Universität Stuttgart
Institut für Flugmechanik und Flugregelung

Blimp Control with Safe Model-Based Reinforcement Learning

Master Thesis

Background:

A blimp is a lighter-than-air vehicle with a potential to execute long duration missions. Reinforcement learning is one possible approach for the blimp control problem. Blimp dynamics are, however, subject to significant parameter uncertainty due to changes in ambient air properties, such as temperature and pressure. This poses a significant challenge in guaranteeing safety [1] with model-based reinforcement learning during the exploration phases of the method. The goal of this thesis is to establish a safe exploration framework for the blimp control task that enables real-world RL training.

Problem definition:

Dyna-styled Model-based reinforcement learning (MBRL) agent has higher sample efficiency compared to the model-free counterpart. Furthermore, the dynamic model can be used for generating safety bounds as well. In this work, the goal is to derive an optimal control policy using MBRL while guaranteeing the exploration safety using the control barrier function (CBF) or forward reachability set. To accelerate the training, we will first derive a nominal dynamic model and apply the Dyna-styled residual MBRL method to further improve this model, safety bound, and control policy.

Task:

- Implement residual MBRL algorithm, e.g. MBPO.
- Implement safety constraints, e.g. CBF.
- Test if the safety guarantees are achieved during training.

Requirements:

- Interested in Robotic Control and Reinforcement Learning
- Experience with C++, Python, Pytorch
- Good academic performance

If interested, please contact:

Yu-Tang Liu yu-tang.liu@ifr.uni-stuttgart.de

[1] https://arxiv.org/abs/1903.11199

iFR