

You gave the AI a perfectly reasonable prompt. It gave you perfectly forgettable prose. Here's what's actually going wrong — and the practical fixes that work.

There's a sentence pattern that shows up in AI-generated writing so reliably you could set a clock by it: *"In today's fast-paced world..."* Or maybe *"Certainly! Here is a comprehensive overview..."* You know the ones. The question most people don't ask is *why* this happens — and what it means for how you prompt.

The short answer: AI models are trained to predict the most statistically probable next token. When your prompt is vague, the model falls back on whatever patterns appear most often in its training data. Business writing is full of hedge-words, filler phrases, and structural clichés — so that's exactly what you get back.

The good news is that this is almost entirely a prompting problem, not a model problem. The same model that generates boilerplate can produce writing with genuine voice, specificity, and personality — if you tell it to.

Why Generic Output Happens: Three Root Causes

Before fixing the problem, it helps to understand exactly where it comes from. There are three main culprits:

1. **Vague prompts.** Without constraints, the model picks the most probable output — statistically average prose.
2. **No defined audience.** Writing to nobody means writing for everybody, which means writing for nobody well.
3. **Missing examples.** Without a reference for tone or style, the model defaults to the most generic register it knows.

The Real Problem Is Under-Specification

When you type “write a blog post about productivity,” you’ve given the model a topic and nothing else. No audience. No tone. No angle. No word count. No format constraints. No prior examples of writing you like. The model isn’t being lazy — it’s filling in every missing decision with its best average guess. And averages, by definition, aren’t memorable.

Think about what a human writer would need before starting that same assignment. They’d want to know: who’s reading this? What do they already know? What should they feel or do differently afterward? Is this punchy and short, or long-form and nuanced? Should it sound like the author talking to a friend, or like a polished magazine feature?

Your prompt should answer those questions before the model has to guess.

The model isn’t being lazy — it’s filling in every missing decision with its best average guess. Averages, by definition, aren’t memorable.

Fix 1: Give It a Persona and an Audience in the Same Breath

The fastest upgrade to any writing prompt is to specify who is writing *and* who they’re writing for. These two constraints alone eliminate enormous swathes of generic output.

Before:

Write a blog post about time management for professionals.

After:

You're a former management consultant who now coaches early-stage founders. Write a blog post for first-time startup CEOs who are drowning in meetings and haven't shipped anything in three weeks. Be direct. No platitudes.

Notice that the “after” version does three things simultaneously: it establishes a credible voice, it makes the audience hyper-specific (not “professionals” — a drowning first-time CEO), and it gives a clear emotional situation the reader is in. Each of those constraints is doing real work.

Fix 2: Use Negative Examples, Not Just Positive Ones

Most people tell the AI what they want. Fewer people tell it what they *don't* want. This is a significant missed opportunity, because exclusion is often easier to specify than inclusion.

Add a “don't” clause to your prompt:

Do not use the phrases "in today's world," "it's important to," "leverage," "game-changer," or "at the end of the day." Do not open with a rhetorical question. Do not use bullet points.

This sounds almost comically simple. It works anyway. The model will actively avoid those patterns rather than defaulting to them.

Tip: Keep a running list of phrases you hate seeing in AI output. Paste it into your prompts.

Fix 3: Paste In an Example of Writing You Actually Like

Describing a writing style is hard. Showing it is easy. If you've ever read a piece and thought “I want this AI to write like *that*” — paste a paragraph of it directly into your prompt. You don't need a full style guide. Two to three sentences of example prose is enough to shift the model's register noticeably.

Style anchoring prompt:

Match the tone and sentence rhythm of this passage: [paste excerpt]. Don't copy it — use it as a calibration for how blunt, short, and direct the writing should feel.

This technique — sometimes called “style anchoring” — works because you’re giving the model a concrete target instead of an abstract description like “casual but professional” (which means something different to everyone).

Fix 4: Ask for a Draft, Then Ask It to Edit Itself

One of the underused patterns in prompt engineering is the two-pass approach. Generate a first draft, then send it back to the model with a specific editing instruction. This forces a second perspective on the output and tends to strip out filler that slipped through the first time.

Two-pass editing prompt:

Now rewrite that draft. Cut every sentence that doesn't either (a) make a specific claim or (b) move the reader forward. If a sentence is throat-clearing, remove it. Target 20% shorter.

The editing pass tends to be more effective than trying to get a perfect draft on the first try, because you can be far more specific about what’s wrong once you have actual text to point at.

Fix 5: Specify Format With Surprising Specificity

Left to its defaults, the model will produce the most common format for whatever you asked for. Blog post? Three headers, five paragraphs each, bullet points under each header. Predictable structure produces predictable reading experiences.

Unusual format constraints force unusual thinking. Try asking for a piece with no headers at all. Or one that opens in the middle of an anecdote. Or a recommendation that leads with the conclusion and then justifies it. Or a listicle where each item is exactly one sentence.

Constraints like these are small — but they interrupt the default pattern, and interruptions tend to produce more interesting output.

None of these fixes require deep technical knowledge. They require specificity — the same specificity you’d use when briefing a junior writer. The more clearly you

define the situation, the audience, the constraints, and the tone, the less the model has to guess. And the less it guesses, the less it sounds like every other AI output on the internet.

The irony is that making AI writing sound less like AI writing is mostly a human skill: the ability to describe what you want with precision. That's something worth developing.

Want more prompting techniques like this? I share practical, copy-paste prompt strategies every week. [Subscribe to the newsletter](#) — or browse the full [prompt library](#) to go deeper.

Share this:

- [Share on Facebook \(Opens in new window\) Facebook](#)
- [Share on X \(Opens in new window\) X](#)