Universität Stuttgart
Institut für Systemtheorie und Regelungstechnik
Prof. Dr.–Ing. Frank Allgöwer

Open Thesis (MA)

Data-driven Policy Gradient for the Linear Quadratic Regulator

Description:

Policy gradient (PG) methods are optimization techniques that iteratively refine control policies by directly maximizing performance objectives through gradient-based updates to achieve optimal decision-making. They represent the basis of many widely used algorithms in **reinforcement learning**. This thesis focuses on analyzing an **indirect data-driven policy gradient** method to solve the **linear quadratic regulator** problem for systems with unknown dynamics. The method integrates **recursive least squares (RLS)** for online model estimation and a certainty-equivalent framework to compute policy updates. The main objectives of this thesis include analyzing the robustness of both **gradient descent** and **natural policy gradient** methods. One promising approach consists of modeling the interaction between RLS and PG as an algorithmic dynamical system for system-theoretic analysis. Based on this analysis, the thesis aims to deliver analytical results on the convergence of such iterative schemes and, informed by these, new algorithmic and computation improvements on the state-of-the-art approaches to PG in optimal control and reinforcement learning

Prerequisites:

- Knowledge in *Optimal Control*
- Strong mathematical background
- Programming skills, *Matlab or Python*

Supervisor:

Bowen Song
Room 2.236

Area:

Optimal Control
Learning-based Control
Reinforcement Learning

Properties:

Type: **MA**

30% literature
40% theory
30% simulation

Beginning:

Any time

Further information on www.ist.uni-stuttgart.de/lehre/bama

Aushang vom November 22, 2024