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
Michael Volpp
michael.volpp@kit.edu

Difficulty

Algorithmic

Math

Application

Variational Bayesian Deep Learning for Model Predictive Control

Description

Model predictive control (MPC) with learned dynamics models is a popular approach to model-based reinforcement learning. Several recent works (e.g., [1]) emphasize the importance of capturing model uncertainty of the learned dynamics. There are multiple approaches to this, all relying on a Bayesian formulation of the learning problem. On the one hand, variational inference approaches have been used to learn Gaussian approximations of the true posterior over parameters. On the other hand, using an ensemble of several neural networks has also proven to be effective. In fact, recent work [2] argues that a combination of both yields the best results.

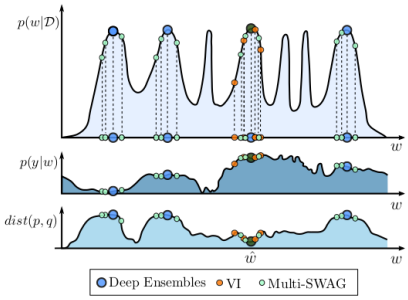

(a) The benefits of variational ensembles for Bayesian deep learning [2].

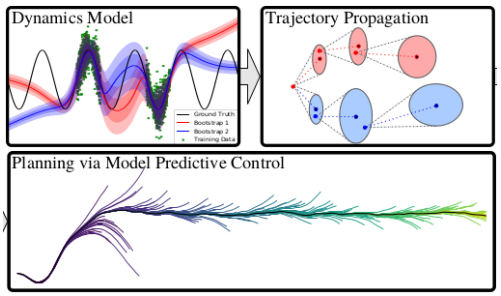

(b) Schematic overview of a model predictive control control approach [1].

This thesis aims at evaluating several variational approaches for Bayesian deep learning in the context of dynamics learning for MPC. The approaches should be evaluated using a single variational approximation as well as ensembles of them.

Tasks

- Get to know and implement several approaches for variational Bayesian deep learning,
- Evaluate and compare the approaches for learning dynamics models for MPC on a set of benchmarks.

Qualifications

- Background in Computer Science, Mathematics, Physics, or similar,
- Good mathematical understanding,
- Experience with programming in Python (PyTorch is a plus).

References

[1] Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. *arXiv preprint arXiv:1805.12114*, 2018.

[2] Andrew Gordon Wilson and Pavel Izmailov. Bayesian deep learning and a probabilistic perspective of generalization. *arXiv preprint arXiv:2002.08791*, 2020.