Universität Stuttgart
Institut für Flugmechanik und Flugregelung

Design of a Temporal Convolutional Neural Network for RL-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 new requirements on the memorization capabilities of the underlying deep neural network, since this represents a long-term correlated decision-making problem: The glider needs to make a trade-off between tracking the path, searching for new updrafts, and exploiting previously detected, memorized updrafts.

Problem Definition

In the past, we have used a Long Short-Term Memory (LSTM) to capture updraft dynamics over time. Currently, we are examining Transformers for this task due to the limited memory timespan of LSTMs. For our use case, Temporal Convolutional Neural Networks (T-CNNs), a network architecture originating from Computer Vision, have the potential to capture long-term dependencies at a lower computational complexity than Transformers. Hence, the goal of this thesis is to design a T-CNN architecture for our RL-based autonomous soaring approach that can efficiently encode a long time series of sensor measurements.

Tasks

- Literature review of T-CNNs and familiarization with our RL-based autonomous soaring approach
- Design of multiple T-CNN variants for our use case (vary network features such as depth, kernel shapes, residual connections, etc.)
- Training the T-CNN variants in our existing RL training setup
- Conducting an ablation study to measure the impact of the varied network features and selection of the best-performing T-CNN design
- Assessment of the T-CNN's training behavior, resulting performance, and computational demand by comparison to our existing approaches

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