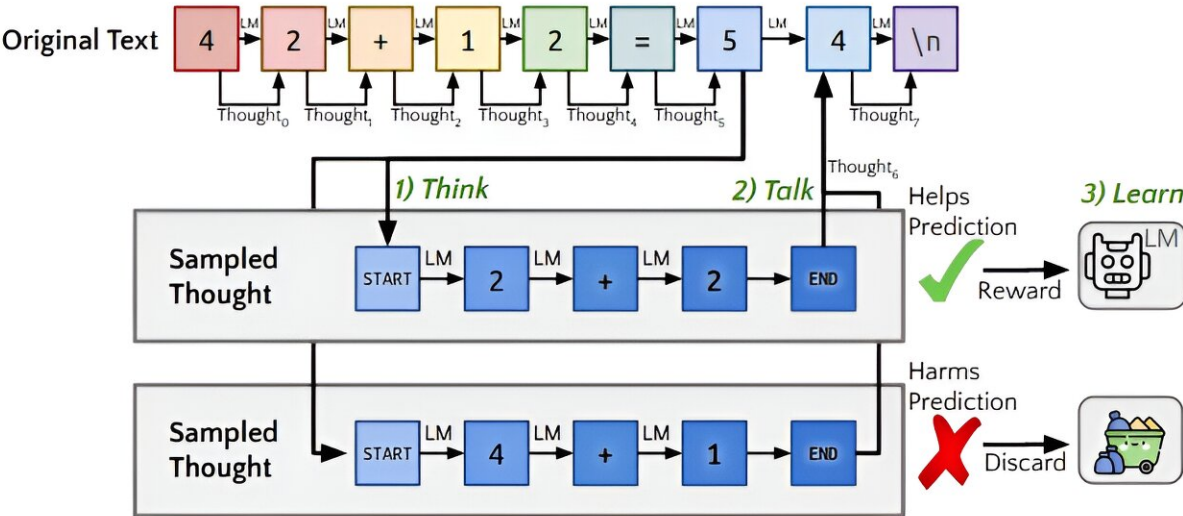


Quiet-STaR algorithm allows chatbot to think over its possible answer before responding

March 21 2024, by Bob Yirka



Quiet-STaR. We visualize the algorithm as applied during training to a single thought. We generate thoughts, in parallel, following all tokens in the text (think). The model produces a mixture of its next-token predictions with and without a thought (talk). We apply REINFORCE, as in STaR, to increase the likelihood of thoughts that help the model predict future text while discarding thoughts that make the future text less likely (learn). Credit: *arXiv* (2024). DOI: 10.48550/arxiv.2403.09629

A collaboration between AI researchers at Stanford University and Notbad AI Inc. has resulted in the development of an algorithm that allows current chatbots to mull over possible responses to a query before giving its final answer. The team has published a [paper](#) on the *arXiv* preprint server describing their new approach and how well their algorithm worked when paired with an existing chatbot.

As the researchers note, the general approach taken by current chatbots is to develop an answer to a query posed by a human using training data. None of the chatbots currently being used by the public stop to ponder multiple possible answers to a query before giving the one it thinks is most likely to be what the human wanted. If a human responded in such a fashion, it would be described as simply blurting out an answer.

In this new study, the research team has given chatbots a means for mulling a bit before answering, and in so doing, claim to have created a way for chatbots to be much more accurate—and to answer questions in more human-like ways.

The algorithm, Quiet-STaR, works by first asking the chatbot to produce multiple answers to a given query. It compares the answers with the original query to decide which appears to be the best. It then directs the chatbot to return that answer to the user. The team also gave the algorithm the ability to learn from its own work, thereby improving its mulling capabilities over time.

To test their algorithm, the researchers added it to the open-source Mistral 7B [chatbot](#) and tested it using a standard reasoning test—it scored 47.2%. Without the algorithm, Mistral 7B scored just 36.3%. It also did much better on a math test.

The research team notes that their [algorithm](#) could be plugged into any of the chatbots currently in use, though it would have to be done by their makers, a move they suggest could improve the accuracy of chatbots in general.

More information: Eric Zelikman et al, Quiet-STaR: Language Models Can Teach Themselves to Think Before Speaking, *arXiv* (2024). [DOI: 10.48550/arxiv.2403.09629](https://doi.org/10.48550/arxiv.2403.09629)

© 2024 Science X Network

Citation: Quiet-STaR algorithm allows chatbot to think over its possible answer before responding (2024, March 21) retrieved 27 April 2024 from <https://techxplore.com/news/2024-03-quiet-star-algorithm-chatbot.html>

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.