AI assisted with brainstorming, organizing, stress testing, and editing this article using the prompt engineering techniques described above. But every judgment call and editorial decision was the lawyer's. That is the point. Although AI assists, the lawyer must remain in charge.

This article is the first in a two-part series. The next installment will examine the ethical issues posed by AI, with an emphasis on ABA Formal Opinion 512, which is currently the leading guidance on AI use. ■

Fundamentals of Title Insurance

Second Edition

JAMES L. GOSDIN

Intended as a broad yet thorough overview, *Fundamentals of Title Insurance* explains what title insurance means, what it can and cannot do, and what the parties in a real estate deal can ask for in a transaction. Written in an accessible, user-friendly manner, this guide provides examples, practice pointers, charts and checklists throughout the text to explain key points. The author, general counsel for a nationwide title insurance company and a leading industry expert, covers all relevant topics, beginning with a general background of title insurance, its variations and common features, and typical title insurance coverage and issues. Additional topics include:



2025, 584 pages, 7x10
Paperback/ebook
PC: 5431142
Price: \$159.95 (list)
\$127.95 (RPTE Members)

- What is title insurance, and why should consumers obtain it
- Explanations of the 2006 ALTA Loan Policy and the 2006 ALTA Owner's Policy
- The expanded coverage residential loan policy
- When to secure a short form title policy
- Title insurance requirements of government-sponsored enterprises and agencies, such as Fannie Mae and Freddie Mac
- Regulation of rate and forms
- Types of endorsements and which ones to request
- Marked-up commitments, pro forma policies, and extended coverage
- Closing protection letters
- Guidance for retention limits and reinsurance in commercial transaction
- Specialized title insurance products, guarantees, and other products



ambar.org/fundtitleinsurance

ABA
AMERICAN BAR ASSOCIATION
Real Property, Trust and
Estate Law Section