Provably Safe Reinforcement Learning Control of a Quadrotor

Background

Quadrotors are highly versatile tools with applications ranging from hobby photography to transport of medicine to rural areas. Recently, they have also been considered for remote power line inspection, which would allow reducing maintenance costs for grid operators significantly. However, controlling a quadrotor is a difficult task requiring a powerful decision-making tool. Deep Reinforcement Learning (RL) has proven its potential to control complex systems from various domains by leveraging the function approximation capacity of neural networks [8], [6]. However, most state-of-the-art RL algorithms have a significant disadvantage which prevents their deployment beyond simulated environments: They do not provide the possibility to consider any safety specifications. This can lead to critical failures, for example, a quadrotor colliding with a power line. Many different approaches for augmenting RL algorithms with safety guarantees have lately been proposed [3]. Most of these approaches are based on modeling the controlled environment to some level of abstraction. To enable efficient yet valid verification of safety constraints, one can use simple models with uncertain parameters that capture unmodeled behavior. The parameters are identified based on data from the actual system such that the abstract model encloses the behavior of the system (so-called „conformant“ models) [7]. However, the system behavior might change during operation due to unexpected disturbances (e.g., wind gusts), requiring a re-identification of the conformant model. The idea of this thesis is therefore to combine safe reinforcement learning approaches with conformant models which are updated online to ensure continuous and provably safe RL control of a quadrotor.

Description

The goal of this project is to integrate a safe RL controller with a conformant system model of a quadrotor which is updated based on data collected online. The approach will be developed and tested using a simulated quadrotor subject to disturbances similar to the setting in [4]. This setting will be implemented using the safe-control-gym library (for modeling the quadrotor) [2] and Stable Baselines3 (for the RL controller) [1]. The first objective of the project is to set up the basic pipeline for a safe RL controller with an abstract model of the quadrotor that is conformant to some offline collected trajectories.

Quadrotor **[5]**

This basic setup will be used to compare different approaches for the conformant model and evaluate performance of the safe and the unsafe controller. The second part of the project will consist of modeling an external force as an additional and unexpected disturbance acting on the quadrotor. Such a change in the system behavior requires a re-identification of the conformant model, for which the student will implement a computationally efficient framework.

Tasks

- familiarization with conformant model identification and safe reinforcement learning using reachability analysis
- identification of a conformant model of the quadrotor based on real trajectories from a tracking task,
- computation of a robust safe region and safe back-up controller (LQR) for the quadrotor with uncertain dynamics,
- integration of the conformant model and the safe back-up controller with the RL controller,
- modelling of an external disturbance on the quadrotor in the safe-control-gym, and
- development of a pipeline for performing online updates of the conformant model and the robust safe region based on system trajectories collected during the learning process.

Technical University of Munich

Department of Informatics

Chair of Robotics, Artificial Intelligence and Real-time Systems

Supervisor:
Prof. Dr.-Ing. Matthias Althoff

Advisor:
Hannah Markgraf, Michael Eichelbeck

Type:
BA/MA

Research area:
Safe Control of Cyber-Physical Systems

Programming language:
Python, Matlab

Required skills:
Basic understanding of cyber-physical systems, control theory, and Reinforcement Learning

Language:
English, German

For more information please contact us:

Phone:

E-Mail:
hannah.markgraf@tum.de

Website:
https://www.ce.cit.tum.de/

What We Offer

- Research in Machine Learning,
- Weekly meetings with your advisors,
- Flexible start and schedule for the thesis project,
- Thesis topics that will be tailored to your interest, and
- Good coffee in case you want to meet in Garching.

Technical University of Munich

Department of Informatics

Chair of Robotics, Artificial Intelligence and Real-time Systems

References

[1] Stable Baselines3, July 2022. original-date: 2020-05-05T05:52:26Z.

[2] Lukas Brunke, Melissa Greeff, Adam W. Hall, Zhaocong Yuan, SiQi Zhou, Jacopo Panerati, and Angela P. Schoellig. safe-control-gym, January 2021. Issue: 1 Pages: 1–35 Publication Title: Annual Review of Control, Robotics, and Autonomous Systems Volume: 1 original-date: 2021-08-04T22:14:15Z.

[3] Lukas Brunke, Melissa Greeff, Adam W. Hall, Zhaocong Yuan, Siqi Zhou, Jacopo Panerati, and Angela P. Schoellig. Safe Learning in Robotics: From Learning-Based Control to Safe Reinforcement Learning. December 2021.

[4] Jaime F. Fisac, Anayo K. Akametalu, Melanie N. Zeilinger, Shahab Kaynama, Jeremy Gillula, and Claire J. Tomlin. A General Safety Framework for Learning-Based Control in Uncertain Robotic Systems. *IEEE Transactions on Automatic Control*, 64(7):2737–2752, July 2019. Conference Name: IEEE Transactions on Automatic Control.

[5] Farhad Goodarzi. *Geometric Nonlinear Controls for Multiple Cooperative Quadrotor UAVs Transporting a Rigid Body*. PhD thesis, August 2015.

[6] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv:1509.02971 [cs, stat]*, July 2019. arXiv: 1509.02971.

[7] Stefan B. Liu and Matthias Althoff. Reachset Conformance of Forward Dynamic Models for the Formal Analysis of Robots. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 370–376, October 2018. ISSN: 2153-0866.

[8] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, February 2015.