



## Reinforcement Learning Benchmark on Puzzles

Ten years ago DeepMind published their work on playing Atari games using Reinforcement Learning. Since then, machines have outmatched humans in games such as chess, Go, Shogi, Poker and DOTA. However, there are still games which have not been mastered by Machine Learning.

We are interested in building a challenging benchmark for Reinforcement Learning. The benchmark will be based on logic puzzles. As the name suggests, these puzzles can be solved elegantly using logical reasoning instead of brute force trial and error.

In this thesis, we will build a Python wrapper around an existing puzzle collection to make it compatible with Reinforcement Learning frameworks such as OpenAI Gym. The goal is to obtain an environment with many challenging puzzles which can be used by researchers to test their Reinforcement Learning algorithms against.

**Requirements:** Knowledge in Python, C/C++ and Reinforcement Learning. Experience with PyTorch (or TensorFlow), PyGame and OpenAI Gym is an advantage.

We will have weekly meetings to address questions, discuss progress and think about future ideas.

### Contact

- Luca Lanzendörfer : [lanzendoerfer@ethz.ch](mailto:lanzendoerfer@ethz.ch), ETZ G93
- Benjamin Estermann: [besterma@ethz.ch](mailto:besterma@ethz.ch), ETZ G60.1

