



MAD-156B
MAD GENIUS

08/07/25
Version 02

Blog Title Model Behavior: How to Get the Most From AI

Author Kiefer Slaton, Web Director

Post Date XX/XX/XX

Keywords Example, example, example

Model Behavior: How to Get the Most From AI

June 2025 was the beginning of a beautiful relationship between developers and artificial intelligence.

A tool called Cursor gained prominence, allowing us to work with the major AI models directly in the codebase. Around the same time, Anthropic dropped their new flagship large-language model, Claude Sonnet 4, a quantum leap ahead of the previous model. Claude Sonnet 4 is the first model to really be able to execute high-level code work with a good degree of accuracy and sense.

Together, these tools spawned a new term: “vibe coding.” People with little to no experience with code started downloading AI-driven tools like Cursor and writing (and shipping) apps and websites from scratch, using only plain text prompts. Countless seasoned developers began wondering if this was the moment that we became irrelevant.

Then, in July, things got dark on r/CursorAI.

The subreddit started getting flooded with developers complaining about significantly choked-back usage limits. Overnight, Cursor (and most of its competitors) significantly raised prices for access to the highest usage and best models. Simultaneously, Claude Sonnet 4 started seeming “dumber.” It started making silly mistakes, dramatically overengineering solutions, and not responding to feedback.

The Renaissance was over as soon as it had begun.

The frustrations of vibe coders are somewhat warranted, especially the price hikes. But there's also an underlying misunderstanding and expectation of what these models are capable of. Despite how it may feel when talking to ChatGPT, it is not a human mind capable of higher-order thinking, prioritization, or learning from its mistakes. These models are just highly sophisticated tools, and like any other tool, they can dramatically improve productivity and quality when used well and can cause frustration and amplify mistakes when used poorly.

Data scientists have a phrase, "Garbage in, garbage out," that essentially means bad data leads to bad conclusions. This principle applies to the use of AI-assisted work (which makes sense, given that generative AI is just a statistical model under the hood).

I'm going to share a few things I've learned from working daily with these tools for the last year and a half. Regardless of your job function or which tool or model you're using, these principles will ensure you get the most out of collaborating with AI.

Tip #1: Know Your Vision and Communicate It Clearly

Society has experienced many painful disappointments in recent history. Google Glass. Theranos. New Coke. But there is no greater modern collective letdown than season eight of "Game of Thrones."

This show revolutionized television and created an entirely new generation of fantasy fans from the moment Jaime Lannister quipped, "The things I do for love," and shoved Bran Stark from a tower on April 17, 2011. For eight years, the show was a foundational tentpole of pop culture, with its most popular episode drawing over 8 million viewers on release day. You're hard-pressed to find anyone with a TV today who couldn't hum the theme song or whose blood doesn't turn cold at the phrase, "The Lannisters send their regards." In its first six seasons, it was included in discussions with "Breaking Bad," "The Sopranos," and "The Wire" as one of the greatest shows to ever be on television.

Then seasons seven and eight came. Everything went wrong.

Spoiler warning (but is it though?).

Daenerys Targaryen, one of the show's primary protagonists, went from a petulant but good-hearted young ruler to a straight-up Dr. Evil-style comic book villain over the course of one or two episodes. Jaime, the show's villain turned sympathetic anti-hero, seemingly forgot everything he learned over the course of the whole show and left behind a healthy relationship to be crushed by rubble while embracing his maniacal sister. Bran Stark, who spent the latter half of the show plumbing the deepest secrets of the mystic arts, was appointed King of Westeros, because reasons. Entire plotlines were dropped, the ones that were kept were rushed to unsatisfying conclusions, and almost 20 million viewers were left stunned as the final episode's credits rolled. Unlike other shows like "Lost" or "The Walking Dead" that had weak

final seasons but are still remembered fondly, the final season of “Game of Thrones” was so trash that it tarnished the show’s legacy.

So, what on earth went wrong?

If you know about “Game of Thrones,” you probably know that it was based on a book series written by George R.R. Martin. Martin’s books are excellent, filled with rich detail, complex characters, and powerful plotlines that are captivating, heartbreaking, and inspiring all in one. As such, when HBO adapted the show, they weren’t interested in hiring showrunners with plenty of original ideas, but rather ones that could faithfully bring Martin’s stories to life on screen. They found these showrunners in David Benioff and D.B. Weiss, and for most of the show’s run, they adapted the source material into award-winning television that kept viewers on the edges of their seats.

Unfortunately, Martin lost his motivation to write the last two books, so roughly around the end of season five, Benioff and Weiss had nothing left to adapt. They had to freestyle a conclusion that faithfully paid off the books’ dozens of divergent storylines, but they didn’t really know where they were heading. Without a clear vision to follow from the story’s original creator, the show lost the spirit that had made it such a phenomenon.

AI is the Benioff and Weiss of your working team. It can and will produce high-quality output on just about anything you ask of it, but only if you can articulate your vision for the task in a vivid and specific way. In the AI community, this concept is called prompt engineering. A clearly thought-out and written prompt is the first step to getting responses that are high-quality, clearly address the issue or question, and fit within the style you had in mind.

Let’s say I wanted to write this article using an LLM like ChatGPT (I didn’t, I promise). I could open it up and demand, “Write a blog article about the best ways to use LLMs.” It would certainly give me the article I asked for, but it likely wouldn’t hit the points I want to hit, sound anything like me, or communicate to the audience I’m trying to reach.

Now, let’s say I gave it this prompt: “Write a 1000-word blog article about best practices for using generative AI that will apply regardless of which specific tool is being used. The article should include an introduction that talks about the volatility of the AI market, three distinct sections on prompt engineering, context management, and quality assurance, and a conclusion that encourages the user to try the tools for themselves. The target audience is creative professionals that may use LLMs in various ways in their daily work. The voice of the article should be casual, even a little sarcastic or snarky at times. For each concept laid out in the article, find a good analogy that can be universally understood that demonstrates the value of that concept.”

This would get me much closer to the final output I want. Yes, I would still expect to have some back and forth to refine it, and I would want to do a final editing round myself (see [Tip #3](#)). But spending just a few minutes clarifying my message would save me hours reworking the vague output it gave before.

This concept is why I take issue with the idea that AI is “dumbing down” the workforce and hampering creativity. Yes, you can generate a generic article, image, or app without needing a lot of skill or understanding. But if you want to get the most out of generative AI and use it to make things that feel unique to you or your business, you still have to know where you’re trying to go. You have to communicate those directions clearly.

Tip #2: Give It Your Memory

We’ve established that a well-written prompt will get you unique and quality output when working with AI. However, you may quickly find yourself exhausted by having to provide an essay every time you require a simple task to be completed. You will likely repeat the same things over and over again about your standards and historical context for a project. This repetitive work will start to leave you questioning if the tool is worth it.

AI is your most talented employee with amnesia. Every time you assign them work, you have to re-explain the same guidelines: “Remember, we always include client testimonials in our proposals,” “Don’t forget our reports use this specific formatting,” and “We never use industry jargon when writing for customers.” After the third or fourth time giving identical instructions, you’d be pulling your hair out. Any reasonable manager would hand them a comprehensive style guide and say, “Here, keep this at your desk and reference it for every project.”

This is precisely what context engineering accomplishes for AI. Almost all AI tools allow you to upload documents, code files, PDFs, or images that become permanent reference materials for your projects. Instead of retyping your brand guidelines in every single prompt, you upload them once, and they’re always available. Instead of repeatedly explaining your reporting format, you provide template examples that eliminate the guesswork.

What goes in your AI’s reference library will depend on your role. As a developer, I upload code style guides and architecture documentation. A copywriter might include brand voice documents and approved messaging frameworks. A project manager could provide meeting templates and stakeholder communication standards. A designer might upload brand guidelines and successful campaign examples. A writer recently abandoned by the book writer who inspired them may upload the first five books that were complete so they don’t go completely off track and ruin everything they’ve built.

The payoff is huge: your prompts become focused task requests rather than lengthy re-explanations of the same standards. You can simply say, “Write a quarterly report,” and trust that it will follow your

established formatting, tone, and content requirements without you having to spell them out for the hundredth time.

Tip #3: Please Check Its Work

I'll keep this one short and sweet, but it's critically important. If there's one thing I can impress upon you about the current state of the frontier AI models, it's this: they are inconsistent. Their "intelligence" and capability to understand context and constraints fluctuate hourly with the latest build of the model, the type of prompt you give, and even capacity constraints on the model server. You can allow it to be the creator, but you need to be the editor. Whether it's proofreading an article, writing automated tests for a code block, or checking the number of fingers on the hands it generated, don't send anything from one of these models out into the world unverified. I think you'll often find that even with a thorough prompt and great context, it will still take a few iterations before you're happy with the result.

Summary

I hope I've impressed upon you that these AI models are, at the end of the day, highly sophisticated tools, no different from the calculator, computer, or iPhone before them. As with all tools, the best way to build proficiency with generative AI is to dive in and practice.

Early in the process, I would ask ChatGPT for suggestions on how to use it in my daily work. A prompt like this can really give you some great ideas on where to start: "I am an <Insert Job Role> at an <Insert Company Type>. Ask me questions about my daily workflow, then provide suggestions on how I could use generative AI to be more efficient."

I also encourage you to try a few of the different flagship models to see which one fits you best, because they all "feel" slightly different. ChatGPT is matter-of-fact, straight, and to the point, while Claude is a bit more conversational. Gemini excels at handling massive batches of context, making it great to have high-level planning discussions, while the other models are much better at handling specific requests. Start with simple, low-stakes requests or personal projects, and over time you will learn its quirks and build confidence using it.

I'll leave you with this parting thought: however you feel about artificial intelligence, it is here to stay.

History books will look back on this time as equally important to the advent of the internet or social media. This will be a step change in the way work is done, and those that don't learn to use it are at risk of getting left behind. The fear that surrounds this technology is largely unwarranted.

Give it a try. I think you'll find that it is neither a competent worker coming for your job nor a HAL9000-style entity plotting to overthrow humanity. It is a new and powerful resource for the workplace, and when used properly and judiciously, it can unlock a whole new set of possibilities for you in your work.