

SDS PODCAST

EPISODE 718

FIVE-MINUTE

FRIDAY:

**CHATGPT CUSTOM
INSTRUCTIONS: A
MAJOR, EASY HACK
FOR DATA
SCIENTISTS**



(00:05):

This is Five-Minute Friday on a powerful ChatGPT Custom Instructions hack.

(00:19):

Welcome back to the Super Data Science Podcast. I'm your host, Jon Krohn. At my machine learning company, Nebula, I have the great pleasure of working with a number of world-class data scientists. One of those whose work I've mentioned on the show before is Shaan Khosla. With both proprietary and open-source large language models, Shaan has become quite a prompt engineer, an LLM whisperer, if you like. Recently he tipped me off to a way that he's been leveraging ChatGPT's Custom Instructions to automatically get improved responses from OpenAI's juggernaut GPT-4 model. And although I haven't tested it thoroughly, probably their GPT-3.5 model as well. If you haven't used Custom Instructions in ChatGPT before, here's how you do it. And you can watch the YouTube version of this episode if you'd like to be able to see this hands-on demo, but it's pretty simple and you'll probably understand just through an audio-only description.

(01:16):

So when you're in ChatGPT, you go down to where your email address is in the bottom left-hand corner. There's kind of three dots down there, and you select Custom Instructions from the options there. And then there's, yeah, these Custom Instructions you can fill in. So there's a box on what would you like ChatGPT to know about you to provide better responses. You can optionally fill that with information about you if you'd like. But for the purposes of today's podcast episode, the hack doesn't require that first field. Instead, we're just going to make use of the second field, which is how would you like ChatGPT to respond? So if you haven't used this before,



here's the hack. You can use this prompt that Shaan put together himself in this field.

(02:06):

And so it says, "I need you to help me with a task. To help me with a task, first, come up with a detailed outline of how you think you should respond. Then critique the ideas in this outline, mention the advantages, disadvantages in ways it could be improved. Then use the original outline and the critiques you made to come up with your best possible solution. Overall, your tone should not be overly dramatic. It should be clear, professional and direct. Don't sound robotic or like you're trying to sell something. You don't need to remind me, you're a large language model. Get straight to what you need to say to be as helpful as possible. Again, make sure your tone is clear, professional and direct, not overly like you're trying to sell something." So that's his prompt. I've got it for you in the show notes so that you can copy and paste it yourself and derive it to your own needs. Even feel free to share with me on social, if you come up with some derivations of this or you have your own kind of trick way of using these Custom Instructions to get awesome results.

(03:10):

So I'm always available. Tag me in a LinkedIn post or a public Twitter post. I'm always happy to have these kinds of conversations in public so that everyone can enjoy and add on. And yeah, so in case it isn't obvious, so the first part of this prompt that Shaan put together, the part about breaking the task down into parts and critiquing the ideas. This allows the model to have an iteration over its kind of initial inclinations, the original kind of output that it prints out and it allows it to come up with, in many circumstances, a much better output than its original inclination by doing this kind of self-critique to come up with the best solution. So that's why we have the first part in there.



(04:02):

And then we have the second part in there because people who are listening to this podcast, data scientists, machine learning engineers, AI researchers, software engineers, other probably technically minded people, we don't need ChatGPT to kind of have its regular persona, which can be maybe a little bit bubbly. It's telling us about how it's a large language model and therefore it's kind of hedging its bets on what it's saying. We get that we know that, we're professionals, so just give us something clear, professional and direct. And yeah, we'll get more, at least in my case, I'm getting more of the kinds of results that I'm looking for, maybe you will as well.

(04:40):

All right, that's it for today's episode. I hope this ChatGPT Custom Instructions hack reaps dividends in your daily workflows. Until next time, my friend, keep on rocking it out there and I'm looking forward to enjoying another round of the Super Data Science Podcast with you very soon.