ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

Prof. R. Wattenhofer

Abstract Reasoning Corpus (ARC) Challenge

The Abstract Reasoning Corpus (ARC) Challenge represents a significant benchmark in the development of AI systems capable of human-like reasoning. The challenge requires models to solve tasks that demand abstract thinking, a skill that current AI struggles to master. State-of-the-art solutions have achieved only 46% accuracy on this challenge, compared to 85% accuracy by humans (see current leaderboard here). Even the most advanced LLMs fall short of human performance, challenging the notion that artificial general intelligence (AGI) exists and underscoring that LLMs are merely *next-word predictors* rather than genuine *reasoning agents*.

As part of your project, you will develop novel approaches to tackle the ARC Challenge, contributing to a field that lies at the core of next-generation AI. Successful advancements in this area not only push the boundaries of AI capabilities but also come with the potential to earn a half-million-dollar prize for solutions that surpass human performance. So let's get to work!

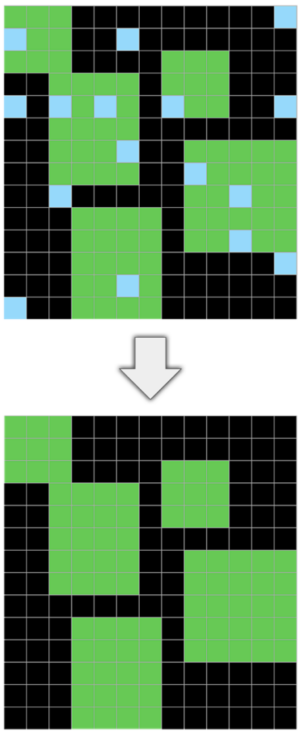

Figure 1: ARC puzzle example.

You are encouraged to bring your own ideas to tackle this problem. Additionally, we recommend exploring unsupervised reinforcement learning (skill discovery) and hierarchical planning [1], as these techniques appear promising for tasks of this nature.

Requirements Strong motivation, ability to work independently, and interest in conducting exciting theoretical and/or empirical research. Solid background in mathematics and artificial intelligence (especially reinforcement learning). Good programming skills in python (and ideally with libraries such as `pytorch` and `gym`). A deep interest in what constitutes *intelligence* at its core is a substantial plus.

Interested? Please get in touch with us for more details!

Contact

- Frédéric Berdoz: `fberdoz@ethz.ch`, ETZ G60.1

References

[1] Shubham Pateria, Budhitama Subagdja, Ah-hwee Tan, and Chai Quek. "Hierarchical Reinforcement Learning: A Comprehensive Survey." *ACM Computing Surveys*, vol. 54, no. 5, article 109, June 2022, pp. 1-35. https://doi.org/10.1145/3453160.