**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

Performance, Precision, and Payloads: Adaptive Nonlinear MPC for Quadrotors

Hanover, Drew ; Foehn, Philipp ; Sun, Sihao ; Kaufmann, Elia ; Scaramuzza, Davide

Abstract: Agile quadrotor flight in challenging environments has the potential to revolutionize shipping, transportation, and search and rescue applications. Nonlinear model predictive control (NMPC) has recently shown promising results for agile quadrotor control, but relies on highly accurate models for maximum performance. Hence, model uncertainties in the form of unmodeled complex aerodynamic effects, varying payloads and parameter mismatch will degrade overall system performance. In this letter, we propose L1 -NMPC, a novel hybrid adaptive NMPC to learn model uncertainties online and immediately compensate for them, drastically improving performance over the non-adaptive baseline with minimal computational overhead. Our proposed architecture generalizes to many different environments from which we evaluate wind, unknown payloads, and highly agile flight conditions. The proposed method demonstrates immense flexibility and robustness, with more than 90% tracking error reduction over non-adaptive NMPC under large unknown disturbances and without any gain tuning. In addition, the same controller with identical gains can accurately fly highly agile racing trajectories exhibiting top speeds of 70 km/h, offering tracking performance improvements of around 50% relative to the non-adaptive NMPC baseline.

DOI: <https://doi.org/10.1109/LRA.2021.3131690>

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

ZORA URL: <https://doi.org/10.5167/uzh-216662>

Journal Article

Accepted Version

Originally published at:

Hanover, Drew; Foehn, Philipp; Sun, Sihao; Kaufmann, Elia; Scaramuzza, Davide (2022). Performance, Precision, and Payloads: Adaptive Nonlinear MPC for Quadrotors. *IEEE Robotics and Automation Letters*, 7(2):690-697.

DOI: <https://doi.org/10.1109/LRA.2021.3131690>

Performance, Precision, and Payloads: Adaptive Nonlinear MPC for Quadrotors

Drew Hanover, Philipp Foehn, Sihao Sun, Elia Kaufmann, Davide Scaramuzza

Abstract—Agile quadrotor flight in challenging environments has the potential to revolutionize shipping, transportation, and search and rescue applications. Nonlinear model predictive control (NMPC) has recently shown promising results for agile quadrotor control, but relies on highly accurate models for maximum performance. Hence, model uncertainties in the form of unmodeled complex aerodynamic effects, varying payloads and parameter mismatch will degrade overall system performance. In this paper, we propose $\mathcal{L}_1$ -NMPC, a novel hybrid adaptive NMPC to learn model uncertainties online and immediately compensate for them, drastically improving performance over the non-adaptive baseline with minimal computational overhead. Our proposed architecture generalizes to many different environments from which we evaluate wind, unknown payloads, and highly agile flight conditions. The proposed method demonstrates immense flexibility and robustness, with more than 90% tracking error reduction over non-adaptive NMPC under large unknown disturbances and without any gain tuning. In addition, the same controller with identical gains can accurately fly highly agile racing trajectories exhibiting top speeds of 70 km/h, offering tracking performance improvements of around 50% relative to the non-adaptive NMPC baseline.

SUPPLEMENTARY MATERIAL

Video: <https://youtu.be/8oB1rG5iYc4>

I. INTRODUCTION

A. Motivation

Unmanned aerial vehicle (UAV) utilization in industrial applications is increasing at an astounding rate [1–3]. There are over 300,000 commercial drones registered in the US alone according to the Federal Aviation Administration [4]. Significant commercial opportunities exist for robust autonomous systems which can safely conduct inspection of sensitive, hazardous, and remote systems under uncertainty. As a result, the global drone service market is expected to grow from 4.4 Billion USD to 63.6 Billion between 2018 and 2025 [5]. In transportation and shipping applications, UAVs can improve efficiency of operations via autonomous missions and high speed maneuvers leading to dramatic time

Fig. 1: Top: Quadrotor carrying a beer payload of unknown mass while flying an aggressive racing trajectory. Our proposed $\mathcal{L}_1$ -NMPC method allows to reduce tracking errors on such agile trajectories by more than 90% compared to non-adaptive methods. Bottom: $\mathcal{L}_1$ -NMPC control diagram.

savings, cost reductions, and higher throughput [6]. On a similar note, emergency scenarios raised the need for agile autonomous systems to conduct search and rescue tasks when time is of the essence [7]. To facilitate these opportunities, UAVs must be able to accurately track agile trajectories in presence of model uncertainties and external disturbances such as unknown drag coefficients, varying payloads, or wind gusts respectively. These uncertainties can significantly degrade the performance and reliability of the system and potentially lead to loss of control if not compensated for. High-fidelity physics-based models can improve control performance, but are often prohibitively expensive to procure and require extensive levels of domain expertise [8–11]. With the advancements of data-driven methods such as those described in [8–10], the costs of obtaining accurate models has been dramatically reduced. Models learned from data, however, have a tendency to overfit and can be intractable to update online [11].

In reality, modeling the system with complete accuracy is impossible, necessitating control algorithms which are robust to uncertainty. The field of robust control has tried to address this problem with varying degrees of success. Classical methods such as H_∞ and more recently stochastic model predictive control may enable safe behavior of the autonomous system, but at the cost of significant performance degradation [12, 13]. Predicting external disturbances such as wind-gusts is often impractical due to the chaotic nature of the dis-

Manuscript received: September, 8, 2021; Accepted November, 10, 2021.

This paper was recommended for publication by Editor Pauline Pounds upon evaluation of the Associate Editor and Reviewers' comments. 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): 10.1109/LRA.2021.3131690

turbance [10, 14]. This necessitates the development of highly adaptive control algorithms which can robustly compensate for unknown model dynamics without prior assumptions, while pushing the system to its dynamic limits.

B. Contribution

In this work, we propose a novel quadrotor control architecture which cascades a nonlinear model predictive controller (NMPC) to an $\mathcal{L}_1$ adaptive controller to fly highly aggressive racing trajectories under various types of model uncertainties and external disturbances with speeds up to 20 m s^{-1} . We show that the addition of the adaptive controller can drive the real system towards the behavior specified by the underlying MPC model with significantly better tracking performance when compared to a non-adaptive baseline. The proposed controller can compensate in real-time, while being transferable to many different applications such as flying in windy environments, carrying completely unknown payloads, and flying aggressive racing trajectories without any re-tuning of control gains. When unknown payloads up to 60% of the quadrotor mass are introduced into the system, our approach demonstrates a reduction of tracking error over 90% compared to non-adaptive NMPC methods onboard a real quadrotor. Additionally, we show that it is even possible to accurately fly racing trajectories with unknown slung payloads. The data indicates we can track these trajectories *with* an unknown payload representing 13% of the quadrotor's mass with 44% higher accuracy than a non-adaptive NMPC *without* a payload attached. Our experiments under nominal model conditions demonstrate at least a 50% tracking performance improvement over state of the art data-driven MPC methods *without* an aerodynamics model on a set of increasing speed circle trajectories which exhibit speeds of up to 36 km/h, indicating that the adaptive component is able to learn and compensate for unmodeled aerodynamic effects in real time with minimal computational overhead.

II. RELATED WORK

Trajectory tracking controllers for quadrotors have been studied extensively over the last decade. A detailed survey on quadrotor control methods can be found in [15–17]. Most of the existing literature focuses on hover conditions or slow speed maneuvers which satisfy small angle assumptions necessary for linear control methods. We are interested in exploring a wider flight envelope, and focus on advanced techniques to push the physical limits of the platform.

Agile flight of aerial vehicles has been a top priority for the aerospace industry for the better part of 80 years, as a part of which NASA developed Model Reference Adaptive Control (MRAC) to deal with large model uncertainties that are difficult to model and measure [18]. Readers interested in the underlying mathematics of adaptive control methods for aerospace vehicles are pointed to the following references [19–22]. Specifically, we focus on the $\mathcal{L}_1$ adaptive control approach due its inherent ability to provide rapid adaptation that is decoupled from the robustness of the controller [23]. Applications of $\mathcal{L}_1$ adaptive controllers have been successfully

demonstrated across a variety of aerial vehicles such as fixed-wings, quadcopters, and octocopters [24–26]. The main feature of the $\mathcal{L}_1$ adaptive controller is to drive a system towards a desired reference model behavior. Typically, this is done using a linear reference model to specify the desired behavior, however this can lead to unrealistic desired dynamics which cannot be achieved by the real system. Adaptive control has successfully demonstrated accurate trajectory tracking using quadrotors in several works [27–30]. However, the maneuvers conducted are typically simple step inputs, or slow speed circles which do not exploit the inherent agility of the quadrotor platform.

Adaptive controllers often act as an augmentation to an existing baseline controller rather than as a standalone controller. Authors in [31] take advantage of the high level planning of linear MPC, cascaded with an adaptive control law to adapt to persistent model mismatch. A cascaded linear MPC with a linear reference model $\mathcal{L}_1$ adaptive controller was demonstrated for the quadrotor trajectory tracking problem with exogenous disturbances in [32], but the trajectories demonstrated were simplistic and slow. They claim a reduced computational cost compared to nonlinear MPC frameworks, however the limitations of linear optimal control when applied to nonlinear problems are well understood, especially when the vehicle exhibits highly aggressive maneuvers [33].

Similarly, in [34] a Model Predictive Path Integral (MPPI) controller was coupled with a nonlinear reference model $\mathcal{L}_1$ adaptive controller for agile quadrotor flight, however the authors have not shown feasibility of the proposed method on real hardware. No analysis of the adaptive control signal is provided, however video footage released of the simulation performance of the proposed $\mathcal{L}_1$ -MPPI architecture indicates highly oscillatory control performance in the first-person camera view¹. Additionally, the high level MPPI controller can only run at a rate of 50 Hz using desktop hardware which makes it infeasible for on-board control.

Previous works demonstrating agile trajectory tracking with physical quadrotors include [35, 9, 36, 37]. The authors of [35] demonstrated accurate tracking of aggressive quadrotor trajectories up to 12.9 m s^{-1} using a cascaded geometric controller with Incremental Nonlinear Dynamic Inversion (INDI). In [9], NMPC leveraging data driven methods to improve model fidelity was used to achieve state of the art tracking performance at speeds up to 14 m s^{-1} . Control commands were calculated off-board the quadrotor and sent via wireless communication due to the increased computational overhead from the learned model. Because the model parameters are obtained offline, it can not adapt to online parametric changes such as payloads or a reduction in actuator efficacy.

As the authors in [9] correctly point out, INDI and adaptive control approaches coupled with traditional geometric controllers are purely reactive and have no ability to plan over a prediction horizon. To address this, we couple an $\mathcal{L}_1$ adaptive control law with a nonlinear MPC, therefore taking advantage of the prediction horizon and fully exploiting the nonlinearities of the system for maximum performance. This selection is further motivated by the recent comparative study between

¹<https://youtu.be/f602VSGIVb0>

geometric and nonlinear MPC methods for agile quadrotor control in [16]. The controller is evaluated against several state of the art controllers and a wide variety of test conditions both in simulation and reality.

III. METHODOLOGY

A. Notation

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. All four rotors are assumed to be located on 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 $\mathbf{q} \odot \mathbf{v} = \mathbf{q} \cdot [0, \mathbf{v}^\top]^\top \cdot \bar{\mathbf{q}}$, where $\bar{\mathbf{q}}$ is the quaternion's conjugate.

B. Quadrotor Vehicle Dynamics

The quadrotor system dynamics are given by

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

where $\mathbf{g} = [0, 0, -9.81 \text{ m/s}^2]^\top$ denotes Earth's gravity, $\mathbf{T}_B$ is the collective thrust from the 4 rotors, $\mathbf{J} = \text{diag}(J_x, J_y, J_z)$ is the diagonal moment of inertia matrix, m is the quadrotor mass, and $\boldsymbol{\tau}_B$ is the body torque:

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

Where $\mathbf{P}$ is the thrust allocation matrix given by

$$\mathbf{P} = \begin{bmatrix} -d_{x0} & -d_{x1} & d_{x2} & d_{x3} \\ d_{y0} & -d_{y1} & -d_{y2} & d_{y3} \\ -c_\tau & c_\tau & -c_\tau & c_\tau \end{bmatrix} \quad (3)$$

where d_{x_i}, d_{y_i}, c_τ for $i = 0, 1, 2, 3$ are the distances from each rotor, i , to the respective Body frame axis and the rotor drag torque constant respectively following the rotor positions and spin configurations denoted in Figure 2 and obeying the right hand rule. For discrete time, an explicit Runge-Kutta method of 4th order is used:

$$\mathbf{x}_{k+1} = f_{RK4}(\mathbf{x}_k, \mathbf{u}_k, \delta t). \quad (4)$$

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

C. MPC Formulation

We construct a quadratic optimization problem using a multi-shooting scheme and solve the following discretized nonlinear optimal control problem

$$\begin{aligned} \min_{\mathbf{u}} \quad & \mathbf{x}_N^T \mathbf{Q} \mathbf{x}_N + \sum_{k=0}^{N-1} \mathbf{x}_k^T \mathbf{Q} \mathbf{x}_k + \mathbf{u}_k^T \mathbf{R} \mathbf{u}_k \\ \text{subject to:} \quad & \mathbf{x}_{k+1} = \mathbf{f}_{RK4}(\mathbf{x}_k, \mathbf{u}_k, \delta t) \\ & \mathbf{x}_0 = \mathbf{x}_{init} \quad \mathbf{u}_{min} \leq \mathbf{u}_k \leq \mathbf{u}_{max} \end{aligned} \quad (5)$$

as a sequential quadratic program (SQP) executed in a real-time iteration scheme [38]. We discretize the system evolution into N steps over a time horizon T and constrain the input to be between $\mathbf{0} \leq \mathbf{u}_k \leq \mathbf{u}_{max}$. The optimal control problem is implemented using the open source ACADO toolkit [39].

D. $\mathcal{L}_1$ -Adaptive Augmentation

We implement the $\mathcal{L}_1$ adaptive controller using a nonlinear reference model [40], which estimates both matched and unmatched uncertainties using a piecewise constant adaptation law [41, 42]. The derivation is similar to [34], however we account for the uncertainties directly at the rotor thrust level. First, define $\mathbf{R}_B^I = [\mathbf{e}_x^B, \mathbf{e}_y^B, \mathbf{e}_z^B]$ as the rotation matrix from the body frame to the inertial frame. We can then rewrite the dynamics to account for both matched and unmatched uncertainties as

$$\dot{\mathbf{v}}_{WB} = \frac{\mathbf{T}_B}{m} \mathbf{e}_z^B + \mathbf{g} + \frac{\mathbf{T}_B}{m} \mathbf{R}_B^I \boldsymbol{\varsigma} - \mathbf{C}_d \mathbf{v}_{WB}, \quad (6)$$

$$\dot{\boldsymbol{\omega}}_B = \mathbf{J}^{-1}[\boldsymbol{\tau}_B - \boldsymbol{\omega}_B \times \mathbf{J}\boldsymbol{\omega}_B + \boldsymbol{\xi}], \quad (7)$$

where $\boldsymbol{\varsigma} = [\varsigma_x, \varsigma_y, \varsigma_z]^\top$ is the uncertainty appearing in the linear accelerations, $\boldsymbol{\xi} = [\xi_x, \xi_y, \xi_z]^\top$ is the uncertainty appearing in the angular accelerations, and $\mathbf{C}_d$ is a matrix of linear drag coefficients. Since a quadrotor is an underactuated system, capable of providing linear acceleration only along its body z -axis, the unmatched uncertainties defined as $\boldsymbol{\sigma}_{um} = [\varsigma_x, \varsigma_y]^\top$, appear purely in the X and Y linear accelerations. This can be thought of as existing in the null space of the controllability matrix and therefore cannot be compensated for directly. Then, what remains are the matched uncertainties $\boldsymbol{\sigma}_m = [\varsigma_z, \xi_x, \xi_y, \xi_z]^\top$ which can be compensated for directly.

Next, consider the reduced state variable, $z = [v_{WB}, \omega_B]$. Its derivative can be broken up as a function of the nominal and uncertain dynamic behavior as follows:

$$\dot{z} = f(R_B^I) + g(R_B^I)(u_{L1} + \sigma_m) + g^\perp(R_B^I)\sigma_{um}, \quad (8)$$

where $f(R_B^I)$ is the desired dynamics defined as

$$f(R_B^I) = \begin{bmatrix} g + \frac{T_{MPC}}{m} e_z^B - C_d v_{WB} \\ J^{-1} \tau_B - \omega_B \times J \omega_B \end{bmatrix}. \quad (9)$$

The single rotor thrusts in expression of τ_B are obtained from the solution of the NMPC at the current time.

Define $g(R_B^I)$ as the uncertainty in the matched component of the dynamics, and $g^\perp(R_B^I)$ the unmatched component:

$$g(R_B^I) = \begin{bmatrix} \frac{e_z^B}{m}, \frac{e_z^B}{m}, \frac{e_z^B}{m}, \frac{e_z^B}{m} \\ J^{-1} P \end{bmatrix} \quad g^\perp(R_B^I) = \begin{bmatrix} \frac{e_x^B}{m}, \frac{e_y^B}{m} \\ 0_{3 \times 1}, 0_{3 \times 1} \end{bmatrix}. \quad (10)$$

Let u_{L1} be the adaptive control input which can act as a standalone controller, or complement the NMPC signal via addition. Define the $\mathcal{L}_1$ observer as

$$\dot{\tilde{z}} = f(R_B^I) + g(R_B^I)(u_{L1} + \hat{\sigma}_m) + g^\perp(R_B^I)\hat{\sigma}_{um} + A_s \tilde{z},$$

where $\tilde{z} = \hat{z} - z$ and noting that z is the state obtained from an estimator, and $\hat{z}$ is the state predicted from the $\mathcal{L}_1$ observer. Define $\Phi = A_s^{-1}(e^{A_s T_s} - I)$ where T_s is the time step and A_s is a Hurwitz matrix which represents the adaptation gains. Then the piecewise-constant adaptation law is given by

$$\begin{bmatrix} \hat{\sigma}_m(iT_s) \\ \hat{\sigma}_{um}(iT_s) \end{bmatrix} = -I_{6 \times 6} G^{-1}(iT_s) \Phi^{-1} \mu(iT_s), \quad (11)$$

where $G(iT_s) = [g(R_B^I), g^\perp(R_B^I)]$ and $\mu = e^{A_s T_s} \tilde{z}(iT_s)$ are evaluated at time step i . Next, define a first order, strictly proper continuous time filter $C(s)$. The $\mathcal{L}_1$ control law is then

$$u_{L1} = -C(s) \hat{\sigma}_m. \quad (12)$$

In practice, we implement the control law in discrete time as

$$u_{L1,k} = u_{L1,k-1} e^{-\omega_{co} T_s} - \hat{\sigma}_{m,k} (1 - e^{-\omega_{co} T_s}), \quad (13)$$

where ω_{co} is the cutoff frequency of the strictly proper first order filter. Finally, the discrete time $\mathcal{L}_1$ observer can be propagated forward in time via

$$\hat{z}_{k+1} = \hat{z}_k + [f_k + g_k(u_{L1,k} + \hat{\sigma}_{m,k}) + g_k^\perp \hat{\sigma}_{um,k} + A_s \tilde{z}_k] T_s.$$

IV. EXPERIMENTS AND RESULTS

Our experiments are designed to answer the following research questions: (i) How does our proposed $\mathcal{L}_1$ -NMPC compare to data-driven MPC methods? (ii) To what extent can $\mathcal{L}_1$ -NMPC react to both parametric and non-parametric disturbances in real world tests? (iii) Does the proposed $\mathcal{L}_1$ -NMPC generalize across test scenarios without the need for gain tuning? We set out to answer these questions through a variety of simulation and real world tests which cover large external disturbances and agile maneuvers. The experiments start out with a set of different model predictive controllers which are gradually eliminated over the course of the experiments based on their tracking performance and computational complexity. We point the reader to our corresponding video for a better understanding of the level of flexibility and robustness our proposed architecture can demonstrate.

Avg. dt [ms]	Model				
	GP-MPC	MPPI	NMPC	INDI-NMPC	L1-NMPC
	4.13	23.13	0.81	0.82	0.82

TABLE I: Average controller update times running on an Intel Core i7-8750H CPU @ 2.20GHz laptop with 16 Gb of RAM and Nvidia GeForce GTX 1060 using CUDA 11.2

A. Simulation

We begin by implementing several state of the art controllers in simulation including a Single Rotor Thrust NMPC (SRT-NMPC), a data-driven NMPC (GP-MPC) from [9], MPPI with Baseline Control from [34], INDI-NMPC from [16], and our proposed $\mathcal{L}_1$ -NMPC with and without an aerodynamic model included in the underlying NMPC. All baseline controllers have an aerodynamic model enabled by default. We use the open source Gazebo simulator [43] with the the AscTec Hummingbird quadrotor model using the RotorS extension [44]. Performance in simulation is measured by comparing positional tracking errors on a set of circular reference trajectories. These trajectories feature a radius of 5 m and vary in peak velocities from 2.5-10 m s⁻¹.

1) *Without Disturbance*: First, we fly these trajectories with a model that best represents our knowledge of the dynamics within the RotorS simulation. In these cases, the mass, inertia, drag, and rotor arm lengths of the simulated quadrotor are perfectly known. The model parameters used in the NMPC are identical to the parameters in the RotorS simulator. These experiments without disturbance serve to understand the maximum achievable performance by each method in case of perfect model identification.

We include the timings for each of the tested controllers in Table I and show the tracking performance on a semi-log scale for the increasing speed circle trajectories in Figure 3. Our $\mathcal{L}_1$ -NMPC *without* an aerodynamic model outperforms the state of the art INDI-NMPC and GP-MPC in all cases except for the fastest circle trajectories, indicating that the adaptation law is able to compensate partially for unmodeled aerodynamics.

Once we enable the drag model, the $\mathcal{L}_1$ -NMPC outperforms all proposed methods, however we note that the performance improvement is less than 1 cm RMSE over the course of a 60 s long trajectory. We expect the performance benefit to be small since the model parameters perfectly match those described in the simulator. In contrast to all other approaches, MPPI performed an order of magnitude worse. We therefore do not consider it as a viable candidate moving forward. SRT-NMPC with a linear drag model performs similarly to GP-MPC, but with 80% less computational overhead and therefore forms the baseline for the remaining simulation trials.

2) *Model Mismatch*: Next, we inject three different forms of model mismatch into the model used for the NMPC. The same increasing speed circle trajectories from the previous section are used in this analysis. First, the mass of the quadrotor is increased by 660 g, representing a 90% increase from the nominal mass of the system. Second, the inertias on all axis are doubled in the simulator. Finally, the right side rotor arm lengths are reduced by 25%, representing a persistent center of gravity offset.

Fig. 3: Simulated tracking accuracy of increasing speed, on a 5 m radius circle trajectory in the RotorS simulation environment. Tracking accuracy is given by RMSE from the reference trajectory. Each point represents the accumulated positional tracking error across the trajectory for a given peak velocity.

		Model							
		NMPC		NMPC+I		INDI-NMPC		L1-NMPC	
Dist.	v_{peak} [m s ⁻¹]	RMSE [m]	RMSE [m]	%↓	RMSE [m]	%↓	RMSE [m]	%↓	
Mass	2.5	0.434	0.029	93	0.434	0	0.007	98	
	4	0.441	0.045	90	0.440	0	0.010	98	
	6	0.471	0.091	81	0.470	0	0.020	96	
	8	0.532	0.167	69	0.537	0	0.032	94	
	10	Crash	Crash	?	Crash	?	Crash	?	
Inertia	2.5	0.02	0.033	-65	0.009	55	0.007	65	
	4	0.016	0.043	-169	0.015	6	0.014	13	
	6	0.037	0.080	-116	0.028	24	0.030	19	
	8	0.070	0.150	-114	0.048	31	0.052	26	
	10	0.081	0.240	-196	0.074	9	0.079	2	
Arm Length	2.5	0.078	0.056	28	0.008	90	0.007	91	
	4	0.049	0.038	22	0.015	69	0.016	67	
	6	0.052	0.088	-69	0.032	38	0.033	37	
	8	0.057	0.151	-165	0.053	7	0.056	2	
	10	0.083	0.241	-190	0.083	0	0.087	-5	

TABLE II: Tracking performance of increasing speed, 5 m radius circle trajectories with mass, inertia, and rotor arm length disturbances in the RotorS simulator. Each table entry represents the tracking RMSE over the course of a single trajectory for different speeds. In the Mass cases, we increase the mass of the quadrotor by 660 g representing a 90% increase in mass. Similarly, in the Inertia cases, we double the inertia of the quadrotor. Finally, in the Arm Length cases, we increase the length of the right side rotor arms by 25% representing a significant shift to the center of gravity.

Since the baseline NMPC does not contain any integrator action, we expect any pure-mass disturbances to result in a steady-state offset in target Z-height. In an attempt to make a fair comparison, we add an additional comparison case which embeds an integrator state on positional error into the NMPC. The gains for all controllers are not adjusted from the nominal tests in the previous section. Table II provides the results for these parametric disturbance cases.

As parametric disturbances are introduced into the test scenarios, the proposed $\mathcal{L}_1$ -NMPC architecture demonstrates a sizeable advantage over the other state of the art methods. Only the NMPC with integrator action and $\mathcal{L}_1$ -NMPC show any ability to compensate for mass mismatch, with $\mathcal{L}_1$ -NMPC reducing the tracking error by over 90% in all cases. We note that all controllers fail to complete the 10 m s⁻¹ circle

trajectory with a 660 g payload. The thrust requirements to fly this case exceed what is available on the simulated quadrotor which leads to the NMPC solution failing to converge in all cases. In the inertia and rotor arm length test cases, INDI-NMPC and $\mathcal{L}_1$ -NMPC perform almost identically. The performance delta between the two methods is less than 5 mm across these 10 cases. Our adaptive architecture demonstrates robustness to these uncertainties in addition to providing a performance benefit when the model is well known. We are able to immediately identify various forms of disturbance and swiftly reject it without any gain tuning or model learning.

B. Real World Experiments

We test our proposed $\mathcal{L}_1$ -NMPC controller performance on a quadrotor with mass 750 g outfitted with a Jetson TX2 on-board computer and Radix flight controller with our own custom low-level flight control firmware. The flight controller accepts single rotor thrust inputs and performs closed loop rotor speed control. The quadrotor has a thrust to weight ratio of about 4.5. We run all controllers completely onboard the quadrotor and solve the optimal control problem at 100 Hz. A Vicon motion capture system² provides pose updates at 400 Hz which are fused with the Inertial Measurement Unit (IMU) via an Extended Kalman Filter (EKF) for state estimation.

To demonstrate our approach can significantly improve tracking performance under model mismatch and aerodynamic disturbance, we conduct experiments in five settings:

- Setting (i): Increasing speed 5 m radius circle trajectories up to 10 m s⁻¹ to compare performance to GP-MPC and INDI-NMPC in a nominal setting.
- Setting (ii): 2 m s⁻¹ circle trajectory with a 450 g unknown payload to show tracking performance when a large mass mismatch is present.
- Setting (iii): 2 m s⁻¹ circle with external aerodynamic forces to show rapid disturbance rejection.
- Setting (iv): Mildly aggressive flight with a top speed of 11.9 m s⁻¹ with a 100 g slung payload to show agile flight is possible with unknown payloads.
- Setting (v): Highly aggressive flight with a top speed of 19.4 m s⁻¹ without payload to demonstrate capability of accurate tracking near the system limits.

In Setting (i), we perform the same increasing speed circle trajectories from the simulation experiments and compare the tracking performance of our method to GP-MPC, SRT-NMPC, and INDI-NMPC. Table III shows the results of this comparison. As can be seen, even without an aerodynamic model, our $\mathcal{L}_1$ -NMPC architecture outperforms GP-MPC by over 70% and has a slight advantage over INDI-NMPC, indicating that the adaptive controller is compensating for both model mismatch and aerodynamic disturbances. Additionally, we show that SRT-NMPC with a linear aerodynamic model matches the performance of GP-MPC and therefore use SRT-NMPC as the baseline moving forward due to its computational advantage over GP-MPC.

Next, in Setting (ii) we perform a 2 m s⁻¹ circle with a 450 g payload attached. This represents a mass increase of over

²<https://www.vicon.com/>

		Model									
		GP-MPC		SRT-NMPC		INDI-NMPC		LI-NMPC No Aero		LI-NMPC w/ Aero	
v_{peak} [m s ⁻¹]		RMSE [m]		RMSE [m]	% ↓	RMSE [m]	% ↓	RMSE [m]	% ↓	RMSE [m]	% ↓
2.5		0.109		0.070	36	0.038	65	0.022	80	0.016	85
4		0.103		0.086	17	0.052	50	0.026	75	0.021	80
6		0.129		0.103	20	0.055	57	0.035	73	0.034	74
8		0.154		0.153	1	0.058	62	0.057	63	0.041	73
10		0.203		0.223	-10	0.048	76	0.088	57	0.047	77

TABLE III: Tracking performance of increasing speed circle trajectories evaluated in real-world scenario (i).

60%. The MPC model is not updated to reflect the changed mass and inertia of the system and the adaptive controller must compensate for a large disturbance from takeoff. Fig. 4 shows the tracking performance of the quadrotor for scenarios (ii) and (iii). Relative to the non-adaptive baseline, the $\mathcal{L}_1$ -NMPC rejects the disturbance and accurately tracks the specified trajectory. It is important to note that the adaptation to the unknown payload takes place rapidly, resulting in less than 1 cm steady state Z position tracking error. In fact, the adaptive controller immediately compensates for the unknown payload by applying higher thrust magnitudes to each rotor upon takeoff, as can be seen in Figure 5. In contrast, the non-adaptive MPC has over a 35 cm steady state Z position error and cannot compensate for the modified inertia of the platform, resulting in overshoots straying away from the reference.

In Setting (iii) we fly the same trajectory without payloads but instead add an external aerodynamic disturbance in the form of a powerful fan which the quad has to fly directly in front of. The baseline NMPC immediately demonstrates sizeable tracking performance degradation, while the $\mathcal{L}_1$ -NMPC is able to compensate for these disturbances in real time and maintain acceptable, albeit reduced performance. In addition to being able to compensate for large parametric disturbances, the $\mathcal{L}_1$ -NMPC substantially reduces the tracking error when an exogenous disturbance is introduced. We believe this has large implications for quadrotors conducting both inspection and shipping tasks outdoors in windy environments.

For Settings (iv) and (v), we fly a highly dynamic race track trajectory generated using the method described in [6]. We compare our proposed approach against the Non-Adaptive MPC and INDI-NMPC in these settings due to the latter's demonstrated success on aggressive racing trajectories. In Setting (iv), we attach an unknown 100 g slung payload to introduce a time varying disturbance and fly a mildly aggressive trajectory with peak velocities of 11.9 m s^{-1} and peak linear accelerations of over 2 g's. With the unknown slung load attached, the $\mathcal{L}_1$ -NMPC controller reduces tracking error by 44% relative to the Non-Adaptive NMPC *without* a payload attached. Additionally, we outperform INDI-NMPC by 34% in tracking error, and show that the performance of $\mathcal{L}_1$ -NMPC with and without payloads are nearly identical. This indicates that the adaptive controller is driving the true system dynamics towards those specified by the reference system, regardless of present disturbances.

In Setting (v), we do not introduce disturbances but increase the aggressiveness of the trajectory. The drone experiences over 4 g's of linear acceleration, and experiences peak linear velocities of around 20 m s^{-1} . The tracking performance for each of these scenarios can be seen in Figure 6. The $\mathcal{L}_1$ -

Fig. 4: Scenarios (ii) and (iii): Slow speed circle trajectory with induced disturbances. (Top) A payload of 450 g is added to the quadrotor without updating the nominal NMPC model. The adaptive controller learns the disturbance and immediately compensates for it. (Bottom) An external aerodynamic disturbance is added in the form of a fan and force the quadrotor to fly directly through the turbulence generated from the fan. The adaptive controller compensates immediately for the disturbance and keep the quadrotor close to the reference trajectory.

Fig. 5: Thrust commands for Scenario (ii). (Top Left) MPC Commands. (Top Right) $\mathcal{L}_1$ -MPC Commands. (Bottom Left) First second of thrust commands from MPC, resulting in large steady state offset in Z-position. (Bottom Right) First second of thrust commands from $\mathcal{L}_1$ -MPC resulting in immediate adaptation to the unknown payload.

NMPC shows a performance improvement of 49% of relative to the baseline SRT-NMPC but is outperformed by INDI-NMPC by less than 5 cm RMSE. The INDI-NMPC controller is a highly capable racing controller, but relies on an accurate aerodynamic model to achieve such levels of performance. With our proposed method we do not require such an accurate model to achieve high tracking accuracy as demonstrated in Scenario (i). Beyond this, INDI-NMPC requires that the model be updated anytime there is a mass change in order to maximize performance, whereas our controller can learn

and compensate the disturbance without the need of updating model parameters or controller gains. Across all of these experiments no adjustments are made to the high level MPC, nor the adaptive controller. Even with this constraint, we are able to accurately track a variety of trajectories with various unknown disturbances. This demonstrates substantial flexibility and robustness when deploying the quadrotor in uncertain scenarios for minimal computational overhead.

V. DISCUSSION

Through extensive experimentation, the $\mathcal{L}_1$ adaptive controller demonstrably enhances a baseline model predictive controller by enabling immediate compensation for model mismatch and external disturbances. Cascading the baseline NMPC with an adaptive controller provides substantial benefits in tracking performance across all test cases, especially when there are model mismatches present. The performance benefit from the adaptation law is at least 50% greater than the data driven model augmentation provided by GP-MPC, without the large computational overhead. When implementing $\mathcal{L}_1$ -MPPI, we saw the same high frequency oscillatory content in the body rates that is evident in the videos released with [34]. While MPPI has shown outstanding results on ground vehicles [45, 46], we believe the sampling approach suffers from the curse of dimensionality problem induced by the 4D space. We conclude this discussion with the following recommendation: If a user requires aggressive maneuvers and accurate tracking on trajectories with over 4g linear accelerations and there is high confidence in the model accuracy, INDI-NMPC offers the best performance over all controllers considered for onboard computation. In fact, INDI-NMPC provides the highest autonomous quadrotor racing performance of any control architecture previously studied in the literature to date as indicated by the results in this study and the results in [16]. On the other hand, if the user is operating in uncertain environments which may cause large disturbances and require transportation of various payloads without the ability to update the model, $\mathcal{L}_1$ -NMPC offers over a 90% performance improvement over non-adaptive NMPC and INDI-NMPC. Like the non-adaptive NMPC, INDI-NMPC does not offer any integrative action on linear accelerations and therefore will always have non-zero steady state error in Z-height tracking when payloads are present. This is evidenced by the results in Scenario (iv) where we showed our proposed $\mathcal{L}_1$ -NMPC outperforming INDI-NMPC when an unknown slung payload is introduced into the system.

VI. CONCLUSION

In this work, we proposed augmenting a multi-shooting NMPC with an $\mathcal{L}_1$ adaptive inner-loop controller to compensate for model mismatch which can significantly degrade the performance of the NMPC. We derived the adaptation law at the individual rotor thrust level and leveraged the same nonlinear dynamics model that is embedded in the model predictive controller. The adaptive controller demonstrated minimal computational overhead, and can compute corrective control signals in only 10 microseconds which drastically

improves the system performance relative to the non-adaptive NMPC baseline. We showed that our proposed method can outperform state of the art data driven MPC methods in both simulation and extensive real world testing on the quadrotor tracking problem. Instead of having to retrain data driven models online, we instead simply adapt in real time to both parametric and non-parametric disturbances. Finally, we have shown that the proposed $\mathcal{L}_1$ -NMPC can fly highly aggressive racing trajectories while also being able to carry payloads with no additional gain tuning. One drawback of our proposed approach is the potential for violating actuator constraints due to the inner loop cascade. In the future, we plan to explore using the uncertainty estimation from the adaptation law to update our nominal dynamics model online. We hypothesize that this could significantly improve the prediction accuracy of the NMPC, thus enabling a highly adaptive optimal controller which can fully obey state and input constraints.

VII. ACKNOWLEDGMENT

The authors thank Leonard Bauersfeld for taking the image used in Figure 1 and his advice on editing the accompanying video. Additionally, the authors thank Thomas Laengle for his hardware troubleshooting in the early stages of this project.

REFERENCES

- [1] G. Loianno and D. Scaramuzza, "Special issue on future challenges and opportunities in vision-based drone navigation," *Journal of Field Robotics*, 2020.
- [2] S. Rajendran and S. Srinivas, "Air taxi service for urban mobility: A critical review of recent developments, future challenges, and opportunities," *Transportation Research Part E: Logistics and Transportation Review*, vol. 143, p. 102090, Nov 2020.
- [3] E. Ackerman, "Skydio demonstrates incredible obstacle-dodging full autonomy with new r1 consumer drone," *IEEE Spectrum*, 2018. [Online]. Available: <http://spectrum.ieee.org/automaton/robotics/drones/skydio-r1-drone>
- [4] "Secondary navigation," Aug 2021. [Online]. Available: https://www.faa.gov/uas/resources/by_the_numbers/
- [5] Newswire, "Global drone service market report 2019: Market is expected to grow from usd 4.4 billion in 2018 to usd 63.6 billion by 2025, at a cagr of 55.9%," Apr 2019. [Online]. Available: <https://bit.ly/3zUewze>
- [6] 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>
- [7] M. Faessler, F. Fontana, C. Forster, E. Mueggler, M. Pizzoli, and D. Scaramuzza, "Autonomous, vision-based flight and live dense 3D mapping with a quadrotor MAV," *Journal of Field Robotics*, vol. 33, no. 4, pp. 431–450, 2016.
- [8] 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.
- [9] G. Torrente, E. Kaufmann, P. Foehn, and D. Scaramuzza, "Data-driven mpc for quadrotors," *IEEE Robotics and Automation Letters*, 2021.
- [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] L. Bauersfeld, E. Kaufmann, P. Foehn, S. Sun, and D. Scaramuzza, "Neurobem: Hybrid aerodynamic quadrotor model," *RSS: Robotics, Science, and Systems*, 2021.
- [12] K. Zhou and J. C. Doyle, *Essentials of robust control*. Prentice hall Upper Saddle River, NJ, 1998, vol. 104.
- [13] I. R. Petersen, V. A. Ugrinovskii, and A. V. Savkin, *Robust control design using H-inf methods*. Springer Science & Business Media, 2012.
- [14] B. Nisar, P. Foehn, D. Falanga, and D. Scaramuzza, "VIMO: Simultaneous visual inertial Model-Based odometry and force estimation," *IEEE Robot. Autom. Lett.*, vol. 4, no. 3, pp. 2785–2792, Jul. 2019.

Fig. 6: Scenarios (iv) and (v): Tracking error over the course of a dynamic race track. The adaptive controller enforces the quadrotor to fly just as the NMPC model describes it should, regardless if there is a payload present. We demonstrate that when an unknown payload is attached, the $\mathcal{L}_1$ -NMPC controller outperforms the non-adaptive baseline *without* a payload attached on an agile racing trajectory with peak speeds of 11.9 m s^{-1} (left), and can achieve high tracking accuracy on extremely agile trajectories with peak velocities of 19.4 m s^{-1} without the need to update controller gains or model parameters (right).

- [15] T. P. Nascimento and M. Saska, "Position and attitude control of multi-rotor aerial vehicles: A survey," *Annual Reviews in Control*, 2019.
- [16] S. Sun, A. Romero, P. Foehn, E. Kaufmann, and D. Scaramuzza, "A comparative study of nonlinear mpc and differential-flatness-based control for quadrotor agile flight," *arXiv preprint arXiv:2109.01365*, 2021.
- [17] A. Romero, S. Sun, P. Foehn, and D. Scaramuzza, "Model predictive contouring control for near-time-optimal quadrotor flight," *arXiv preprint arXiv:2108.13205*, 2021.
- [18] I. M. Mareels, B. D. Anderson, R. R. Bitmead, M. Bodson, and S. S. Sastry, "Revisiting the mit rule for adaptive control," in *Adaptive Systems in Control and Signal Processing 1986*, ser. IFAC Workshop Series. Pergamon, 1987.
- [19] E. Lavretsky and K. A. Wise, "Robust adaptive control," in *Robust and adaptive control*. Springer, 2013, pp. 1 – 449.
- [20] K. J. Åström and B. Wittenmark, *Adaptive control*. Courier Corporation, 2013.
- [21] P. A. Ioannou and J. Sun, *Robust adaptive control*. Courier Corporation, 2012.
- [22] N. Hovakimyan and C. Cao, *L1 adaptive control theory: Guaranteed robustness with fast adaptation*. SIAM, 2010.
- [23] C. Cao and N. Hovakimyan, "Design and analysis of a novel l1 adaptive control architecture with guaranteed transient performance," *IEEE Transactions on Automatic Control*, vol. 53, no. 2, pp. 586–591, 2008.
- [24] R. Beard, C. Cao, and N. Hovakimyan, "An l1 adaptive pitch controller for miniature air vehicles," in *AIAA guidance, navigation, and control conference and exhibit*, 2006, p. 6777.
- [25] I. Gregory, C. Cao, E. Xargay, N. Hovakimyan, and X. Zou, "L1 adaptive control design for nasa airstar flight test vehicle," in *AIAA guidance, navigation, and control conference*, 2009, p. 5738.
- [26] S. Mallikarjunan, B. Nesbitt, E. Kharisov, E. Xargay, N. Hovakimyan, and C. Cao, "L1 adaptive controller for attitude control of multirotors," in *AIAA guidance, navigation, and control conference*, 2012, p. 4831.
- [27] Z. T. Dydek, A. M. Annaswamy, and E. Lavretsky, "Adaptive control of quadrotor uavs: A design trade study with flight evaluations," *IEEE Trans. Control Sys. Tech.*, vol. 21, no. 4, pp. 1400–1406, 2013.
- [28] C. Nicol, C. Macnab, and A. Ramirez-Serrano, "Robust adaptive control of a quadrotor helicopter," *Mechatronics*, vol. 21, no. 6, pp. 927–938, 2011.
- [29] M. Schreier, "Modeling and adaptive control of a quadrotor," in *2012 IEEE International Conference on Mechatronics and Automation*, 2012, pp. 383–390.
- [30] P. Kotaru, R. Edmonson, and K. Sreenath, "Geometric L1 Adaptive Attitude Control for a Quadrotor Unmanned Aerial Vehicle," *Journal of Dynamic Systems, Measurement, and Control*, vol. 142, no. 3, 12 2019, 031003. [Online]. Available: <https://doi.org/10.1115/1.4045558>
- [31] M. Bujarbaruah, X. Zhang, H. E. Tseng, and F. Borrelli, "Adaptive mpc for autonomous lane keeping," *International Symposium on Advanced Vehicle Control (AVEC)*, 2018.
- [32] K. Pereida and A. P. Schoellig, "Adaptive model predictive control for high-accuracy trajectory tracking in changing conditions," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2018, pp. 7831–7837.
- [33] M. Kamel, M. Burri, and R. Siegwart, "Linear vs nonlinear mpc for trajectory tracking applied to rotary wing micro aerial vehicles," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 3463–3469, 2017, 20th IFAC World Congress. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2405896317313083>
- [34] J. Pravitra, K. A. Ackerman, C. Cao, N. Hovakimyan, and E. A. Theodorou, "L1-adaptive mppi architecture for robust and agile control of multirotors," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 7661–7666.
- [35] 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.
- [36] 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.
- [37] T. Lee, M. Leok, and N. H. McClamroch, "Geometric tracking control of a quadrotor uav for extreme maneuverability," *IFAC Proceedings Volumes*, 2011.
- [38] 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.
- [39] B. Houska, H. Ferreau, and M. Diehl, "ACADO Toolkit – An Open Source Framework for Automatic Control and Dynamic Optimization," *Optimal Control Applications and Methods*, vol. 32, no. 3, pp. 298–312, 2011.
- [40] X. Wang and N. Hovakimyan, "L1 adaptive controller for nonlinear time-varying reference systems," *Systems & Control Letters*, vol. 61, no. 4, pp. 455–463, 2012.
- [41] Z. Li and N. Hovakimyan, "L1 adaptive controller for mimo systems with unmatched uncertainties using modified piecewise constant adaptation law," in *2012 IEEE 51st IEEE conference on decision and control (CDC)*. IEEE, 2012, pp. 7303–7308.
- [42] E. Xargay, N. Hovakimyan, and C. Cao, "L1 adaptive controller for multi-input multi-output systems in the presence of nonlinear unmatched uncertainties," in *Proceedings of the 2010 American Control Conference*, 2010, pp. 874–879.
- [43] 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.
- [44] 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.
- [45] G. Williams, A. Aldrich, and E. A. Theodorou, "Model predictive path integral control: From theory to parallel computation," *Journal of Guidance, Control, and Dynamics*, vol. 40, no. 2, pp. 344–357, 2017.
- [46] 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.