Karlsruhe Institute of Technology

Autonome Lernende Roboter (ALR)

Prof. Gerhard Neumann

Project Type

- [x] Master Thesis
- [ ] Bachelor Thesis
- [ ] Research Project

Supervisors

Philipp Becker

philipp.becker@kit.edu

Difficulty

Algorithmic

Math

Application

Improving Uncertainty Estimation for Model-Based Reinforcement-Learning

Description

The main idea of Model-Based Reinforcement-Learning (MBRL) is to learn a model of the environment and its dynamics, which subsequently can be used for planning and control. The resulting approaches are usually very sample efficient and, once a good model is learned, it can be used to tackle multiple tasks. Many more traditional approaches emphasize the importance of the model appropriately representing uncertainty. Here, we differentiate between uncertainty in model parameters, i.e., uncertainty caused by lack of data (*epistemic*-uncertainty) and uncertainty induced by the ambiguity of the data (*aleatoric*-uncertainty).
One way of formalizing the system dynamics, especially in a partially-observable setting, are state-space models. Those are the back-bone of a recent line of works [1], Yet, this line of research also mostly ignored the role of uncertainty representation, working with sub-optimal assumptions and simplifications.

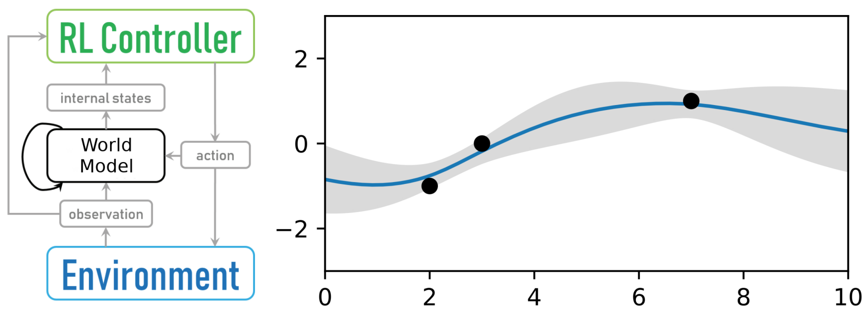

Figure 1: In model-based RL we learn a world model which serves as the basis for an RL controller. Capturing the uncertainty of the world model is crucial for success.

In recent work we disentangled the modeling of alleatoric and epsitemic uncertainty in State Space Model approaches for MBRL. This resulting approach builds on well-understood, theoretically founded components to address the different types of uncertainty. In this thesis we now want to build on this by investigating different ways of improving upon those individual components

Tasks

- Getting familiar with model-based RL, State-Space Models, and how to represent epistemic and aleatoric uncertainty.
- Implement different ways of representing the epistemic uncertainty, such as ensembles [2], variational approximations, or others.
- Implement different ways of representing the aleatoric uncertainty, e.g., by working with a discrete state-space model as proposed in [1]).
- Evaluate everything on standard benchmarks and benchmarks, where appropriate uncertainty representation is crucial.

References

[1] Danijar Hafner, Timothy Lillicrap, Mohammad Norouzi, and Jimmy Ba. Mastering atari with discrete world models. *arXiv preprint arXiv:2010.02193*, 2020.

[2] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.