Prompt engineering: Is being an AI 'whisperer' the job of the future or a short-lived fad?

August 28 2023, by Cameron Shackell

Credit: AI-generated image ([disclaimer](#))

As generative AI settles into the mainstream, growing numbers of [courses](#) and [certifications](#) are promising entry into the "[hot job](#)" of [prompt engineering](#).

Having skills in using natural language (such as English) to "prompt" useful content out of AI models such as ChatGPT and Midjourney seems like something many employers would value. But is it as simple as doing a short course and riding the wave to a six-figure salary?

The prompt engineering hype

A Washington Post article published in February did a lot to seed the notion that prompt engineers are "AI whisperers" who "program in prose." It dropped some big salary numbers and quoted a job ad by Silicon Valley company Anthropic calling for people who have "a creative hacker spirit and love solving puzzles."

Similar articles in Time, Forbes and Business Insider further fueled the frenzy.

And to complete the transition from geek to chic, several influencers jumped on board to portray prompt engineering as a gold rush open for anyone willing to study and learn a few tricks.

Are there really that many jobs?

That Anthropic ad is still hanging around. Six months later, it seems more like a corporate publicity stunt than a search for talent.

As many commentators predicted, prompt engineering hasn't exploded as a standalone career. At the time of writing this article, there wasn't a single advertisement for a "prompt engineer" role on the main job sites in Australia. And only four listings mentioned prompt engineering in the job description.

The situation seems better in the United States. But even there, the new

profession has largely been subsumed into other roles such as machine learning engineer or AI specialist.

There are few reliable statistics on the growth (or lack of growth) in prompt engineering. Most data are anecdotal. The reality is further clouded by consulting firms such as Deloitte promoting it as "the dawn of a new era" as part of their AI business drive.

What's the reality?

A lot of the confusion about whether prompt engineering is useful comes from not recognizing that there are two different types of value creators: domain experts and technical experts.

Domain experts

The germ of truth in the "anyone can do it" narrative is that experts in a particular subject are often the best prompters for a defined task. They simply know the right questions to ask and can recognize value in the responses.

For example, in branding and marketing, generative AI is taking off for what I have dubbed generic or "G-type" creative tasks (such as making the Pepsi logo in the style of Picasso). When advertising experts start hacking away at prompting, they quickly invent ways to do things even the most skilled AI gurus can't. That's because technical gurus often don't know much about copyrighting or marketing.

Technical experts

On the other hand, tech gurus who grapple "under the hood" with the enormous complexity of AI models can also add value as prompt

engineers. They know arcane things about [how AI models work](#).

They can use that knowledge, for example, to improve results for everyone using AI to obtain data from a company's internal documents. But they typically have little domain knowledge outside of AI.

Both domain expert and technical expert prompt engineers are valuable, but they have different skill sets and goals. If an organization is using generative AI at scale, it probably needs both.

Why is prompting hard?

Generative AI ultimately produces outputs for people. Advertising copy, an image or a poem is not useful or useless until it succeeds or fails in the [real world](#). And in many real-world scenarios, domain experts are the only ones who can judge the usefulness of AI outputs.

Nonetheless, these evaluations are ultimately subjective. We know 2 + 2 = 4. So it's simple to test prompts that stop AI from [hallucinating](#) that [the answer is 5](#). But how long does it take to work out if an AI-designed ad campaign is more or less effective than a human-designed one (even if you do have a domain expert on hand)?

In [my past research](#), I have suggested the evaluation of generative AI should move [closer to semiotics](#)—a field that can connect [natural language](#) to the real world. This could help narrow the evaluation gap over time.

Is prompt engineering worth learning?

Beyond playing with some [tips and tricks](#), formally learning how to write prompts seems a bit pointless for most people. For one thing, AI models

are constantly being updated and replaced. [Specific prompting techniques](#) that work now may only work in the [short term](#).

People looking to get rich from prompt engineering would be better advised to focus on pairing AI and [problem formulation](#) in their area of expertise. For example, if you're a pharmacist you might try using generative AI to double check warning labels on prescriptions.

Along the way you'll sharpen your [expository writing](#), acquire the basic generative AI skills (which employers might appreciate), and *maybe* strike gold with a killer application for the right audience.

Eventually, boasting that you know how to prompt AI will become resumé furniture. It will be comparable to boasting you know how to use a search engine (which wasn't always [so intuitive](#))—and may paint you as a dinosaur if mentioned.

This article is republished from [The Conversation](#) under a Creative Commons license. Read the [original article](#).

Provided by The Conversation