Reinforcement-Learning-Based Drone Racing for Nanocopters Using a Self-Developed Motion Capture System

Zemp, Maurice (School: Kollegium St. Fidelis Stans)

Autonomous drone research has become increasingly important, yet there are many burdens which make its access limited. Especially for autonomous drone racing, high-end motion-capture systems and high-performance simulation environments remain mostly closed-source and expensive. Thus, this project addresses this gap by investigating whether reinforcement learning (RL) together with a cheap, self-developed Motion Capture (MoCap) Framework still enables aggressive, high-speed nanocopter drone racing. A high-performance Rust simulation environment was developed to support training the Proximal Policy Optimization (PPO) Algorithm to complete arbitrary sequences of gates. For real-world deployment, a custom Rust MoCap Framework was developed on top of OpenCV to use off-the-shelf infrared cameras to determine the position and orientation of the nanocopter in space. A unified User Interface supports direct Sim2Real deployment. Both the simulation and the tracking system were evaluated on performance and accuracy. The trained policy was able to complete any track in near-optimal time after less than an hour of training. The MoCap achieved an average deviation of 5 mm, with the flight volume being classroom-sized, costing no more than 100 Swiss Francs. In the real world, the nanocopter was able to complete the racing tracks while passing A3-sized gates at speeds of up to 25 km/h. The results demonstrate that RL-based controllers can reliably transfer from simulation to reality while using affordable hardware. This approach allows precise autonomous flight while lowering costs for robotics research and education.

Awards Won:

Third Award of $1,200