Universität Stuttgart

INSTITUT FÜR
KOMMUNIKATIONSNETZE
UND RECHNERSYSTEME
Prof. Dr.-Ing. Andreas Kirstädter

Master thesis No. 1083

Designing and Implementing Model Based Reinforcement Learning for Networking Problems

Methods
Simulation
Machine Learning

Topics
Network control
Reinforcement Learning

Background

Reinforcement learning for network control faces strict requirements on sample efficiency, stability, and constraint satisfaction. Model-free methods often require prohibitive interaction time and struggle to generalize across traffic regimes and topologies. Model-based reinforcement learning addresses these limitations by explicitly learning a predictive world model of network dynamics and using this model for planning. Recent approaches such as latent-space world models and imagination-based rollouts have demonstrated that planning over learned dynamics can drastically reduce real-environment interactions. In parallel, planning-centric RL methods that combine learned models with search, such as those inspired by AlphaZero, have shown that coupling value and policy learning with lookahead planning yields robust and high-performing decision-makers. For networking problems with structured dynamics and well-defined constraints, model-based RL with planning is a particularly promising paradigm.

Problem Description

The objective of this student work is to investigate model-based reinforcement learning and planning for dynamic network control. A world model of the network dynamics is learned from interaction or offline data and subsequently used for planning-based decision making. The design is inspired by latent-dynamics models and imagination rollouts as well as search-based policy improvement. The work consists of the following tasks:

- Formulation of the network control problem as a sequential decision process
- Design and training of a predictive world model of network dynamics
- Implementation of planning in model space (e.g., rollout-based or tree search methods)
- Coupling planning with policy and value learning
- Evaluation against model-free RL and static baselines in terms of sample efficiency, performance, and stability

Acquired Knowledge and Skills

The student will gain a deep understanding of model-based reinforcement learning and planning under uncertainty. They will acquire practical experience with learning world models, latent-state representations, and planning-based policy optimization. The work further builds expertise in advanced RL architectures, sample-efficient learning, and dynamic network control, as well as strong skills in experimental design and evaluation of learning-based systems.

Requirements
Programming Experience
Basic Machine Learning Knowledge

Desirable knowledge
Communication Networks Architecture and Design
Neural Networks

Contact

M.Sc. Nicolas Hornek
room 1.402 (ETI II), phone 685-67992, E-Mail nicolas.hornek@ikr.uni-stuttgart.de