Universität Stuttgart
Institut für Flugmechanik und Flugregelung

Design of a Curriculum for Reinforcement Learning-Based Autonomous Soaring

Master‘s Thesis

Background

Exploiting thermal updrafts can enable fixed-wing UAVs to reduce their energy consumption significantly while extending their endurance. At the iFR, we have successfully demonstrated an end-to-end reinforcement learning (RL) approach for autonomous updraft localization and exploitation, i.e. a glider learned how to process sensor inputs and compute control commands in order to detect updrafts and gain altitude within them. Our next goal is to combine this with a path tracking ability. This poses a more complex task, since the glider must now solve a three-fold decision-making problem: It needs to make a trade-off between tracking the path, looking for new updrafts, and exploiting previously detected updrafts.

Problem Definition

Similar to human learning, it was found that RL agents can learn complex tasks better if they are learned as a sequence of simpler, consecutive sub-tasks. For instance, humans usually learn about addition and functions before they learn how to compute integrals. Hence, the goal of this thesis is to develop such curricula for the autonomous soaring task as well. This should include both a hand-crafted curriculum (e.g., (1) learn updraft localization, (2) learn path tracking, (3) learn …) as well as an automatized curriculum learning approach, in which the curriculum gets automatically adjusted based on the agent’s learning progress.

Tasks

- Literature review of curriculum learning (CL) and familiarization with our RL-based autonomous soaring approach
- Comparative analysis of published automatized CL methods and selection of a promising method for our use case
- Design of a hand-crafted curriculum for the autonomous soaring task
- Design and implementation of an automatized CL approach based on the selected method from literature
- Deployment of the two CL approaches in our exiting RL training setup
- Assessment by comparing the resulting training behaviors and agents’ performances among the two CL approaches and to a non-CL baseline

Requirements

- Strong interest in machine learning, basic understanding is required
- Programming experience, Python and PyTorch knowledge is ideal
- Good academic performance

Contact

Christian Gall, M.Sc.
E-Mail: christian.gall@ifr.uni-stuttgart.de
Tel.: +49 711 685-67333

iFR