

Hippos, Corduroy, and Hallucinations

How I Tamed the AI Hallucination Monster



Picture this: you ask your new AI image tool for a friendly holiday scene, and it delivers...a hippo lounging in a corduroy chair beside a Christmas tree. Huh? You'd be forgiven for thinking I stumbled into a Dr. Seuss story. But no, the culprit was a poorly constructed prompt—yep, my bad! While the result was undeniably entertaining (and might even look great on a holiday sweater), it's also a cautionary tale about the "hallucination monster" lurking in every AI. Buckle up, because in this article we'll explore how you can keep your AI from creating wacky, hippo-infested dreams, and instead guide it to accurately produce the answers you need, every time.

You use AI, whether you know it or not. Maybe it's your phone finishing your sentences, or your kids chatting up a virtual assistant to help with homework. Whether you view AI as a shiny new toy or the villain in Hollywood's new apocalyptic plot, one thing's for sure: it's here to stay. Think of AI like a hammer, an incredibly powerful tool that comes with a simple instruction: "Don't whack your thumb!" In the pages ahead, we'll walk through the best ways to avoid AI's equivalent of a thumb smash – hallucinations - and harness its power for clear, rational, and truly useful answers. Let's start with a quick example of how the AI magic works.



So many choices!

When you provide that lazy AI with an easy-to-follow path, you get a direct strike, guided to your desired answer. Let's talk about how to do this, and I'll share examples of what happened when I didn't take my own advice.



Provide an easy path

Think of a bolt of lightning which starts with your request and strikes on your answer. Like real lightning, AI takes the easiest path to an answer. Where you have many possibilities, you can't predict which fork the lightning will take.



Poor prompts are rewarded with kale smoothies

Start With a Clear and Specific Objective

Recently my wife and I visited a trendy "healthy donut shop." I asked for "something sweet, but healthy," envisioning a guilt-free treat. The barista handed me a kale blast with a whipped cream topping (YUCK!!). My wife, on the other hand, is laser-focused: "I'd like a glazed donut with Bavarian cream in the center, topped with a chocolate drizzle." The clerk nodded and handed her an irresistibly decadent pastry that looked like it came straight from a gourmet bakery. Moral of the story? Vague requests get vague (and often disappointing) results. Specific requests get you exactly what you want—just like when you're crafting prompts for AI.

How to do it:

- Before asking a question, ask yourself: "What do I want out of this response?"
- State your goal explicitly:
 - ✓ Poor prompt: "Give me a recipe for a healthy dessert."
 - ✓ Good prompt: "Give me a low-carb dessert recipe that uses allulose instead of sugar, maintains a rich caramelized flavor, and pairs well with coffee."

Just like ordering at a donut shop, **the more precise you are, the better the results!**

Set Clear Context to Avoid Misinterpretation

My father-in-law, a passionate bowler and retired Air Force officer, asked me to show him an AI demo. What better way than having AI generate an image that combines his two great interests, an Air Force officer landing a perfect strike in a retro bowling alley? Amusingly, the AI took the path of least resistance, drawing connections between “air force” and “strike” rather than linking “strike” to bowling!



Prompt gone rogue: Turns out, military precision doesn't apply at the bowling alley.

The takeaway? **AI needs clear context and constraints to deliver accurate results. Without them, it may focus on unintended connections.**

How to do it:

Provide necessary background—Who is the audience? What format do you need?

- ✓ Specify constraints—Word limits, key details, or exclusions.
- ✓ Use negative prompts—Explicitly tell AI what to avoid preventing off-topic or misleading results.

Example Prompt:

- ✓ P o o r : *“Create an image of an Air Force officer throwing a strike.”*
- ✓ *“Create an image of an Air Force officer at a 1950s-style bowling alley, releasing a bowling ball down the lane. The focus should be on the sport—avoid any military combat elements.”*

By guiding AI with clear instructions, you ensure it delivers the right strike—whether in bowling or in generating responses.

Use Unambiguous Language and Structured Prompts

During a conversation with my global team, we were discussing interesting food, fried alligator! Talking about large lizards, someone mentioned Komodo dragons, so I asked AI for a mean-looking one. Unfortunately, a typo (“commodo” instead of “Komodo”) resulted in this horrifying picture. Lesson learned—**precision matters!**



- ✓ AI can misinterpret vague references, typos, or ambiguous terms.
- ✓ Structured prompts help keep responses organized and relevant.


How to Do It:

- ✓ Be specific—Avoid unclear pronouns like *it* or *they* without a clear subject. *Typos matter! One letter off, one nightmare created.*
 - Break down complex prompts—Use bullet points or numbered steps for clarity.
 - Use consistent keywords—Prevent AI from making unintended associations.
- ✓ Guide AI with Examples and Clear Formatting
Generative AI mirrors the structure and style of the prompt.
 - Want a table? Provide an example layout.
 - Want bullet points? Use them in your prompt.
- ✓ Keep It Professional and Clear
AI often reflects the tone of your prompt. If you want a neutral, authoritative response:
 - Avoid sarcasm or overly casual language.
 - Use direct, professional wording.
- ✓ Proofread for typos (or risk summoning toilet dragons).
- ✓ Read complex prompts aloud to ensure clarity.
- ✓ AI follows what you type—**garbage in, garbage out!**

Challenge the Answers

Why it matters:

- ✓ AI models can produce confident but wildly incorrect responses.
- ✓ Cross-referencing or verifying information helps identify mistakes.
- ✓ Sometimes the first answer is not exactly what you need. **Use your RI (Real Intelligence) to challenge and question the answers.**



YOU are the judge!

How to do it:

- ✓ Ask the model to show its sources or to provide references.
- ✓ Prompt the AI to double-check its own answers or to phrase uncertain information as “based on the data I have” rather than stating it as fact.
- ✓ Strange as it sounds, **argue with your AI!**
- ✓ You can refine your prompt based on the AI’s initial response for improved results.
- ✓ Ask follow-up questions. Challenge the answers!
- ✓ If the AI’s answer is too long, ask it to summarize. If it’s unclear, request more details or clarifications.
- ✓ Keep narrowing or broadening the prompt until you get the desired outcome

Example verification prompt: “Provide a factual summary of the event and include links to reputable sources for verification. If no reputable source is available, indicate that you are unsure.”

Limits, Limits, Limits!

Generative AI is all artificial and NOT intelligent

- ✓ *NEVER* rely on it blindly. Always verify information independently.
- ✓ Ask for alternative sources or fact-check if the response covers important information (e.g., medical, legal, financial data).

Quick Guide to Monster Taming

1. Define Your Goal – State the outcome you want.
2. Provide Context – Give relevant details and constraints. Set Guardrails – Specify what is and isn’t allowed.
3. Use Clear Language – Avoid ambiguous phrasing. Maintain Tone – Use the style you want mirrored. Give Examples – Show formats or templates.
4. Verify, Check, and Iterate – Ask for sources or disclaimers. Refine your prompt based on the result.
5. Proofread – Check your prompt before submitting.
6. Know AI’s Limits – **Always** fact-check critical information.

Thanks for reading – and may all your dreams be of hippos and dragons!