High-Speed Drone Flight with On-Board Sensing and Computing

Antonio Loquercio, Elia Kaufmann, Yunlong Song, and 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. Agile Flight in the Wild

We have developed an approach to fly a quadrotor at high speeds in a variety of environments with complex obstacle geometry while having access to only onboard sensing and computation. By predicting navigation commands directly from sensor measurements, we decrease the latency between perception and action while simultaneously being robust to perception artifacts, such as motion blur, missing data, and sensor noise. To deal with sample complexity and not endanger the physical platform, we train the policy *exclusively* in simulation. We leverage abstraction of the input data to transfer the policy from simulation to reality [1], [2]. To this end, we utilize a stereo matching algorithm to provide depth images as input to the policy. We show that this representation is both rich enough to safely navigate through complex environments and abstract enough to bridge simulation and reality. Our choice of input representation guarantees a strong similarity of the noise models between simulated and real observations and gives our policy robustness against common perceptual artifacts in existing depth sensors. We train our policy with imitation learning on a priviledged expert and deploy the policy on a physical quadrotor. A qualitative example of flight in the wild is shown in Figure 1.

Fig. 1: Autonomous flight in the wild with on-board sensing and computing.

III. 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.

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

The authors are with the Robotics and Perception Group, Department of Informatics, University of Zurich, and Department of Neuroinformatics, 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 European Union's Horizon 2020 Research and Innovation Programme under grant agreement No. 871479 (AERIAL-CORE) and the European Research Council (ERC) under grant agreement No. 864042 (AGILEFLIGHT).

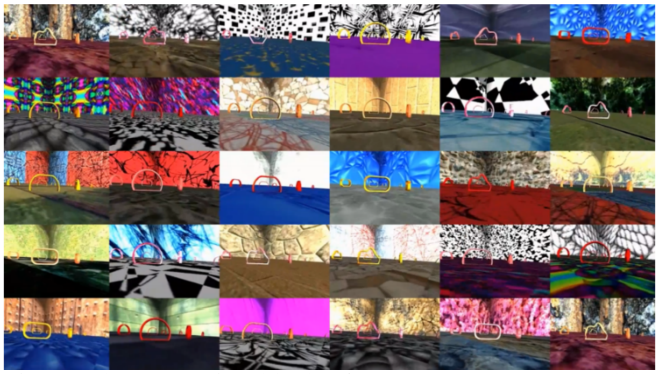

Fig. 2: The perception block of our system, represented by a convolutional neural network (CNN), is trained only with non-photorealistic simulation data.

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 [3]. 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 [4]. As before, the prediction of the next gate is provided by a neural network. However, in contrast to the previous work, the neural network also predicts a measure of uncertainty.

In spite of the representation of the next waypoint, training neural networks requires a large amount of data. Collecting such data is generally a very tedious and time consuming process, which represents a limitation of the two previous works. In addition, when something changes in the environment, or the appearance of the gates changes significantly, the data collection process needs to be repeated from scratch. For this reason, in our most recent work [5] we have proposed to collect data *exclusively* in simulation. To enable transfer between the real and the physical world, we randomized all the features which are unimportant for predictions, *i.e.* illumination, gate shape, floor texture, and background. A sample of the training data generated by this process, called domain randomization, can be observed in Fig. 2. Our approach was the first to demonstrate zero-shot sim-to-real transfer on the task of agile drone flight. A collection of the ideas presented in the above works has been used by our team to compete in the alpha-pilot competition [6].

Fig. 3: A quadrotor performs a Matty Flip. The drone is controlled by a deep sensorimotor policy, and uses only onboard sensing and computation.

IV. Deep Drone Acrobatics

Acrobatic flight with quadrotors is extremely challenging. Human drone pilots require many years of practice to safely master maneuvers such as power loops and barrel rolls. For aerial vehicles that rely only on onboard sensing and computation, the high accelerations that are required for acrobatic maneuvers together with the unforgiving requirements on the control stack raise fundamental questions in both perception and control. For this reason, we challenged our drone with the task of acrobatics maneuvers [2]. In order to achieve this task, we trained a neural network to predict actions from raw sensor observations. Training is done by imitating an optimal controller with access to privileged information in the form of the exact drone's state. Since this information is not available in the physical world, we trained the neural network to predict actions from inertial and visual observations.

Similarly to previous work, all of the training is done in simulation, without the need of any data from the real world. We achieved this by abstraction of sensor measurements, which are more similar between domains than the raw observation themselves. Both theoretically and experimentally, we have shown that abstractions strongly favours sim-to-real transfer. The learned policy allowed our drone to go faster than ever before and successfully fly maneuvers with accelerations of up to 3g, as the Matty flip in Fig. 3.

V. The Way Forward

Our work shows that neural networks have a strong potential to control agile platforms like quadrotors. In comparison to traditional methods, neural policies are more robust to noise in the observations and can deal with imperfection in sensing and actuation. However, traditional methods are still better in generalizing to different tasks, and are more transparent and easy to interpret. Improving the generalization of neural sensorimotor policies, as well as their interpretability by measuring prediction uncertainty [7], is a very interesting venue for future work. Solving those challenges could potentially bring autonomous drones closer in agility to human pilots and birds.

REFERENCES

[1] M. Müller, A. Dosovitskiy, B. Ghanem, and V. Koltun, “Driving policy transfer via modularity and abstraction,” in *Conference on Robot Learning*, 2018, pp. 1–15.

[2] E. Kaufmann, A. Loquercio, R. Ranftl, M. Müller, V. Koltun, and D. Scaramuzza, “Deep drone acrobatics,” *RSS: Robotics, Science, and Systems*, 2020.

[3] E. Kaufmann, A. Loquercio, R. Ranftl, A. Dosovitskiy, V. Koltun, and D. Scaramuzza, “Deep drone racing: Learning agile flight in dynamic environments,” *Conference on Robot Learning (CORL)*, 2018.

[4] E. Kaufmann, M. Gehrig, P. Foehn, R. Ranftl, A. Dosovitskiy, V. Koltun, and D. Scaramuzza, “Beauty and the beast: Optimal methods meet learning for drone racing,” in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 690–696.

[5] A. Loquercio, E. Kaufmann, R. Ranftl, A. Dosovitskiy, V. Koltun, and D. Scaramuzza, “Deep drone racing: From simulation to reality with domain randomization,” *IEEE Transactions on Robotics*, vol. 36, no. 1, pp. 1–14, 2019.

[6] P. Foehn, D. Brescianini, E. Kaufmann, T. Cieslewski, M. Gehrig, M. Muglikar, and D. Scaramuzza, “Alphapilot: Autonomous drone racing,” *Robotics: Science and Systems (RSS)*, 2020.

[7] A. Loquercio, M. Segu, and D. Scaramuzza, “A general framework for uncertainty estimation in deep learning,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3153–3160, 2020.