RYAN K. BOETTGER

WHITE PAPER

Becoming an AI Whisperer: A Field Guide to Prompt Engineering

Ryan K. Boettger · June 12, 2026 · White Paper

> This is the field-guide version of my *Intercom* article, *From Technical Communicator to AI Whisperer* — a little longer, a lot more hands-on, and with a few prompting techniques borrowed from my THECB assessment workshop folded in.

A few years ago, a survey in *Intercom* asked technical communicators whether technology had taken over the tasks they do at work. More than 70% said rarely or never — and the same share doubted their jobs could be automated by 2030. Then the pandemic hit, AI showed up in everyone's tools, and "experience with AI" started appearing in job postings everywhere.

So here's the good news, because there is good news: the robots aren't coming for your job. But they *are* changing it, and the writers who thrive are the ones who realize something quietly reassuring — **the skill that makes you good at this work is exactly the skill a machine needs from you.** Clarity. Concision. Knowing your audience. Getting a complicated idea to land. We have a new name for applying those skills to an AI: prompt engineering.

What prompt engineering actually is

Strip away the hype and it's this: **writing clear, purposeful instructions that get an AI to produce what you actually want.** Not magic words, not secret incantations — just good communication, aimed at a strange new audience that happens to be made of math. If you've ever rewritten a set of instructions because a reader misunderstood them, you already have the instinct. You're just learning a new reader.

The core move: Do-What-How

When people get frustrated with AI output, it's almost always because the prompt was vague. The fix I teach is a little three-part scaffold — **Do-What-How**:

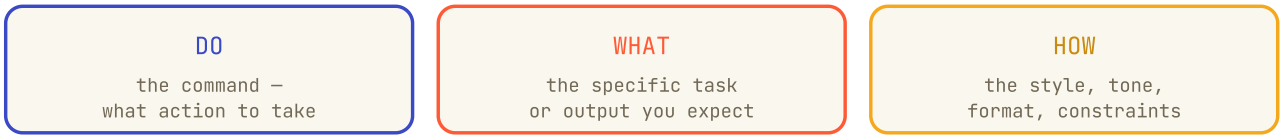

Watch what it does to a real prompt. Start basic:

```
Instruct a general user how to create a style in MS Word in five steps
or fewer, using clear and concise language suited for beginners.
```

That's **DO** (instruct a user), **WHAT** (create a style in MS Word, five steps or fewer), and **HOW** (clear, concise, beginner-friendly). Now keep the Do and What, and change only the How:

```
Instruct a general user how to create a style in MS Word in five steps
or fewer, in the voice of Deadpool.
```

Same task, wildly different result — because **How** is where tone and style live. Once you see the three parts, you start noticing which one your prompt forgot.

Prime it first: Chain-of-Thought

Here's a technique I lean on in my assessment work that pairs beautifully with Do-What-How. Before you ask for the thing, **ask the model to think out loud about it.** Instead of demanding a rubric (or a memo, or a test) cold, you warm it up:

```
I'm a college professor. What can you tell me about assessment rubrics?
```

Then — *after* it's surfaced what it knows — you give it the Do-What-How prompt. This is **Chain-of-Thought**: you get measurably better output when the model reasons before it produces, instead of leaping straight to an answer. Think of it as letting your collaborator clear their throat before the big speech.

Interview the AI

A related trick, and one of my favorites: **interview the model to learn how it was trained.** I once wanted to apply Kitty Locker's research on goodwill to some AI-generated customer-service replies, so I simply asked the AI whether it knew her work, and what it called those principles. It told me the term it associated with her research — and I used *that* term in my prompt. The lesson: align your language with the model's training instead of imposing your own vocabulary and hoping it sticks.

Learn to spot the tells

Generate enough text and you start seeing the seams. Synthetic writing tends toward repetitive phrasing, overwritten flowery prose, and — my favorite name for a real problem — **Bow-itis**: the model's compulsion to tie everything up in a

neat little bow at the end, whether the moment calls for it or not. Once you can name the tropes, you can prompt them away (and edit out the ones that slip through).

Get under the hood

The web versions of Claude and ChatGPT hide the dials. But in a developer interface — like OpenAI's Playground — you can actually turn them, and it's worth knowing what they do:

- **Temperature** — the creativity knob. Low (say, 0.7) gives predictable, conservative output that's perfect for business and technical writing; high gives you variety and surprise.
- **Top P** — how wide a net the model casts for its next word. Lower is narrow and focused; higher opens the door to more possibilities.
- **Frequency penalty** — discourages repeating the same words and phrases.
- **Presence penalty** — discourages circling the same topics, nudging the model to move the ideas forward.

A temperature of 0.7 with a Top P of 0.8 is a reliable starting point for most professional writing. Creative work? Turn them up and see what happens.

The ethics you can’t skip

No honest conversation about AI skips this part. A few rules I hold to:

- **Keep it original, accurate, and human-edited.** You are the last line of defense.
- **Don't feed proprietary or client data into public models.** Once it's in, it's in.
- **Describe a voice; don't impersonate one.** Rather than "write in Jane's voice," name the *features* that make Jane's voice hers — which, conveniently, is exactly the language-and- rhetoric skill that makes technical communicators good at this in the first place.
- **Be transparent.** Audiences deserve to know when content is machine-made. If your organization has no AI-writing policy yet, now's the time.
- **Watch for bias.** Models inherit it from their training data; you're responsible for catching it and keeping the work inclusive.

You own the output. Every time.

Where to start

If you take nothing else from this: prompt engineering isn't a new profession bolted onto ours — it's our same old craft, pointed at a new collaborator. Study how language works in your field. Make Do-What-How a habit. Prime with Chain-of-Thought. Learn to spot the tropes. Turn the dials when you can. And never let the machine have the last word — that's still your job.