Learning & Adaptive Systems

ETH zürich

Improving the Ability of Large Language Models

Project Proposal for Master Theses / Semester Projects

Motivation

This project aims to advance the reasoning capabilities of Large Language Models (LLMs), enabling them to handle increasingly complex tasks that require deep, structured thinking. Current LLMs have demonstrated [Brown et al., 2020, Wei et al., 2022, Bubeck et al., 2023] strong performance in diverse domains, but there is significant potential to enhance their ability to think systematically before responding [Goyal et al., 2023], improving accuracy and problem-solving across a wide range of applications. The availability of strong open-weight LLMs provides a strong foundation to pursue new ideas to improve their performance.

Projects

We offer various topics aimed at improving LLMs' abilities, such as selecting informative data for targeted instruction tuning, estimating uncertainty in long-form generation, automatic model selection via mixtures, computationally efficient context windows, and dynamic inference for systematic thinking.

Contact

Contact us to discuss these projects in detail.
Jonas Hübotter (jonas.huebotter@inf.ethz.ch) & Ido Hakimi (ido.hakimi@ai.ethz.ch).

References

[Brown et al., 2020] Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. *arXiv preprint ArXiv:2005.14165*.

[Bubeck et al., 2023] Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y., Lundberg, S., et al. (2023). Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*.

[Goyal et al., 2023] Goyal, S., Ji, Z., Rawat, A. S., Menon, A. K., Kumar, S., and Nagarajan, V. (2023). Think before you speak: Training language models with pause tokens. *arXiv preprint arXiv:2310.02226*.

[Wei et al., 2022] Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q. V., Zhou, D., et al. (2022). Chain-of-thought prompting elicits reasoning in large language models. In *NeurIPS*.