




ARC - Turing Test for AI?

The Abstraction and Reasoning Corpus (ARC) is like an IQ test for machine intelligence, allowing us to quantify progress towards human-level AI. It was introduced in 2019 by François Chollet as a [Kaggle](#) competition. Rather than focusing on a single specific challenge, the aim of ARC is to provide an ensemble of challenges to measure the ability to handle and adapt to new unseen environments. One task of the ARC usually consists of three given samples with input and output pairs. For a fourth sample, only the input is given and the output must be produced. These challenges seem easy for humans to solve, however machines, at least for now, struggle with them a lot. The implicit assumption is that we humans can rely on existing "priors" and knowledge that help us tackle these tasks. Therefore, a central aspect of ARC is to make these priors more explicit and try to incorporate them into learning systems.

Over the past years, Large Language Models have seen a large increase in their capabilities and reach near human level performance across various standardized benchmarks. However, the ARC currently seems beyond their reach - [for now ...](#)

In previous work, we have developed an approach which tries to include existing priors and knowledge to mimic human intuition and leverage the expertise of pretrained models, i.e. large language models. The naive application of this approach already yields promising results, which begs the question: What strategies exist to boost our baseline? Could finetuning, thinking step-by-step or even prompt engineering yield benefits? Moreover, over the past year many simpler versions of ARC have popped up such as mini-ARC, Concept-ARC, 1D-ARC. Do these benchmarks make sense? What do they assess and how does our baseline fare on it?



Requirements: Strong motivation, knowledge in deep learning, or a solid background in machine learning, Python and libraries such as TensorFlow or PyTorch. Experience with language models is an advantage. We will have weekly meetings to discuss open questions and determine the next steps.

Interested? Please contact us for more details!

Contact

In a few short sentences, please tell us why you are interested in the project and about your coding and machine learning background (i.e., your own projects or courses).

- Joël Mathys: jmathys@ethz.ch, ETZ G63
- Benjamin Estermann: besterma@ethz.ch, ETZ G60.1