




Abstraction and Reasoning Challenge

Can a computer learn complex, abstract tasks from just a few examples?



Current machine learning techniques are data-hungry and brittle—they can only make sense of patterns they’ve seen before. Using current methods, an algorithm can gain new skills by exposure to large amounts of data, but cognitive abilities that could broadly generalize to many tasks remain elusive. This makes it very challenging to create systems that can handle the variability and unpredictability of the real world, such as domestic robots or self-driving cars.

However, alternative approaches, like inductive programming, offer the potential for more human-like abstraction and reasoning. The Abstraction and Reasoning Corpus (ARC) provides a benchmark to measure AI skill-acquisition on unknown tasks, with the constraint that only a handful of demonstrations are shown to learn a complex task. It provides a glimpse of a future where AI could quickly learn to solve new problems on its own. The Kaggle Abstraction and Reasoning Challenge invites you to try your hand at bringing this future into the present!

This competition is hosted by François Chollet, creator of the Keras neural networks library. Chollet’s paper on measuring intelligence provides the context and motivation behind the ARC benchmark.

In this project, we want to apply the recent methods on AI that can solve reasoning tasks it has never seen before. Each ARC task contains 3-5 pairs of train inputs and outputs, and a test input for which you need to predict the corresponding output with the pattern learned from the train examples.

If successful, you’ll help bring computers closer to human cognition and you’ll open the door to completely new AI applications!

Requirements: Prior experience or a strong interest in programming. Creativity skills are advantageous.

Interested? Please contact us for more details!

Contact

- Ard Kastrati: kard@ethz.ch, ETZ G61.3