JULY 13, 2026

The Best AI Outputs Come from a Better Process, Not Perfect Prompts

by Starling Underwood

We've all been there. You type what feels like a brilliant, comprehensive prompt, hit enter, and wait, only for the AI to return something that completely misses the mark. It's frustrating, it wastes time, and it usually turns into a game of prompt-engineering whack-a-mole.

That experience reinforces a simple truth: the best AI outputs do not come from perfect initial prompts. They come from a better process.

If you want to stop guessing and start getting elite results, stop treating AI like a magic 8-ball—shaking it, asking a question, and hoping the answer floats up. Instead, treat yourself as the senior attorney directing a capable but literal-minded team member.

Here is a three-step framework to build into your prompts today.

1. Demand Clarifying Questions (The Discovery Step)

When we give humans a task, they usually ask questions to fill in the blanks. AI, however, is a people-pleaser. It will generally hallucinate context just to give you an immediate answer. To prevent that, explicitly tell the AI to question you before it generates the final work product.

The catch is that if you simply say, "ask me questions," the AI may respond with dozens of queries, many of them tangential. Narrow the scope by instructing it to ask only the critical questions needed to proceed.

Sample prompt language: "Before beginning the task, ask me clarifying questions to fill in any gaps in understanding to improve the output. Focus only on information critical to completing the task or to improving your understanding of the context or task. Skip any nice-to-have questions."

Why it works: It forces the AI to identify its own blind spots and gives you a chance to provide context you may not have realized was missing. It also prompts you to consider your own blind spots. The questions returned may surface issues you had not considered, but that are important to the task at hand.

2. Request an Action Plan (The Blueprint Step)

Before the AI executes a complex task, require it to show its work. Whether it is drafting a motion for summary judgment, reviewing a due diligence document set, or outlining a client advisory, request an Action Plan first.

Ask the AI to outline the exact steps, structure, and methodology it intends to use.

- **Review**: Does the logic hold up?
- **Modify**: Is it missing a crucial angle?
- **Approve**: Give the green light once the blueprint looks solid.

Sample prompt language: "After I have answered any clarifying questions but before drafting anything, provide an Action Plan that outlines: (1) the structure you intend to use, (2) the key points or arguments you will address, (3) your proposed methodology, and (4) the sources you intend to search to accomplish your task. Do not proceed until I approve the plan."

Why it works: It is much easier to redirect the AI's approach at the planning stage than to revise a completed draft that went in the wrong direction.

3. Insist on a Verification Checklist (The Quality Control Step)

The final step ensures accountability through human-in-the-loop verification. When the AI delivers the final output, instruct it to include a corresponding checklist mapped directly to the approved Action Plan. That checklist should show how and where each step of the plan was executed.

Sample prompt language: "After the task is completed, provide a Verification Checklist that maps each element of the output back to the approved Action Plan. For each step in the plan, confirm that it was completed and identify where in the output it appears."

Why it works: It holds the AI accountable to its own blueprint. It prevents prompt drift, where the AI starts strong but strays from the original instructions during a multi-step task or extended sequence. It also lets you quickly verify that the tool took the required steps.

Putting It into Practice: From Prompt Writer to Active Director

This framework is more than a set of prompting tips. It reflects a fundamental shift in how you interact with AI.

Instead of viewing yourself as a prompt writer trying to anticipate every piece of context the AI might need, you become an Active Director who sets specific output goals and defines the parameters within which the AI operates.

This approach flips the traditional prompting dynamic. Rather than trying to provide all the context up front, you let the AI surface the information it actually needs. This makes the interaction more efficient, more responsive, and more controlled.

The lawyers who get the best results from AI will not be the ones who write the longest prompts. They will be the ones who build the best process.

That is how you turn AI from a guessing game into a real productivity tool.