Learning-Based Motion Planning for High-Speed Quadrotor Flight

Elia Kaufmann, Davide Scaramuzza

***Abstract*— In this extended abstract, we present our latest research in learning deep sensorimotor policies for agile vision-based quadrotor flight. Instead of relying on an accurate 3D representation of the environment, our research focuses on learning-based approaches that directly map onboard sensory observations to control commands. We show methodologies for successful transfer of such policies from simulation to the real world. 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. One of the main challenges that have blocked progress towards human-level autonomous flight performance is *perception*. This challenge is highlighted by the fact that there are methods that demonstrate impressive feats with quadrotors in controlled environments [1]–[3], while approaches that rely on onboard sensing are constrained to substantially lower agility [4], [5].

Data-driven approaches in the form of neural networks have recently gained a lot of interest due to their potential to make perception robust to noisy real-world data. Recent approaches propose to learn representations of scenes and objects [6], [7] to improve down-stream tasks such as motion planning, manipulation, or pose estimation [8].

In this extended abstract, we summarize our latest research on learning deep sensorimotor policies for agile vision-based quadrotor flight. In contrast to learning a representation of the robot's surrounding that is then combined with a navigation module, our policies represent a holistic approach to autonomous navigation that directly maps onboard sensory observations to commands, without enforcing any intermediate representation of geometry. We train our policies exclusively in simulation and achieve successful transfer to the real platform by leveraging a combination of domain randomization and abstraction of sensory observations. Our policies enable autonomous quadrotors to fly faster and more agile than what was possible before with only onboard sensing and computation, raising the question if accurate knowledge of the 3D surroundings is actually needed for high-speed navigation.

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 (https://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. 1: Our drone flying at high speed through a forest using only onboard sensing and computing.

II. High-Speed Flight in the Wild

We have developed an end-to-end approach that can autonomously fly quadrotors through complex natural and human-made environments at high speeds, with purely onboard sensing and computation. The key principle is to directly map noisy sensory observations to collision-free trajectories in a receding-horizon fashion. This direct mapping drastically reduces processing latency and increases robustness to noisy and incomplete perception. The sensorimotor mapping is performed by a convolutional network that is trained *exclusively* in simulation via privileged learning: imitating an expert with access to privileged information. We leverage abstraction of the input data to transfer the policy from simulation to reality [9], [10]. 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 the gap between simulation and real world. By combining abstraction with simulating realistic sensor noise, our approach achieves zero-shot transfer from simulation to challenging real-world environments that were never experienced during training: dense forests, snow-covered terrain, derailed trains, and collapsed buildings. Our work demonstrates that end-to-end policies trained in simulation enable high-speed autonomous flight through challenging environments, outperforming traditional obstacle avoidance pipelines. A qualitative example of flight in the wild is shown in Figure 1.

Fig. 2: Our quadrotor performs a Matty Flip. The drone is controlled by a deep sensorimotor policy and uses only onboard sensing and computation.

III. Acrobatic Flight

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 performing acrobatic maneuvers [10]. In order to achieve this task, we trained a neural network to predict actions directly from visual and inertial sensor observations. Training is done by imitating an optimal controller with access to privileged information in the form of the exact robot state. Since such information is not available in the physical world, we trained the neural network to predict actions instead 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 using abstraction of sensor measurements, which reduces the simulation-to-reality gap compared to feeding raw observations. Both theoretically and experimentally, we have shown that abstraction strongly favours simulation-to-reality transfer. The learned policy allowed our drone to go faster than ever before and successfully fly maneuvers with accelerations of up to 3g, such as the Matty flip illustrated in Figure 2.

IV. Autonomous Drone Racing

Drone racing is an emerging sport where pilots race against each other with remote-controlled quadrotors while being provided a first-person-view (FPV) video stream from a camera mounted on the drone. Drone pilots undergo years of training to master the skills involved in racing. In recent years, the task of autonomous drone racing has received substantial attention in the robotics community, which can mainly be attributed to two reasons: (i) The sensorimotor skills required for autonomous racing 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. (ii) The task of autonomous racing provides a simple and objective comparison between both robotic and human baselines, which makes it an ideal candidate for a new robotics benchmark.

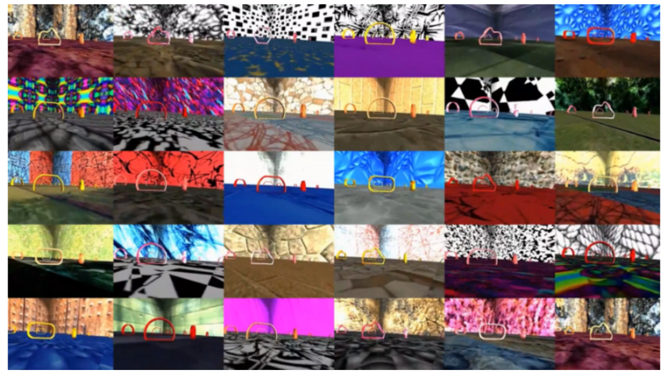

Fig. 3: The perception component of our system, represented by a convolutional neural network (CNN), is trained only with non-photorealistic simulation data.

One approach to autonomous racing is to fly through the course by tracking a precomputed, potentially time-optimal, trajectory [2]. However, such an approach 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 priori what should be the representation of the next waypoint. In our works, we have explored different solutions. In our preliminary work, the neural network predicts a fixed distance location from the drone [11]. 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 a track-dependent global trajectory representing the optimal path through all gates.

For this reason, a follow-up version of this work proposed to use as waypoint the location of the next gate [12]. 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.

Even though the waypoint representation proposed in [12] allowed for efficient transfer between track layouts, it still required substantial amount of real-world data to train.

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 substantially, the data collection process needs to be repeated from scratch. For this reason, in our recent work [13] 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 predictions should be robust against, *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 Figure 3. 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 successfully compete in the alpha-pilot competition [14].

V. Future Work

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, the approaches presented in this extended abstract fall still short of the performance of professional drone pilots. To further push the capabilities of autonomous drones, a specialization to the task is required, potentially through real-world adaptation and online learning. Solving those challenges could potentially bring autonomous drones closer in agility to human pilots and birds.

References

[1] D. Falanga, E. Mueggler, M. Faessler, and D. Scaramuzza, "Aggressive quadrotor flight through narrow gaps with onboard sensing and computing using active vision," in *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2017, pp. 5774–5781.

[2] P. Foehn, A. Romero, and D. Scaramuzza, "Time-optimal planning for quadrotor waypoint flight," *Science Robotics*, vol. 6, no. 56, 2021. [Online]. Available: https://robotics.sciencemag.org/content/6/56/eabh1221

[3] S. Sun, L. Sijbers, X. Wang, and C. de Visser, "High-speed flight of quadrotor despite loss of single rotor," *IEEE Robot. Autom. Lett.*, vol. 3, no. 4, pp. 3201–3207, 2018.

[4] B. Zhou, J. Pan, F. Gao, and S. Shen, "RAPTOR: robust and perception-aware trajectory replanning for quadrotor fast flight," *IEEE Trans. Robot.*, 2021.

[5] H. Oleynikova, C. Lanegger, Z. Taylor, M. Pantic, A. Millane, R. Siegwart, and J. Nieto, "An open-source system for vision-based micro-aerial vehicle mapping, planning, and flight in cluttered environments," *J. Field Robot.*, vol. 37, no. 4, pp. 642–666, Jun. 2020.

[6] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, "Deepsdf: Learning continuous signed distance functions for shape representation," in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2019, pp. 165–174.

[7] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," in *European conference on computer vision*. Springer, 2020, pp. 405–421.

[8] L. Yen-Chen, P. Florence, J. T. Barron, A. Rodriguez, P. Isola, and T.-Y. Lin, "inerf: Inverting neural radiance fields for pose estimation," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 1323–1330.

[9] 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.

[10] E. Kaufmann, A. Loquercio, R. Ranftl, M. Müller, V. Koltun, and D. Scaramuzza, "Deep drone acrobatics," in *Robotics: Science and Systems (RSS)*, 2020.

[11] E. Kaufmann, A. Loquercio, R. Ranftl, A. Dosovitskiy, V. Koltun, and D. Scaramuzza, "Deep drone racing: Learning agile flight in dynamic environments," in *Conf. on Robotics Learning (CoRL)*, 2018.

[12] 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," *IEEE Int. Conf. Robot. Autom. (ICRA)*, pp. 690–696, 2019. [Online]. Available: https://doi.org/10.1109/ICRA.2019.8793631

[13] A. Loquercio, E. Kaufmann, R. Ranftl, A. Dosovitskiy, V. Koltun, and D. Scaramuzza, "Deep drone racing: From simulation to reality with domain randomization," *IEEE Trans. Robot.*, vol. 36, no. 1, pp. 1–14, 2019.

[14] 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. [Online]. Available: https://link.springer.com/article/10.1007/s11370-018-00271-6