Practical Ways to Get More Reliable Output from Generative AI

Why Large Language Models Sometimes Get Things Wrong

Generative AI tools have become a familiar part of everyday office work. They draft emails, summarise documents, generate ideas, and help with research. But anyone who has used them regularly has seen the same problem: the model confidently produces something that is not true. It invents a reference, misstates a fact, or draws a conclusion that does not follow from the data. These errors are often called hallucinations, and they can undermine trust in the tool. The good news is that you can reduce their frequency by adjusting how you prompt and how you review the output.

Understanding What Causes Hallucinations

An AI language model does not know facts the way a database does. It predicts the next word based on patterns it learned during training. When it generates a sentence, it is not checking a source of truth. It is assembling text that looks plausible given the prompt and the millions of documents it has seen. This means that when the model does not have enough context to predict accurately, it fills the gap with something that sounds reasonable but may be wrong. The key is to give the model enough structure and constraint so that it does not need to invent details.

Write Clear, Structured Prompts

The single most effective technique for reducing AI hallucinations is to write prompts that are specific and structured. Vague prompts invite the model to guess. For example, asking "Tell me about the history of Singapore" will produce a broad narrative that may include errors or generalisations. Instead, break the request into parts: "List the main events in Singapore's history from 1965 to 1975, with one sentence per event." By narrowing the scope and specifying the format, you limit the room for invention. The model has less freedom to wander into uncertain territory.

Include Context and Examples

When you give the model a frame of reference, it stays closer to the desired topic. If you want a summary of a report, paste the relevant section into the prompt and say "Summarise the following text in three bullet points." The model works from the text you provide, not from its own memory. Similarly, if you ask for a business email, include the key points you want covered

and the tone you need. The more the prompt resembles a template, the more predictable the output becomes.

Ask for Citations or Sources

One practical trick is to ask the model to include its sources or to state when it is unsure. You can phrase it as: "If you are not certain about a fact, say 'I am not sure' instead of guessing." This will not completely eliminate errors, but it often makes the model more cautious. Some tools also allow you to instruct the model to respond only from a given document or set of notes, which drastically reduces hallucinations because the model stays within a closed set of facts.

Break Complex Tasks into Steps

Many hallucinations happen when the model tries to answer a complex question in one go. If you ask a multi-part question, the model may conflate information or skip a step. Instead, break the task into a series of smaller prompts. For instance, first ask: "List the main arguments in this article." Then follow up: "For each argument, give one supporting point from the text." By guiding the model step by step, you keep it focused and reduce the chance of straying into invented content. This approach also makes it easier to spot errors at each stage before combining the results.

Review Output Systematically

No prompt technique is perfect. You should always review generated text with a critical eye. Read for plausibility, not just fluency. If the model states a specific number or a date, check it against a reliable source. If it names a person or a publication, verify that the reference exists. Over time, you will learn which topics your tool handles well and where it tends to slip. Some office tasks, such as summarising internal documents, are less prone to hallucination because the model works from provided text. Other tasks, such as generating facts about niche subjects, require more caution.

One common pitfall is treating the model's output as final. The best workflow treats the AI as a first draft generator. You edit, correct, and verify before sending anything to a colleague or a client. This is not a sign of failure; it is how professionals use the tool. The model saves time on structure and phrasing, and you save it from making mistakes.

Use the Right Tool for the Job

Not all generative AI tools are equally reliable. Some are designed for creative writing and will happily invent details. Others are built with stronger guardrails and better grounding in source

documents. When you choose a tool for office work, look for ones that allow you to upload reference materials or that explicitly cite sources. The more control you have over what the model can reference, the less it will need to guess. A tool that can point to the paragraph it used for a claim is far more trustworthy than one that simply outputs text with no provenance.

Iterate on the Prompt

If the first output contains errors, do not just ask the same question again. Revise the prompt. Add constraints: "Only use information from this document." Specify the format: "Output a table with two columns." Ask for a shorter answer: "Limit your response to 100 words." Each change narrows the model's possible responses and makes it less likely to hallucinate. Over time, you will build a library of prompt patterns that work reliably for common office tasks. This iterative process is where the real skill of using AI develops.

The Practical Role of Reducing AI Hallucinations in Daily Work

For anyone who uses generative AI regularly, reducing AI hallucinations is not an abstract goal. It directly affects how much you trust the tool and how much time you spend fixing its output. A well-structured prompt can cut the error rate noticeably, but it does not eliminate mistakes entirely. The combination of clear instructions, step-by-step reasoning, and thorough review makes the difference between a tool that wastes time and one that saves it. The goal is not to achieve perfect accuracy every time. That is unrealistic. The goal is to get output that is useful enough that your own editing effort is smaller than the effort of writing from scratch.

In practice, this means treating each generation as a collaboration. You provide the structure and the boundaries. The model provides the words. If the words drift outside the boundaries, you adjust the prompt and try again. Over a few rounds, the output converges on something accurate and usable. This is not a magical process. It is a routine skill that any office worker can learn with a little practice.

Build Your Own Prompt Library

A simple way to improve consistency is to save prompts that work well. When you find a prompt that produces reliable summaries, accurate data extraction, or clear explanations, store it in a document or a shared folder. Over time, you build a set of templates for the tasks you do most often. This reduces the need to invent new prompts each time and lowers the chance of introducing vague instructions that lead to hallucinations. It also helps new team members get up to speed faster. They can use the same prompts and get similar results, because the structure is already tuned for accuracy.

None of these techniques require technical expertise. They are about how you ask, what you ask, and how you check the answer. The most important habit is simply to treat the model as a junior assistant whose work needs review, not as an oracle. With that mindset, you can get a lot of value from generative AI while keeping the mistakes under control.