

Peer Review

Review of: "Transformers Can Navigate Mazes With Multi-Step Prediction"

Niranjan Sitapure¹

1. Texas A&M University - College Station, United States

This work is a great way to show the impact of different "objectives" in training LLMs for other tasks (i.e., next-token vs multi-step prediction), and the results from the different Maze structures highlight the differences very well. A couple of distinct questions that may aid future work or brainstorming on the topic:

1. Given that RLs are very prodigious at Maze / strategy games / Go type of problems, can somehow RL be incorporated into either the training methodology in addition to the diffusion process or in the objective function (MLM U)?
2. How does the LLM handle the first prediction / cold start? Does improving the cold start by 1% increase overall accuracy by say 10%; hypothesizing this because decisions/turns at the start of the Maze are seemingly more important than down the line in the Maze path.

Declarations

Potential competing interests: No potential competing interests to declare.