

With encouragement, large language models devise more efficient prompts

September 21 2023, by Peter Grad



An overview of the OPRO framework. Given the meta-prompt as the input, the LLM generates new solutions to the objective function, then the new solutions and their scores are added into the meta-prompt for the next optimization step. The meta-prompt contains the solution-score pairs obtained throughout the optimization process, as well as a natural language description of the task and (in prompt optimization) a few exemplars from the task. See Figure 3 for a sample meta-prompt for prompt optimization. Credit: *arXiv* (2023). DOI: 10.48550/arxiv.2309.03409

One of the principal drivers of efficient large language model (LLM)



tasks is the prompt.

To be most effective, a prompt must be clear and well-tailored to the task.

Researchers devote significant resources to ensure that prompts are optimized to yield the best outcome. Improper use of keywords, awkward phrasing, vague instructions or lack of appropriate context can degrade the quality of results.

Computer programmers are always trying to craft better ways to formulate prompts. Researchers at Google's DeepMind recently considered a novel approach: What if large language models helped construct prompts?

They came up with a process called OPRO, Optimization by PROmpting.

In a paper, published Sept. 7 on the pre-print server *arXiv*, DeepMind researcher Chengrun Yang explained OPRO is "a simple and effective approach" to assign optimization tasks to LLMs in natural language.

"In each optimization step," Yang said, "the LLM generates new solutions from the prompt that contains previously generated solutions with their values, then the new solutions are evaluated and added to the prompt for the next optimization step."

Such iterative solutions are commonly used in optimization tasks, but the formulation has generally been devised by humans relying heavily on mathematical models.

OPRO capitalizes on LLMs' novel ability to understand natural language instructions.



It creates prompts, clearly defining the challenge, and provides examples of similar problems and instructions for an iterative approach to a solution. That is, as the LLM proposes solutions for each step in the optimization process, the prompt is modified to incorporate those results. The process is repeated until an optimal solution is reached.

"Optimization with LLMs enables quick adaptation to different tasks by changing the problem description in the prompt, and the optimization process can be customized by adding instructions to specify the desired properties of the solutions," said Yang.

The researchers tested their approach on two popular types of challenges: linear regression and the <u>traveling salesman problem</u>. The results were promising, but with an added touch—they found significant improvement.

The linear approach is a statistical model displaying a relationship between text-based and numeric variables. It can be used in financial forecasting, for example, by predicting <u>stock prices</u> based on news reports from Wall Street, or it can recommend Netflix movies based on a user's reviews of programming.

The traveling salesman scenario is a classic <u>optimization</u> problem that provides a list of cities and then determines the shortest and fastest route a salesman would need to take to visit each city without repeat.

OPRO performed admirably. It achieved results "on par with some handcrafted heuristic algorithms," Yang said.

"But with an extra boost, optimized prompts outperform[ed] humandesigned prompts ... by a significant margin, sometimes over 50%," Yang said.



What was the extra boost?

Encouragement.

The DeepMind team discovered that when phrases expressing encouragement were attached to prompts, better results were achieved.

Such phrases included, "Take a deep breath and work on this problem step-by-step," "Let's work this out in a step-by-step way to be sure we have the right answer," and "Let's calculate our way to the solution."

The researchers did not elaborate on why such supportive expressions yielded better results, though it may be assumed LLMs were trained on data containing numerous instances of the expressions associated with careful examination and processing of relevant data.

"Optimization is ubiquitous," Yang said. "While derivative-based algorithms have been powerful tools for various problems, the absence of gradient imposes challenges on many real-world applications... With the advancement of prompting techniques, LLMs have achieved impressive performance on a variety of domains."

More information: Chengrun Yang et al, Large Language Models as Optimizers, *arXiv* (2023). DOI: 10.48550/arxiv.2309.03409

© 2023 Science X Network

Citation: With encouragement, large language models devise more efficient prompts (2023, September 21) retrieved 11 May 2024 from <u>https://techxplore.com/news/2023-09-large-language-efficient-prompts.html</u>

This document is subject to copyright. Apart from any fair dealing for the purpose of private



study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.