Learning Vehicle Models for Agile Flight

Elia Kaufmann, Antonio Loquercio, Yunlong Song, and Davide Scaramuzza

***Abstract*— In this extended abstract, we present our latest research on learning dynamics models for agile quadrotor flight. Access to accurate dynamics models allows to improve performance of model-based controllers, perform controller tuning tasks entirely in simulation, and facilitates simulation-to-reality transfer of learning-based control policies.**

I. Introduction

In recent years, research on fast navigation of autonomous quadrotors has made tremendous progress, continually pushing the vehicles to more aggressive maneuvers (Figure 1). To further advance the field, several competitions have been organized, such as the autonomous drone race series at the recent IROS and NeurIPS conferences [1], [2] and the AlphaPilot challenge [3]. In the near future, estimation and control algorithms will reach the level of maturity necessary to push autonomous quadrotors to the bounds of what is physically possible. This presents the need for quadrotor models that can predict the behaviour of the platform even during highly aggressive maneuvers. Accurately modeling quadrotors flying at their physical limits is extremely challenging and requires to capture complex effects due to aerodynamic forces, motor dynamics, and vibrations. Especially aerodynamic forces pose a challenge, as they depend on hidden state variables like airflow, which cannot be easily measured. Furthermore, the individual downwash induced by the rotors interacts with both the frame and the blades depending on the current state of the platform. The repeatability of tracking errors observed in prior work [4]–[6] and in this work when performing aggressive maneuvers suggests that the difficulty of learning quadrotor dynamics is not caused by stochasticity in the dynamics, but rather by unobserved state variables such as airflow.

Traditional approaches to quadrotor modeling limit the captured effects to simple linear drag approximations and quadratic thrust curves [7]–[9]. Such approximations are computationally efficient and describe the platform well in low-speed regimes, but exhibit increasing bias at higher velocities as they neglect the influence of the inflow velocity on the generated thrust. More elaborate models based on blade-element-momentum (BEM) theory manage to accurately model single rotors at high wind velocities, but they do not account for the aerodynamic interactions between rotors and the frame. Parametric gray-box models [10] aim to overcome these limitations by describing the forces and torques as a linear combination of library functions. While these models can perform well, their performance hinges on the appropriate choice of basis functions, which require human expert knowledge to design. Recent research has investigated computational fluid dynamics [11] to model the aerodynamic effects at play during different flight conditions. While being very accurate, such approaches are computationally expensive and need hours of processing on a compute cluster, rendering them impractical for experiments spanning more than a few seconds.

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).

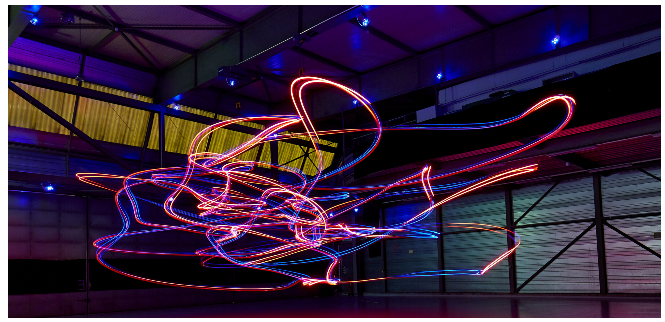

Fig. 1: We developed a quadrotor modeling method that can accurately predict aerodynamic forces and torques acting on the platform. The model is identified from a rich set of maneuvers covering the full performance envelope of the platform.

We have investigated the usage of data-driven methods to model the *residual* dynamics of the quadrotor platform and combining such learned model with a nominal model based on first principles. Our results indicate that such combination results in very accurate dynamics models that are also efficient to compute, while requiring less training data than approaches that learn the entire model from data.

II. Model Learning for Control

Incorporating a learned model into a control pipeline enforces tight constraints on the computational complexity of the model and its inference time. We have developed a lightweight approach that combines Gaussian Processes (GPs) with a nominal quadrotor model that does not account for aerodynamic drag effects [12]. As the nominal model already captures the characteristics of the platform well in the slow-flight regime, the GPs mainly need to model aerodynamic effects that are encountered during fast and aggressive trajectories. This allows to learn an accurate dynamics model from a small number of inducing data points.

As the model is lightweight and efficient to compute, we can integrate it directly into a model predictive controller,

Fig. 2: Our quadrotor platform reaches its physical limits at a pitch angle of 80 degrees while performing a lemniscate trajectory in our experiments. Throughout the trajectory, the platform reaches speeds of up to $14\,\mathrm{m\,s^{-1}}$ and accelerations beyond 4g.

which runs at frequencies greater than 100 Hz. We show that the augmented MPC improves trajectory tracking by up to 70% with respect to its nominal counterpart. We verify our method by extensive comparison to a state-of-the-art linear drag model in synthetic and real-world experiments at speeds of up to $14\,\mathrm{m\,s^{-1}}$ and accelerations beyond 4g.

III. Learning a Simulator

While Gaussian Processes provide an efficient way to augment the dynamics model with real-world flight data, they scale poorly to large amounts of data which imposes a constraint on their predictive capability. Furthermore, the approach presented in [12] does not model residual torques, effectively neglecting moments caused by rotor-to-rotor interactions.

To maximize the predictive power of the quadrotor dynamics model, we developed a quadrotor dynamics model that can accurately capture complex aerodynamic effects by augmenting a state-of-the-art rotor model based on blade-element momentum (BEM) theory with learned residual force and torque terms represented by a deep neural network. The resulting hybrid model (BEM+NN) benefits from the expressive power of deep neural networks and the generalizability of first-principles modeling. The latter reduces the need for extreme amounts of training data. The model is identified using data collected from a large set of maneuvers performed on a real quadrotor platform. Leveraging one of the biggest optical tracking volumes in the world, the platform's state as well as the motor speeds are recorded during flight.

The proposed model is compared against state-of-the-art modeling approaches [13] (PolyFit) as well as BEM without augmentation on unseen test maneuvers. The comparison is done in terms of both evaluation of predicted aerodynamic forces and torques and closed-loop integration of the model in a simulator, each evaluated against real-world reference data. In both categories, a significant performance increase is observed.

This work proposes a novel method to model quadrotors by combining modeling based on first principles with a learning-based residual term represented by a neural network.

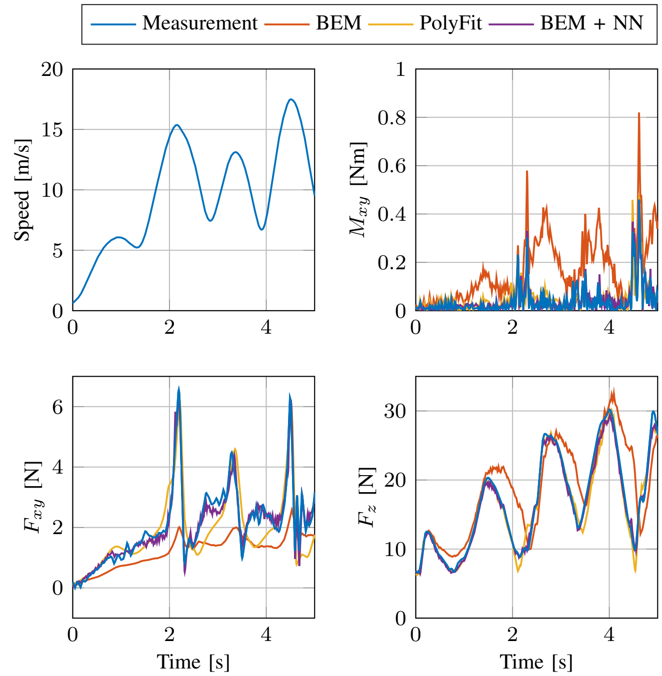

Fig. 3: The plots show a highly aggressive maneuver (from the test dataset) where only models with neural net augmentation predict the forces and torques well.

The proposed method is able to accurately model quadrotors even throughout aggressive trajectories pushing the platform to its limits. This hybrid model outperforms its compositional modules with up to 50% error reduction, including baseline methods that utilize only first-principles modeling, as well as purely learning-based methods. The method shows strong generalization beyond the training set used to identify the model and predicts accurate forces and torques where other methods break down. Controlled experiments indicate that the fusion of learned dynamics with first-principles is a powerful combination. Applied to simulations, our approach enables unprecedented accuracy, reducing potsitional RMSE from $\sim$0.8 m for state-of-the-art approaches, down to below 0.3 m. This could tremendously speed up development and testing of advanced control and navigation strategies for quadrotors, without the need of the tedious and crash-prone trial-and-error strategy on real systems.

IV. Future Work

Our recent works show that combining modeling based on first principles with learning-based residual terms is able to accurately model quadrotors even throughout aggressive trajectories that push the platform to its limits. Our methods show strong generalization beyond the training set used to identify the model parameters and predict accurate forces and torques where other methods break down. Access to such high-fidelity models could tremendously speed up development and testing of advanced control and navigation strategies for quadrotors, such as using policy search methods based on deep reinforcement learning.

REFERENCES

[1] H. Moon, J. Martinez-Carranza, T. Cieslewski, M. Faessler, D. Falanga, A. Simovic, D. Scaramuzza, S. Li, M. Ozo, C. De Wagter *et al.*, "Challenges and implemented technologies used in autonomous drone racing," *Intelligent Service Robotics*, vol. 12, no. 2, pp. 137–148, 2019.

[2] R. Madaan, N. Gyde, S. Vemprala, M. Brown, K. Nagami, T. Taubner, E. Cristofalo, D. Scaramuzza, M. Schwager, and A. Kapoor, "Airsim drone racing lab," in *PLMR Post Proceedings of the NeurIPS 2019 Competition Track*, 2020.

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

[4] A. Punjani and P. Abbeel, "Deep learning helicopter dynamics models," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 3223–3230.

[5] G. Ryou, E. Tal, and S. Karaman, "Multi-fidelity black-box optimization for time-optimal quadrotor maneuvers," in *Robotics: Science and Systems (RSS)*, 2020.

[6] S. Bansal, A. K. Akametalu, F. J. Jiang, F. Laine, and C. J. Tomlin, "Learning quadrotor dynamics using neural network for flight control," in *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 4653–4660.

[7] F. Furrer, M. Burri, M. Achtelik, and R. Siegwart, "Rotors—a modular gazebo mav simulator framework," in *Robot Operating System (ROS)*. Springer, 2016, pp. 595–625.

[8] S. Shah, D. Dey, C. Lovett, and A. Kapoor, "Airsim: High-fidelity visual and physical simulation for autonomous vehicles," in *Field and service robotics*. Springer, 2018, pp. 621–635.

[9] M. Faessler, A. Franchi, and D. Scaramuzza, "Differential flatness of quadrotor dynamics subject to rotor drag for accurate tracking of high-speed trajectories," *IEEE Robot. Autom. Lett.*, 2017.

[10] S. Sun, C. C. de Visser, and Q. Chu, "Quadrotor gray-box model identification from high-speed flight data," *Journal of Aircraft*, vol. 56, no. 2, pp. 645–661, 2019.

[11] P. Ventura Diaz and S. Yoon, "High-fidelity computational aerodynamics of multi-rotor unmanned aerial vehicles," in *2018 AIAA Aerospace Sciences Meeting*, 2018, p. 1266.

[12] G. Torrente, E. Kaufmann, P. Foehn, and D. Scaramuzza, "Data-driven mpc for quadrotors," *IEEE Robotics and Automation Letters*, 2021.

[13] S. Sun, R. Schilder, and C. C. de Visser, "Identification of quadrotor aerodynamic model from high speed flight data," in *2018 AIAA Atmospheric Flight Mechanics Conference*, 2018, p. 0523.