**University of
Zurich**^{UZH}

**Zurich Open Repository and
Archive**

University of Zurich
University Library
Strickhofstrasse 39
CH-8057 Zurich
www.zora.uzh.ch

Year: 2022

Data-Driven MPC for Quadrotors

Torrente, Guillem ; Kaufmann, Elia ; Föhn, Philipp ; Scaramuzza, Davide

DOI: <https://doi.org/10.1109/LRA.2021.3061307>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-216588>

Journal Article

Accepted Version

Originally published at:

Torrente, Guillem; Kaufmann, Elia; Föhn, Philipp; Scaramuzza, Davide (2022). Data-Driven MPC for Quadrotors. IEEE Robotics and Automation Letters, 6(2):3769-3776.

DOI: <https://doi.org/10.1109/LRA.2021.3061307>

Data-Driven MPC for Quadrotors

Guillem Torrente*, Elia Kaufmann*, Philipp Föhn, Davide Scaramuzza

Abstract—Aerodynamic forces render accurate high-speed trajectory tracking with quadrotors extremely challenging. These complex aerodynamic effects become a significant disturbance at high speeds, introducing large positional tracking errors, and are extremely difficult to model. To fly at high speeds, feedback control must be able to account for these aerodynamic effects in real-time. This necessitates a modeling procedure that is both accurate and efficient to evaluate. Therefore, we present an approach to model aerodynamic effects using Gaussian Processes, which we incorporate into a Model Predictive Controller to achieve efficient and precise real-time feedback control, leading to up to 70% reduction in trajectory tracking error at high speeds. 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 14m/s and accelerations beyond 4g.

SUPPLEMENTARY MATERIAL

Video: <https://youtu.be/FHvDghUUQtC>

Code: https://github.com/uzh-rpg/data_driven_mpc

I. INTRODUCTION

Accurate trajectory tracking with quadrotors in high-speed and high-acceleration regimes is still a challenging research problem. While autonomous quadrotors have seen a significant gain in popularity and have been applied in a variety of industries ranging from agriculture to transport, security, infrastructure, entertainment, and search and rescue, they still do not exploit their full maneuverability. The ability to precisely control drones during fast and highly agile maneuvers would allow to not only fly fast in known-free environments, but also close to obstacles, humans, or through openings, where already small deviations from the reference have catastrophic consequences.

Operating a quadrotor at high speeds and controlling it through agile, high-acceleration maneuvers requires to account for complex aerodynamic effects acting on the platform. These effects are difficult to model, since they consist of a combination of propeller lift and drag dependent on the induced airstream velocity, fuselage drag, and complex or even turbulent effects due to the interaction between the propellers, the downwash of other propellers, and the fuselage.

Manuscript received: October 15th, 2020; Revised: December 26th, 2020; Accepted: February 2nd, 2021

This paper was recommended for publication by Editor Clement Gosselin upon evaluation of the Associate Editor and Reviewers' comments.

*These two authors contributed equally.

The authors are with the Robotics and Perception Group, Dep. of Informatics, University of Zurich, and Dep. 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).

Digital Object Identifier (DOI): see top of this page.

Fig. 1: 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 14m s^{-1} and accelerations beyond 4g.

Furthermore, in the context of model-based feedback control, the model complexity is constrained by the feedback time-scale and computational capabilities of the executing platform. Therefore, it is not sufficient to find the most accurate model, but required to find an applicable trade-off between model accuracy and complexity.

Very little work exists on agile control of quadrotors at speeds beyond 5m s^{-1} and accelerations above 2g, [1–8]. Even though these works show agile control at various levels, none of them accounts for aerodynamic effects. This is not a limiting assumption when the quadrotor is controlled close to hover conditions, but introduces significant errors when tracking fast and agile trajectories. Other approaches use iterative learning control to perform highly aggressive trajectories [9], but they are constrained to a single maneuver and do not generalize.

The main challenge when performing aggressive flight is to identify a dynamics model of the platform that is capable of describing the aerodynamic effects while still being lightweight enough to guarantee real time performance. While there exist sophisticated computational fluid dynamics simulations that are able to model turbulent aerodynamic effects [10], they require hours of processing on a compute cluster, and still need to be abstracted in simplified models to be tractable in a control loop running at high frequency.

In this work, we propose to learn the aerodynamic effects acting on the platform from data. Inspired by [11, 12], we use Gaussian Processes to learn the residual dynamics with respect to a simplified quadrotor model that does not account for aerodynamic effects. Learning the residual dynamics simplifies the learning problem and allows describing the model augmentation using only a small number of inducing points

for Gaussian Processes. Using such a small model allows leveraging the combined dynamics formulation in a Model Predictive Control (MPC) pipeline.

Our experiments, performed in simulation and in the real world, show that the proposed approach can significantly improve control performance for agile trajectories with speeds up to 14 m s^{-1} and accelerations exceeding $4g$. We show that the method generalizes between different trajectories and outperforms methods relying on simplified correction terms.

Contributions

In this paper, we present a Model Predictive Control pipeline that is augmented with learned residual dynamics using Gaussian Processes. We extend the approach presented in [11, 12] to three-dimensional GP predictions for the quadrotor platform. By combining the learned GP corrections with the nominal quadrotor dynamics, we can learn an accurate dynamics model from a small number of inducing data points. Such a small model can be efficiently optimized within an MPC pipeline and allows for control 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 m s^{-1} and accelerations beyond $4g$.

II. RELATED WORK

Performing fast and agile maneuvers with an autonomous robot requires knowledge of an accurate dynamics model of the platform. However, especially at high speeds and accelerations, such a description of the system is difficult to obtain due to hard-to-model effects caused by friction, aerodynamics or varying battery voltage. As modeling these effects significantly increases the problem complexity, controlling a robot under such conditions requires to find a trade-off between model expressivity and computational tractability. Most prior work on control of autonomous robots does not account for higher-order effects at all and treats them as external disturbances [1, 13–16]. While this allows for very efficient and lightweight controller implementations, tracking performance progressively decreases for higher speeds. In [17], Nonlinear Incremental Dynamics Inversion is used to achieve robust tracking of fast trajectories. However, the reactive nature of this approach does not allow to account for future disturbances as the controller does not optimize over a horizon of actions.

A recent line of work [18–21] investigates the application of dynamics learned entirely from data for a variety of applications such as robot arms, cars or fluids. These learned dynamics representations take the form of deep neural networks and substitute the nominal dynamics in the MPC. While the resulting dynamics models are very expressive, their optimization is often intractable due to local minima. A common way to overcome this challenge is to use sampling-based optimizers, which in turn scale poorly to high-dimensional input spaces.

Instead of learning the full dynamics from data, [11, 12, 22, 23] combine a nominal model with a learned correction term.

This allows to limit the learned dynamics to have different dimensionality than the nominal system and provides the possibility to learn only specific effects that are difficult to capture with the nominal model.

For the particular case of quadrotor flight, the most prominent source of disturbances are aerodynamic effects originating from drag caused by the rotors and the fuselage, as well as lift effects that act on the platform at high speeds. By conducting controlled experiments in both wind tunnels as well as instrumented tracking volumes, previous work has shown a significant effect of aerodynamic forces already at linear speeds of 5 m s^{-1} [24, 25].

Neglecting other aerodynamic effects, previous work mainly studies the effect of rotor drag [26, 24, 27]. Rotor drag effects originate from blade flapping and induced drag of the rotors. These effects are typically combined as linear effects in a lumped parameter dynamical model [28]. In [24], the authors identify a linear model for the rotor drag and use it to compute feedforward terms of a PID controller. Even though they show improved trajectory tracking performance, they evaluate their linear model only up to 5 m s^{-1} . At these speeds, the linear effect of rotor drag dominates the fuselage drag. We integrate the model of [24] in an MPC pipeline to act as baseline of our approach.

Similar to our work, in [11, 12, 29–31], the authors use Gaussian Processes to improve the control performance of a robotic platform. In [29], Gaussian Processes are used on a quadrotor to correct for wind disturbances. Since instead of platform states only observed disturbances are fed to the GPs, this approach does not learn a dynamics model and can only react to disturbances once they have been observed. In [30], the authors separately learn the translational and rotational dynamics of a quadrotor platform. As this approach learns the full model from data, it requires a large number of training points and is computationally very expensive. As a result, the prediction horizon of the MPC needs to be reduced to a single point to achieve near real-time performance. In [31], a robust experience-driven predictive controller (EPC) is proposed that uses Gaussian belief propagation to account for uncertainties in the state estimate. The controller demonstrates robust constraint satisfaction on a quadrotor platform, where it is integrated into an MPC that controls the translational dynamics of the vehicle. In [11, 12], the authors use the predictions of the Gaussian Processes to improve the tracking performance of an autonomous race car by learning the *residual* dynamics of a nominal model. Learning such residual dynamics instead of the full model allows them to simplify the learning problem and as a result reduce the number of training points in the GP.

Our work is inspired by these approaches, but extends [11, 12] to three-dimensional GP predictions for the quadrotor platform. Instead of learning a mapping from observed disturbances to future disturbance as in [29], we focus on fast flight and correct for aerodynamic effects that arise due to the fast ego-motion of the platform. We tightly integrate the predictions of the Gaussian Processes into the MPC formulation. Instead of using virtual inputs such as bodyrates and collective thrust [1], our MPC models the dynamics down to the motor inputs and can therefore account for the true

Fig. 2: Diagram of the quadrotor model with the world and body frames and propeller numbering convention.

actuation limits of the platform. Our work is the first to combine Gaussian Processes with a full-state quadrotor MPC formulation to model aerodynamic drag effects while still being able to account for the true actuation limits of the platform.

III. METHODOLOGY

A. Notation

We denote scalars in lowercase s , vectors in lowercase bold $\mathbf{v}$, and matrices in uppercase bold $\mathbf{M}$. We define the World W and Body B frames with orthonormal basis i.e. $\{\mathbf{x}_W, \mathbf{y}_W, \mathbf{z}_W\}$. The frame B is located at the center of mass of the quadrotor. Note that we assume all four rotors are situated in the xy -plane of frame B , as depicted in Fig. 2. A vector from coordinate $\mathbf{p}_1$ to $\mathbf{p}_2$ expressed in the W frame is written as: ${}_W\mathbf{v}_{12}$. If the vector's origin coincide with the frame it is described in, we drop the frame index, e.g. the quadrotor position is denoted as $\mathbf{p}_{WB}$. Furthermore, we use unit quaternions $\mathbf{q} = (q_w, q_x, q_y, q_z)$ with $\|\mathbf{q}\| = 1$ to represent orientations, such as the attitude state of the quadrotor body $\mathbf{q}_{WB}$. Finally, full SE3 transformations, such as changing the frame of reference from body to world for a point $\mathbf{p}_{B1}$, can be described by ${}_W\mathbf{p}_{B1} = {}_W\mathbf{t}_{WB} + \mathbf{q}_{WB} \odot \mathbf{p}_{B1}$. Note the quaternion-vector product denoted by $\odot$ representing a rotation of the vector by the quaternion as in $\mathbf{q} \odot \mathbf{v} = \mathbf{q}\mathbf{v}\bar{\mathbf{q}}$, where $\bar{\mathbf{q}}$ is the quaternion's conjugate.

B. Nominal Quadrotor Dynamics Model

We assume the quadrotor is a 6 degree-of-freedom rigid body of mass m and diagonal moment of inertia matrix $\mathbf{J} = \text{diag}(J_x, J_y, J_z)$. Our model is similar to [1, 3] but we write the nominal dynamics $\dot{\mathbf{x}}$ up to second order derivatives, leaving the quadrotors individual rotor thrusts $T_i \forall i \in (0, 3)$ as control inputs $\mathbf{u}$. The state space is thus 13-dimensional and its dynamics can be written as:

$$\dot{\mathbf{x}} = \begin{bmatrix} \dot{\mathbf{p}}_{WB} \\ \dot{\mathbf{q}}_{WB} \\ \dot{\mathbf{v}}_{WB} \\ \dot{\boldsymbol{\omega}}_B \end{bmatrix} = \mathbf{f}_{dyn}(\mathbf{x}, \mathbf{u}) = \begin{bmatrix} \mathbf{v}_W \\ \mathbf{q}_{WB} \cdot \begin{bmatrix} 0 \\ \boldsymbol{\omega}_B/2 \end{bmatrix} \\ \frac{1}{m} \mathbf{q}_{WB} \odot \mathbf{T}_B + \mathbf{g}_W \\ \mathbf{J}^{-1} (\boldsymbol{\tau}_B - \boldsymbol{\omega}_B \times \mathbf{J} \boldsymbol{\omega}_B) \end{bmatrix}, \quad (1)$$

where $\mathbf{g}_W = [0, 0, -9.81 \text{ m/s}^2]^\top$ denotes Earth's gravity, $\mathbf{T}_B$ is the collective thrust and $\boldsymbol{\tau}_B$ is the body torque as in:

$$\mathbf{T}_B = \begin{bmatrix} 0 \\ 0 \\ \sum T_i \end{bmatrix} \quad \text{and} \quad \boldsymbol{\tau}_B = \begin{bmatrix} d_y(-T_0 - T_1 + T_2 + T_3) \\ d_x(-T_0 + T_1 + T_2 - T_3) \\ c_\tau(-T_0 + T_1 - T_2 + T_3) \end{bmatrix} \quad (2)$$

where d_x, d_y are the rotor displacements and c_τ is the rotor drag torque constant. To incorporate these dynamics in discrete time algorithms, we use an explicit Runge-Kutta method of 4th order $\mathbf{f}_{RK4}(\mathbf{x}, \mathbf{u})$ to integrate $\dot{\mathbf{x}}$ given an initial state $\mathbf{x}_k$, input $\mathbf{u}_k$ and integration step δt by:

$$\mathbf{x}_{k+1} = \mathbf{f}_{RK4}(\mathbf{x}_k, \mathbf{u}_k, \delta t). \quad (3)$$

C. Gaussian Process-Augmented Dynamics

Inspired by [11, 12], we use Gaussian Processes to complement the nominal dynamics of the quadrotor in an MPC pipeline. In this setting, the GPs predict the error of the dynamics and correct them at every time instance t_k . Similar to most GP-based learning problems, we assume the existence of the inaccessible true dynamics $\mathbf{f}_{true}$ of the quadrotor, which we measure as $\tilde{\mathbf{y}}_{k+1}$ through a noisy process at discrete time instances t_k :

$$\tilde{\mathbf{y}}_{k+1} = \mathbf{f}_{true}(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{w}_k \quad (4)$$

We further assume that $\mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$ is Gaussian noise, where $\boldsymbol{\Sigma}$ is the time-invariant and diagonal covariance matrix. This means we can effectively treat each dimension of $\mathbf{y}_k$ independently through a separate 1-dimensional output GP. We use the Radial Basis Function (RBF) kernel defined by:

$$\kappa(\mathbf{z}_i, \mathbf{z}_j) = \sigma_f^2 \exp\left(-\frac{1}{2}(\mathbf{z}_i - \mathbf{z}_j)^\top \mathbf{L}^{-2}(\mathbf{z}_i - \mathbf{z}_j)\right) + \sigma_n^2 \quad (5)$$

where $\mathbf{L}$ is the diagonal length scale matrix and σ_f, σ_n represent the data and prior noise variance, respectively, and $\mathbf{z}_i, \mathbf{z}_j$ represent data features.

We redefine the system dynamics as a (corrected, $\mathbf{f}_{corr}$) combination of the dynamics (1) plus the mean posterior of a GP, $\boldsymbol{\mu}$. The GP only corrects a subset of the state, determined in the selection matrix $\mathbf{B}_d$, using the feature vector $\mathbf{z}_k$, determined by selection matrix $\mathbf{B}_z$:

$$\mathbf{f}_{corr}(\mathbf{x}_k, \mathbf{u}_k) = \mathbf{f}_{dyn}(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{B}_d \boldsymbol{\mu}(\mathbf{z}_k) \quad (6)$$

$$\mathbf{z}_k = \mathbf{B}_z [\mathbf{x}_k^\top \quad \mathbf{u}_k^\top]^\top. \quad (7)$$

Given the concatenated training feature samples $\mathbf{Z}$ and the query feature samples $\mathbf{Z}_k$, the mean and covariance of the GP prediction can be recovered as follows:

$$\boldsymbol{\mu}(\mathbf{Z}_k) = \mathbf{K}_k^\top \mathbf{K}^{-1} \mathbf{Z} \quad \boldsymbol{\Sigma}_{\mu k} = \mathbf{K}_{kk} - \mathbf{K}_k^\top \mathbf{K}^{-1} \mathbf{K}_k$$

with $\mathbf{K} = \kappa(\mathbf{Z}, \mathbf{Z}) + \sigma_n^2 \mathbf{I}$ (8)

$$\mathbf{K}_k = \kappa(\mathbf{Z}, \mathbf{Z}_k) \quad \mathbf{K}_{kk} = \kappa(\mathbf{Z}_k, \mathbf{Z}_k).$$

where $\mathbf{K}_{ij}$, the entry of $\mathbf{K}$ with index i, j , is $\mathbf{K}_{ij} = \kappa(\mathbf{z}_i, \mathbf{z}_j)$.

Given the mean and covariance of the GP, not only is it possible to learn the corrected dynamics, but in addition we can also propagate the corrected model forward in time to use it in an MPC. To propagate the state we simply substitute

the nominal dynamics $\mathbf{f}_{dyn}$ with the corrected model $\mathbf{f}_{cor}$ in the Runge-Kutta integration. For the propagation of the covariance, we refer to the formulation in [11, 12].

D. MPC Formulation

In its most general form, MPC stabilizes a system subject to its dynamics $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u})$ along a reference $\mathbf{x}^*(t), \mathbf{u}^*(t)$, by minimizing a cost $\mathcal{L}(\mathbf{x}, \mathbf{u})$ as in:

$$\begin{aligned} & \min_{\mathbf{u}} \int \mathcal{L}(\mathbf{x}, \mathbf{u}) \\ & \text{subject to} \quad \dot{\mathbf{x}} = \mathbf{f}_{dyn}(\mathbf{x}, \mathbf{u}) \quad \mathbf{x}(t_0) = \mathbf{x}_{init} \\ & \quad \quad \quad \mathbf{r}(\mathbf{x}, \mathbf{u}) = 0 \quad \quad \quad \mathbf{h}(\mathbf{x}, \mathbf{u}) \leq 0 \end{aligned} \quad (9)$$

where $\mathbf{x}_0$ denotes the initial condition and $\mathbf{h}, \mathbf{r}$ can incorporate (in-)equality constraints, such as input limitations.

For our application, and as most commonly done, we specify the cost to be of quadratic form $\mathcal{L}(\mathbf{x}, \mathbf{u}) = \|\mathbf{x} - \mathbf{x}^*\|_Q^2 + \|\mathbf{u} - \mathbf{u}^*\|_R^2$ and discretize the system into N steps over time horizon T of size $dt = T/N$. We account for input limitations by constraining $0 \leq \mathbf{u} \leq u_{max}$, and optionally include the GP predictions within the system dynamics.

$$\begin{aligned} & \min_{\mathbf{u}} \mathbf{x}_N^T \mathbf{Q} \mathbf{x}_N + \sum_{k=0}^N \mathbf{x}_k^T \mathbf{Q} \mathbf{x}_k + \mathbf{u}_k^T \mathbf{R} \mathbf{u}_k \\ & \text{subject to} \quad \mathbf{x}_{k+1} = \mathbf{f}_{RK4}(\mathbf{x}_k, \mathbf{u}_k, \delta t) \\ & \quad \quad \quad \mathbf{x}_0 = \mathbf{x}_{init} \\ & \quad \quad \quad u_{min} \leq \mathbf{u}_k \leq u_{max} \end{aligned} \quad (10)$$

where $\mathbf{f}_{RK4}$ can be extended to the corrected dynamics $\mathbf{f}_{cor}$.

To solve this quadratic optimization problem we construct it using a multiple shooting scheme [32] and solve it through a sequential quadratic program (SQP) executed in a real-time iteration scheme (RTI) [32]. All implementations are done using ACADOS [33] and CasADi [34].

E. Practical Implementation

The implementation of the learned dynamics of our GP-MPC must be designed to maximize the performance while minimizing the computational cost added to the optimization. Note that aerodynamic effects operate on the body reference frame. Likewise, the training dataset is adjusted such that the learning problem setup is to identify such a mapping from body frame velocities ${}_B\mathbf{v}$ to body frame acceleration disturbances ${}_B\mathbf{a}_e$, so that ${}_B\mathbf{a}_e = \boldsymbol{\mu}({}_B\mathbf{v})$. Furthermore, to reduce the need for additional training samples, we reduce the dimensionality of our input space such that the mappings are learned axis wise. We can thus write the GP prediction as:

$${}_B\mathbf{a}_{ek} = \boldsymbol{\mu}({}_B\mathbf{v}_k) = \begin{bmatrix} \mu_{vx}({}_Bv_{xk}) \\ \mu_{vy}({}_Bv_{yk}) \\ \mu_{vz}({}_Bv_{zk}) \end{bmatrix} \quad (11)$$

$$\boldsymbol{\Sigma}_{\boldsymbol{\mu}({}_B\mathbf{v}_k)} = \text{diag} \left(\begin{bmatrix} \sigma_{vx}^2({}_Bv_{xk}) \\ \sigma_{vy}^2({}_Bv_{yk}) \\ \sigma_{vz}^2({}_Bv_{zk}) \end{bmatrix} \right) \quad (12)$$

F. Data Collection and Model Learning

To fit the GPs, real-world flight data is collected (as detailed in Sec. IV) using the nominal dynamics model. For each sample at time t_k , the velocity at the next sample point ${}_B\mathbf{v}_{k+1}$ and the predicted velocity at the next sample point ${}_B\hat{\mathbf{v}}_{k+1}$ are recorded, together with the timestep δt_k . We can then compute the time-normalized velocity error, corresponding to the acceleration error:

$${}_B\mathbf{a}_{ek} = \frac{{}_B\mathbf{v}_{k+1} - {}_B\hat{\mathbf{v}}_{k+1}}{\delta t_k} \quad (13)$$

To select the hyperparameters of the kernel function (l, σ_n, σ_f) , we perform maximum likelihood optimization on the collected dataset. Being a non-parametric method, the complexity of GP regression depends on the number of training points. As using the full dataset would make MPC optimization intractable in real time, we subsample the dataset and only use a small number of inducing points. To this end, we leverage the smooth nature of the aerodynamic effects by sampling these points at regular intervals in the ranges of the training set.

IV. EXPERIMENTS AND RESULTS

We design our evaluation procedure to address the following questions: (i) What is the contribution of the learned dynamics of our GP-MPC in a closed-loop tracking task? (ii) How does our GP-MPC compare to an MPC with linear aerodynamic effect compensation, as proposed in [24]? (iii) How does the learned model generalize to unseen trajectories? Finally, we validate our design choices with ablation studies. We refer the reader to the attached video to understand the dynamic nature of our experiments.

A. Experimental Setup

We conduct experiments both in simulation as well as on a real quadrotor platform. To assess our proposed approach, the quadrotor executes three different trajectories (Random, Circle, Lemniscate), illustrated in Fig. 3. The lemniscate trajectory lies in the horizontal plane and is defined by $[x(t) = 2 \cos(\sqrt{2}t); y(t) = 2 \sin(\sqrt{2}t) \cos(\sqrt{2}t)]$.

We compare the tracking performance on these trajectories using our MPC with the *Nominal* quadrotor model (1), and the improvement after adding different correction terms identified from data. We study two possible augmentations: *GP-MPC* (ours) and *RDRv*. The RDRv approach was proposed in [24] as a feed-forward PID controller term, which identifies a set of linear drag coefficients along the body axes. We instead incorporate this linear compensation into the nominal dynamics of our MPC pipeline. Note that the three control approaches only differ in the dynamics model used by the MPC, i.e. they use the same control frequency and cost matrices. Furthermore, both the GP-MPC and the RDRv are always trained on the same dataset.

For the simulation experiments, we perform a Nominal run on random polynomial trajectories of high aggressiveness to collect training samples for the GP and the coefficient

Fig. 3: Trajectories considered in this work. Left: single-loop randomly-generated polynomial trajectory with motion along all axes. The tracking starts and ends at the upper-right corner. Center and right: circular and lemniscate trajectories respectively. Both have zero translation along the z axis, tracking starts at 0 velocity, ramps up until reaching a peak, and ramps down back to hover. The position references remain as shown in the figures in all cases.

Fig. 4: Closed-loop position tracking error as a function of maximum velocity achieved in our custom simulator. *Ideal* denotes the nominal MPC performance in a disturbance-free scenario. *Nominal* corresponds to the un-augmented MPC, and *GP-MPC 15* and *GP-MPC 100* are our GP-augmented controllers where the GP's have been trained with 15 and 100 training samples.

identification for RDRv. With both models fitted, we deploy all three controllers, Nominal, RDRv and GP-MPC, on the test trajectories without retraining. For the real-world experiments, we first perform a run of the Nominal baseline on both circle and lemniscate trajectories, which is also used for GP training and RDRv coefficient identification. In the subsequent rollouts, we test RDRv and GP-MPC on these two trajectories in different permutations.

In our experimental setup, we only consider obstacle-free scenarios and aim to minimize the tracking error that arises mainly due to aerodynamic disturbance effects. For this reason, the predicted covariance of the GPs is not used in our experiments.

B. Experiments in Simulation

We first evaluate the performance for individual maneuvers in simulation. To isolate the effects of varying MPC computation times for different models, we divide the simulation experiments into two parts: *simplified simulation* and *Gazebo simulation*. This setup allows to compare the predictive per-

formance of arbitrary sized models without the need to correct for varying computation times.

Simplified Simulation This simulation is constituted of a simple forward integration of the system dynamics (1) using an explicit Runge-Kutta method of 4th order with a step size of 0.5 ms. We assume to have access to perfect odometry measurements of the quadrotor, ideal tracking of the commanded single-rotor thrusts, and that the MPC computation is instantaneous. The simulator models drag effects caused both by the rotors as well as the fuselage. Additionally, zero-mean Gaussian noise forces and torques are simulated that act on the quadrotor body, as well as asymmetric noise on the motor voltage signals.

In the simplified simulation, we investigate the influence of the number of training points of the GP on the predictive performance. As the choice of this hyperparameter constitutes a trade-off between model accuracy and computation time, we seek the model with the minimum number of inducing points that surpasses a desired performance threshold. This effect is investigated with two experiments: first, we analyze the trade-off between GP performance and optimization time.

Fig. 5: Trade-off between number of training samples, GP performance (top) and optimization time (bottom). Models are trained and evaluated on data collected in our Simplified Simulation.

In the second experiment, we extend the comparison of different-sized GPs to closed-loop experiments on the three test trajectories.

Having identified the optimal size of the GP, we perform an additional set of experiments to compare the tracking performance on the circle and lemniscate test trajectories between the Nominal and RDRv baselines and our approach. Both trajectories are executed up to varying maximum speeds, where the highest speed pushes the platform to its physical limits. The training set for both the RDRv and the GP models in this simulator is collected by executing random polynomial trajectories of high aggressiveness (such as Fig. 3 left) with the quadrotor, up to 16 m s^{-1} axis-wise. This technique works well in simulation since it allows to explore densely all the ranges of operating points without risk of breaking the aircraft if tracking fails.

The results of the GP size analysis are summarized in Fig. 5. As can be seen, the complexity of the optimization problem approximately follows a linear function with respect to the number of training points. Since the predictive performance barely increases when adding more than 20 inducing points, we identify the optimal range to be between 15 to 25 samples, corresponding to 4 – 5 ms of optimization time. For comparative purposes, we illustrate in Fig. 4 how two of our GP models perform in closed-loop tracking for 15 and 100 training samples. It can be verified that a larger number of samples is strictly beneficial, but comes at the expense of an increase in optimization time. In fact, such a large model is not usable for a real time application of our pipeline. Based on this evidence, we chose to use 20 inducing points for the rest of this work.

The main results of the closed-loop tracking experiments are summarized in Table I. The table reports position tracking error in millimeters for both maneuvers at varying maximum speeds. While the *Ideal* column indicates the tracking error in case of no unmodelled disturbances (i.e. the MPC dynamics model perfectly fit the actual system), the *Nominal* column

represents the baseline when no model augmentations are enabled. Note that even though the MPC controller performs very well in the Ideal scenario, it does not achieve zero tracking error due to discretization effects. Both the RDRv baseline as well as the GP-MPC significantly improve the tracking error compared to the non-augmented Nominal case. However, while RDRv performs comparably to our approach up to speeds of 4 m s^{-1} , it starts to fail for higher speeds due to its inability to model higher-order aerodynamic effects such as body drag. Our approach also captures these effects very well and shows consistent improvement for the full range of tested speeds.

Gazebo Simulation To verify the results obtained in the simplified simulation in a well-known quadrotor simulator, we also perform closed-loop tracking experiments in Gazebo [35]. We employ the AscTec Hummingbird quadrotor model using the RotorS extension [36]. To properly evaluate the performance of our pipeline, we also use ground truth odometry measurements instead of a state estimator. We collect a dataset containing velocities in the range $[-12, 12] \text{ m s}^{-1}$ for training our models. This dataset is obtained by tracking randomly generated aggressive trajectories, as with the simplified simulator. We execute the circle and lemniscate trajectories at increasing speeds and compare the tracking performance of the Nominal and RDRv baselines, as well as our approach. Note that once again we use completely independent training and test sets in our setup to ensure our models can generalize to new trajectories.

The main results of the Gazebo experiments are summarized in Fig. 6. In this case, both RDRv as well as our GP-MPC achieve very similar performance over the full range of tested speeds. This is expected, as the RotorS package implementation only simulates rotor drag as aerodynamic effect [36], which follows a linear mapping with respect to the body frame velocity [27]. The true aerodynamic effects acting on a quadrotor however are a combination of rotor drag, body drag and turbulent effects caused by the propellers. We are analyzing these effects in more detail in the following section.

C. Experiments in the Real World

Lastly, we compare the performance of our GP-MPC against both Nominal as well as RDRv controllers on a real quadrotor. We use a custom quadrotor that weighs 0.8kg and has a thrust-to-weight ratio of 5:1. We run the controller on a laptop

TABLE I: Comparison of closed-loop tracking errors on the circle and lemniscate trajectories in simulation.

		Model					
		Ideal	Nominal	RDRv		GP-MPC	
	v_{peak} [m s ⁻¹]	RMSE [mm]	RMSE [mm]	RMSE [mm]	%↓	RMSE [mm]	%↓
Circle	4	0.1	114.1	18.1	84	16.1	85
	8	0.4	241.2	56.9	76	25.4	89
	12	1.2	338.3	93.0	72	28.4	93
Lemn.	4	0.3	104.0	15.5	85	16.3	84
	8	1.5	157.7	32.3	79	20.3	87
	12	4.2	212.4	60.6	71	24.4	88
Opt. dt [ms]		1.32		1.76		4.13

Fig. 6: Closed-loop position tracking error as a function of maximum velocity achieved in the RotorS Gazebo simulator in the circle and the lemniscate trajectories.

Fig. 7: The quadrotor used for real-world experiments.

computer and send control commands in the form of collective thrust and desired bodyrates at 50 Hz to the quadrotor through a Laird RM024 radio module. A PID controller running on-board the quadrotor tracks the sent commands. The quadrotor flies in an indoor arena equipped with an optical tracking system that provides pose estimates at 100 Hz. Note that our control method also works with state estimates that are obtained differently than with a motion capture system. As in the simulation experiments, we compare the tracking error along both circle and lemniscate trajectories with speeds up to 14 m s^{-1} .

To demonstrate that our approach can correct for complex aerodynamic effects, we perform the real world experiments in two settings: in setting (i) we perform all maneuvers with the quadrotor as pictured in Fig. 7, while in setting (ii) we extend the quadrotor body with a vertical drag board. This drag board introduces additional asymmetric aerodynamic disturbance as can be seen in Fig. 8. For the real world experiment, we use 20 inducing points on our GP's.

Table II summarizes the results of our real world experiments in setting (i). We train two GP models on the circle and lemniscate trajectories, and use them at test time in all permutations. As can be seen, our methods as well as the RDRv baseline improve tracking performance by up to 50%, with our approach slightly outperforming the RDRv

Fig. 8: Aerodynamic effects observed in the real quadrotor platform along body axes x (left column) and y (right column) as a function of body frame velocity. The platform was studied in its default configuration (upper row), and with an additional flat board attached along the body x axis (lower row), resulting in a significantly increased body drag effect in the y direction.

baseline. This result can be explained by the fact that the quadrotor platform used in setting (i) is very compact and powerful, rendering the main source of disturbance being the rotor drag. Rotor drag is a linear effect, which can be well compensated for by the linear RDRv model augmentation.

TABLE II: Comparison of the RDRv and GP-MPC methods in the real world experiments.

	Model RMSE [mm]						
Ref.	Nomin.	GP (circle)	%↓	GP (lemn.)	%↓	RDRv	%↓
Circle	319.7	172.9	46	141.0	56	168.3	47
Lemn.	396.2	254.2	36	266.3	33	269.3	33

TABLE III: Velocity-dependent tracking performance of the augmented MPC methods on the circle trajectory.

Config.	v_{range} [m s ⁻¹]	Model				
		Nominal	RDRv		GP-MPC 20	
		RMSE [m]	RMSE [m]	%↓	RMSE [m]	%↓
Default	0-2	0.087	0.130	-49	0.109	-25
	2-4	0.233	0.119	49	0.103	56
	4-6	0.329	0.177	46	0.129	61
	6-8	0.458	0.210	54	0.154	66
	8-10	0.531	0.192	64	0.203	62
Drag board	0-2	0.197	0.132	33	0.060	69
	2-4	0.346	0.287	17	0.078	77
	4-6	0.564	0.381	32	0.141	75
	6-8	0.837	0.463	45	0.219	74
	8-10	0.912	Crash	?	0.379	59

The slight improvement of GP-MPC over RDRv in this setting can be explained by the ability of the GPs to also account for imperfect thrust mappings.

Finally, we compare the circle tracking for settings (i) and (ii) in Table III. As can be seen, our approach significantly outperforms RDRv in setting (ii), where the latter fails to capture the full nonlinearity of aerodynamic effects. This can be verified also in Fig. 8, where the linear fit leads to significant bias.

V. CONCLUSION

In this work, we propose the usage of Gaussian Processes to augment the nominal dynamics of a quadrotor to compensate for aerodynamic effects. This GP-based model augmentation is integrated in a Model Predictive Controller and the resulting system significantly improves positional tracking error, both in simulation and on a real quadrotor. Using data from previously recorded flights, the GP's are trained to predict the acceleration error of the nominal model given its current velocity in body frame.

In extensive experiments in simulation and the real world, we show that our approach outperforms a state-of-the-art linear drag model. Furthermore, our GP-augmented controller opens up interesting lines of follow-up research for future development. On one hand, we plan to make use of the predicted uncertainty to perform safe agile trajectories close to obstacles. On the other hand, leveraging the fast fitting time of our GP models (in the order of seconds), training and control loop can be executed in parallel on separate threads in real time during flight. This would enable to adapt the dynamics model to varying external or internal conditions such as wind disturbance or battery voltage.

VI. ACKNOWLEDGMENT

The authors thank Leonard Bauersfeld for his help with the photographic illustrations.

REFERENCES

- [1] D. Falanga, P. Foehn, P. Lu, and D. Scaramuzza, "PAMPC: Perception-aware model predictive control for quadrotors," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2018.
- [2] M. Faessler, F. Fontana, C. Forster, and D. Scaramuzza, "Automatic re-initialization and failure recovery for aggressive flight with a monocular vision-based quadrotor," in *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2015.
- [3] M. Kamel, T. Stastny, K. Alexis, and R. Siegwart, "Model predictive control for trajectory tracking of unmanned aerial vehicles using robot operating system," in *Robot Operating System (ROS)*. Springer, 2017, pp. 3–39.
- [4] M. Neunert, C. de Crousaz, F. Furrer, M. Kamel, F. Farshidian, R. Siegwart, and J. Buchli, "Fast nonlinear model predictive control for unified trajectory optimization and tracking," in *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2016.
- [5] D. Mellinger, N. Michael, and V. Kumar, "Trajectory generation and control for precise aggressive maneuvers with quadrotors," *Int. J. Robot. Research*, 2012.
- [6] E. Kaufmann, A. Loquercio, R. Ranftl, M. Müller, V. Koltun, and D. Scaramuzza, "Deep drone acrobatics," *RSS: Robotics, Science, and Systems*, 2020.
- [7] 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.
- [8] Y. Song and D. Scaramuzza, "Learning high-level policies for model predictive control," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2020.
- [9] S. Lupashin, A. Schöllig, M. Sherback, and R. D'Andrea, "A simple learning strategy for high-speed quadcopter multi-flips," in *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2010.
- [10] P. Ventura Diaz and S. Yoon, "High-fidelity computational aerodynamics of multi-rotor unmanned aerial vehicles," in *AIAA Aerospace Sciences Meeting*, 2018.
- [11] L. Hewing, J. Kabzan, and M. N. Zeilinger, "Cautious model predictive control using gaussian process regression," *IEEE Trans. Control Sys. Tech.*, 2019.
- [12] J. Kabzan, L. Hewing, A. Liniger, and M. N. Zeilinger, "Learning-based model predictive control for autonomous racing," *IEEE Robot. Autom. Lett.*, 2019.
- [13] D. Bicego, J. Mazzetto, R. Carli, M. Farina, and A. Franchi, "Nonlinear model predictive control with enhanced actuator model for multi-rotor aerial vehicles with generic designs," *J. Intell. Robot. Syst.*, 2020.
- [14] C. Liu, H. Lu, and W.-H. Chen, "An explicit mpc for quadrotor trajectory tracking," in *IEEE Chin. Control Conf. (CCC)*, 2015.
- [15] P. Ru and K. Subbarao, "Nonlinear model predictive control for unmanned aerial vehicles," *Aerospace*, 2017.
- [16] D. Lunni, A. Santamaria-Navarro, R. Rossi, P. Rocco, L. Bascetta, and J. Andrade-Cetto, "Nonlinear model predictive control for aerial manipulation," in *IEEE Int. Conf. Unm. Aircraft Syst. (ICUAS)*, 2017.
- [17] E. Tal and S. Karaman, "Accurate tracking of aggressive quadrotor trajectories using incremental nonlinear dynamic inversion and differential flatness," *IEEE Transactions on Control Systems Technology*, 2020.
- [18] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou, "Information theoretic MPC for model-based reinforcement learning," in *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2017.
- [19] D. Fan, A. Aghamohammadi, and E. Theodorou, "Deep learning tubes for tube mpc," *Robotics: Science and Systems (RSS)*, 2020.
- [20] K. Bieker, S. Peitz, S. L. Brunton, J. N. Kutz, and M. Dellnitz, "Deep model predictive flow control with limited sensor data and online learning," *Theoretical and Computational Fluid Dynamics*, 2020.
- [21] I. Lenz, R. A. Knepper, and A. Saxena, "DeepMPC: Learning deep latent features for model predictive control," in *Robotics: Science and Systems (RSS)*, 2015.
- [22] M. Saveriano, Y. Yin, P. Falco, and D. Lee, "Data-efficient control policy search using residual dynamics learning," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2017.
- [23] A. Carron, E. Arcari, M. Wermelinger, L. Hewing, M. Hutter, and M. N. Zeilinger, "Data-driven model predictive control for trajectory tracking with a robotic arm," *IEEE Robot. Autom. Lett.*, 2019.
- [24] 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.
- [25] S. Sun, C. C. de Visser, and Q. Chu, "Quadrotor gray-box model identification from high-speed flight data," *Journal of Aircraft*, 2019.
- [26] P.-J. Bristeau, P. Martin, E. Salaün, and N. Petit, "The role of propeller aerodynamics in the model of a quadrotor uav," in *IEEE Eur. Control Conf. (ECC)*, 2009.
- [27] P. Martin and E. Salaün, "The true role of accelerometer feedback in quadrotor control," in *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2010.
- [28] R. Mahony, V. Kumar, and P. Corke, "Multirotor aerial vehicles: Modeling, estimation, and control of quadrotor," *IEEE Robot. Autom. Mag.*, 2012.
- [29] M. Mehndiratta and E. Kayacan, "Gaussian process-based learning control of aerial robots for precise visualization of geological outcrops," in *IEEE Eur. Control Conf. (ECC)*, 2020.
- [30] G. Cao, E. M.-K. Lai, and F. Alam, "Gaussian process model predictive control of an unmanned quadrotor," *J. Intell. Robot. Syst.*, 2017.
- [31] V. R. Desaraju, A. E. Spitzer, C. O'Meadhra, L. Lieu, and N. Michael, "Leveraging experience for robust, adaptive nonlinear mpc on computationally constrained systems with time-varying state uncertainty," *Int. J. Robot. Research*, 2018.
- [32] M. Diehl, H. G. Bock, H. Diedam, and P. B. Wieber, "Fast direct multiple shooting algorithms for optimal robot control," in *Fast motions in biomechanics and robotics*. Springer, 2006, pp. 65–93.
- [33] R. Verschuere, G. Frison, D. Kouzoupis, N. van Duijkeren, A. Zanelli, R. Quirynen, and M. Diehl, "Towards a modular software package for embedded optimization," *IFAC-PapersOnLine*, 2018.
- [34] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi – A software framework for nonlinear optimization and optimal control," *Math. Prog. Comp.*, 2019.
- [35] N. Koenig and A. Howard, "Design and use paradigms for gazebo, an open-source multi-robot simulator," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2004.
- [36] 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.