Universität Stuttgart
Institut für Flugmechanik und Flugregelung

Meta Reinforcement Learning based Blimp Control

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. This work aims at synthesizing optimal blimp control policy that can adapt to the environmental changes, e.g.. different wind distribution, buoyancy, etc.

Problem definition:

Classic optimal control design pipeline usually assumes an accurate dynamic model of the plant. This assumption barely holds due to the high parameter uncertainty of the blimp. As a solution, we propose a model-free deep residual RL controller (DRRL) to learn a safe control policy by interaction. Prior work has shown the potential of such approach. However, due to the low sample efficiency nature of the on-policy RL algorithm, online learning and adaptation remain impossible. In this work, the goal is to investigate if a meta-learning approach can be leveraged to reduce the impact of the environmental changes by limiting the training only to the feature extractor while keeping the control policy frozen.

Task:

- Implement prior on-policy deep residual RL (DRRL) baselines [1].
- Implement an off-policy agent with different feature extractors.
- Test the adaptive properties of different feature extractors.

Requirements:

- Interested in robotic control and reinforcement learning
- Experienced with C++, Python, and Pytorch or Jax
- Good academic performance

If interested, please contact:

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

[1] https://ps.is.mpg.de/publications/liu_iros_22

iFR