Agile Autonomy: High-Speed Flight with On-Board Sensing and Computing

Antonio Loquercio Davide Scaramuzza

Abstract—**In this extended abstract, we present our latest research in learning deep sensorimotor policies for agile vision-based quadrotor flight. In addition, we discuss the open research questions that still need to be answered to improve the agility and robustness of autonomous drones.**

I. Introduction

Quadrotors are among the most agile and dynamic machines ever created. However, developing fully autonomous quadrotors that can approach or even outperform the agility of birds or human drone pilots with only onboard sensing and computing is very challenging and still unsolved. Current state-of-the-art works tackled this problem by splitting the task into a series of consecutive blocks: perception, map building, and planning. Although simple and effective, such an approach typically discards interactions among the different blocks and requires each block to make over-simplifying assumptions. Additionally, due to the presence of sequential processing blocks between sensors and actuators, the time to go from observation to action increases at the cost of agility. In this extended abstract, we summarize our latest research in learning deep sensorimotor policies for agile vision-based quadrotor flight. Learning sensorimotor controllers represents a holistic approach that is more resilient to noisy sensory observations and imperfect world models. Training robust policies requires however a large amount of data. However, we will show that simulation data, combined with randomization and abstraction of sensors' observations, is enough to train policies that generalize to the real world. Such policies enable autonomous quadrotors to fly faster and more agile than what was possible before with only onboard sensing and computation.

II. Deep Drone Racing

Drone racing is a popular sport in which professional pilots fly small quadrotors through complex tracks at high speeds. Drone pilots undergo years of training to master the sensorimotor skills involved in racing. Such skills would also be valuable to autonomous systems in applications such as disaster response or structure inspection, where drones must be able to quickly and safely fly through complex dynamic environments. Developing a fully autonomous racing drone is difficult due to challenges that span dynamics modeling, onboard perception, localization and mapping, trajectory generation, and optimal control. For this reason, autonomous drone racing has attracted significant interest from the research community, giving rise to multiple autonomous drone racing competitions.

The authors are with the Dep. of Informatics, University of Zurich, and Dep. of Neuroinformatics, the University of Zurich and ETH Zurich, Switzerland (http://rpg.ifi.uzh.ch). This work was supported by the National Centre of Competence in Research (NCCR) Robotics through the Swiss National Science Foundation (SNSF) and the SNSF-ERC Starting Grant.

Fig. 1. Autonomous Drone Racing with on-board sensing and computing.

One approach to autonomous racing is to fly through the course by tracking a precomputed global trajectory. However, global trajectory tracking requires to know the race-track layout in advance, along with highly accurate state estimation, which current methods are still not able to provide. Indeed, visual inertial odometry is subject to drift in estimation over time. SLAM methods can reduce drift by relocalizing in a previously-generated, globally-consistent map. However, enforcing global consistency leads to increased computational demands that strain the limits of on-board processing.

Instead of relying on globally consistent state estimates, we deploy a convolutional neural network to identify the next location to fly to, also called waypoints. However, it is not clear a prior what should be the representation of the next waypoint. In our works, we have explored different solutions.

In our preliminary work, best system paper award at the Conference on Robot Learning (CORL) 2018, the neural network predicts a fixed distance location from the drone [1]. Training was done by imitation learning on a globally optimal trajectory passing through all the gates. Despite being very efficient and easy to develop, this approach cannot efficiently generalize between different track layouts, given the fact that the training data depends on the track-dependent globally optimal trajectory.

For this reason, a follow-up version of this work proposed to use as waypoint the location of the next gate [2]. As before, the prediction of the next gate is provided by a neural network.