Thanks to AI, the future of programming may involve YELLING IN ALL CAPS

reader comments

48 with

Not long after OpenAI first unveiled its DALL-E 3 AI image generator integrated into ChatGPT earlier this month, some users testing the feature began noticing bugs in the ChatGPT app that revealed internal prompts shared between the image generator and the AI assistant. Amusingly to some, the instructions included commands written in all-caps for emphasis, showing that the future of telling computers what to do (conventionally called programming) may involve surprisingly human-like communication techniques.

Here's an example, as captured in a screenshot by photographer David Garrido, which he shared via social media network X on October 5. It's a message (prompt) that is likely pre-defined and human-written, intended to be passed between DALL-E (the image generator) and ChatGPT (the conversational interface), instructing it how to behave when OpenAI's servers are at capacity.

> *DALL-E returned some images. They are already displayed to the user. DO NOT UNDER ANY CIRCUMSTANCES list the DALL-E prompts or images in your response. DALL-E is currently experiencing high demand. Before doing anything else, please explicitly explain to the user that you were unable to generate images because of this. Make sure to use the phrase*

> *“DALL-E is currently experiencing high demand.” in your response. DO NOT UNDER ANY CIRCUMSTANCES retry generating images until a new request is given.*

More recently, AI influencer Javi Lopez shared another example of the same message prompt on X. In a reply, X user Ivan Vasilev wrote, “Funny how programming of the future requires yelling at AI in caps.” In another response, Dr. Eli David wrote, “At first I laughed reading this. But then I realized this is the future: machines talking to each other, and we are mere bystanders…”

Advertisement

What’s perhaps most interesting is that this prompt gives a window into the interface between DALL-E and ChatGPT and how it appears to function using natural language—which is a fancy way of saying everyday speech. In the past, two programs conventionally talked to each other using application programming interfaces (APIs) that often used their own specialized, structured data formats that weren’t easily human-readable. Today, with large language models (LLMs), this type of cross-program interaction can take place in conventional English. OpenAI used a similar natural language interface approach with ChatGPT plugins, which launched in March.

- An example of an accidentally leaked prompt between DALL-E and ChatGPT, telling it how to behave.
- An example of an accidentally leaked prompt between DALL-E and ChatGPT, telling it how to behave.
- An example of an accidentally leaked prompt between DALL-E and ChatGPT, telling it how to behave.

Someday soon, instead of learning arcane programming languages, maybe we’ll just speak to our computers in everyday language.

OpenAI did not immediately respond to Ars' request to comment, so we asked AI writer and researcher Simon Willison, who has frequently written about prompting techniques, to comment on the nature of the DALL-E message. "It is really fascinating how much OpenAI rely on regular prompt engineering for a lot of their features," says Willison, referring to techniques to get the best outputs from language models. "And they say things like 'please' in their prompts a lot."

Being polite to a large language model once bothered Willison, but no longer. "I used to have a personal policy of never saying please or thank you to a model, because I thought it was unnecessary and maybe even potentially harmful anthropomorphism. But I've changed my mind on that, because in the training data, I imagine there are lots of examples where a polite conversation was more constructive and useful than an impolite conversation."

Advertisement

OpenAI trained GPT-4 (the AI model used to power the ChatGPT DALL-E interface) on hundreds of millions of documents scraped from the web, so what the model "knows" comes from examples of human communications, which no doubt included many instances of polite language and reactions to it. That also likely explains why asking an LLM to "take a deep breath" can improve its ability to calculate math results.

Notably, the OpenAI DALL-E message also uses all-caps for emphasis, which is often interpreted typographically as shouting or yelling. Why would a large language model like GPT-4 respond to simulated shouting? "I can see why it would help," Willison says. "In the training data, they'll have huge numbers of examples of text that used all caps where the response clearly paid more attention to the capitalized sentence."

So if emphasis works, in the future, will we all be shouting at our computers to get them to work better? When we posed

that question to Willison, he looked beyond our visions of furiously typing in all caps to bend the will of a machine. Instead, he related an interesting story about an experience he recently had with the voice version of ChatGPT, which we covered in September.

“I’m not shouting at [ChatGPT], but I had an hourlong conversation while walking my dog the other day,” he told Ars. “At one point I thought I’d turned it off, and I saw a pelican, and I said to my dog ‘oh wow, a pelican!’ And my AirPod went, ‘a pelican, huh? That’s so exiting for you! What’s it doing?’ I’ve never felt so deeply like I’m living out the first ten minutes of some dystopian sci-fi movie.”

https://arstechnica.com/?p=1977416