Karlsruhe Institute of Technology

Autonome Lernende Roboter (ALR)
Prof. Gerhard Neumann

Project Type

- ☐ Master Thesis
- ☑ Bachelor Thesis
- ☐ Research Project

Supervisors

☞ Philipp Becker
@ philipp.becker@kit.edu

Difficulty

Algorithmic

Math

Application

Tackling Offline Reinforcement Learning by Model Learning

Description

Many recent successes of Deep Learning techniques are enabled by an abundance of data, such as text, which is available in a virtually unlimited amount via the internet. The goal of offline Reinforcement Learning (RL)[3] is to facilitate data, which was previously collected by any agent, to learn an RL agent without interaction with the environment. One possible approach to this problem is learning a world model, i.e., a model of the environment the agent is supposed to act in. Such a model can then be used for planning and control. This approach is also the underlying paradigm of Model-Based Reinforcement-Learning (MBRL), which, however, assumes the classical online RL setting of an agent that interactively collects new data.

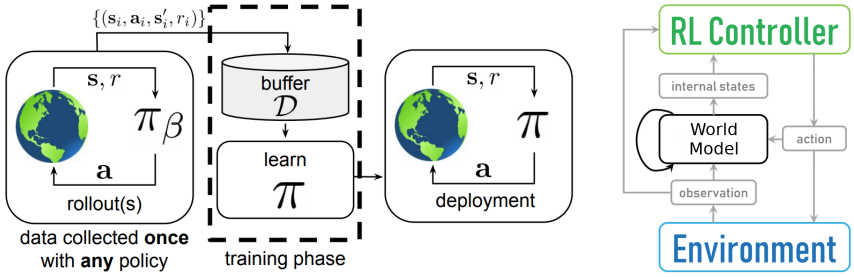

Figure 1: Left: The idea of offline RL is to use previously collected data without further environment interaction. Right: In model-based RL we learn a world model which serves as the basis for an RL controller, yet most current approaches assume an online RL setting, i.e., the data is collected by the agent itself during training. Here we are going to investigate how we can use offline data for this instead.

Here we are going to investigate how recent MBRL approaches[2] can be used to learn world models from offline data. The key challenge is obtaining models that are still well suited for planning and control and we will investigate how to include common offline RL wisdom into the approaches to improve their performance in the offline RL setting.

Tasks

- Get familiar with the general ideas behind offline and model-based RL, as well as, the particular MBRL methods used.
- Evaluate the usage of MBRL methods for recent offline RL benchmarks [1].
- Include common ideas from offline RL into the approach to get an approach that's better tailored to the challenges of offline RL.

References

[1] Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4rl: Datasets for deep data-driven reinforcement learning, 2020.

[2] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019.

[3] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.