

A Simulation Evaluation Suite for Robust Adaptive Quadcopter Control

Dingqi Zhang¹, Ran Tao², Sheng Cheng², Naira Hovakimyan², and Mark W. Mueller¹.

Abstract—Robust adaptive control methods are essential for maintaining quadcopter performance under external disturbances and model uncertainties. However, fragmented evaluations across tasks, simulators, and implementations hinder systematic comparison of these methods. This paper introduces an easy-to-deploy, modular simulation testbed for quadcopter control, built on *RotorPy*, that enables evaluation under a wide range of disturbances such as wind, payload shifts, rotor faults, and control latency. The framework includes a library of representative adaptive and non-adaptive controllers and provides task-relevant metrics to assess tracking accuracy and robustness. The unified modular environment enables reproducible evaluation across control methods and eliminates redundant reimplementations of components such as disturbance models, trajectory generators, and analysis tools. We illustrate the testbed’s versatility through examples spanning multiple disturbance scenarios and trajectory types, including automated stress testing, to demonstrate its utility for systematic analysis. Code is available at <https://github.com/Dz298/AdaptiveQuadBench>.

I. INTRODUCTION

Tracking trajectories in the presence of external disturbances and model uncertainties is a fundamental challenge in quadcopter control. Disturbances such as wind gusts, payload variations, rotor degradation, and communication latency are common in real-world settings, especially in applications like aerial logistics and infrastructure inspection. While accurate modeling can improve control performance, capturing all relevant dynamics, including time-varying forces and unmodeled aerodynamics, is often infeasible in practice. As a result, robust and adaptive control strategies have been widely explored to ensure reliable performance under uncertainty [1].

A range of control approaches have been developed to address these challenges. Classical methods such as PID [2], geometric control on $SE(3)$ [3], and model predictive control (MPC) [4], [5] offer strong performance under nominal conditions but can degrade in the presence of model mismatch. To improve robustness, techniques such as \mathcal{L}_1 adaptive control [6], adaptive geometric control [7], [8], and adaptive incremental nonlinear dynamic inversion (INDI) [9] have been proposed to compensate for uncertainty in mass, inertia, aerodynamic drag, and external forces. In parallel, learning-based controllers trained in simulation using reinforcement learning or imitation learning have shown promise in transferring to real-world flight through techniques such as

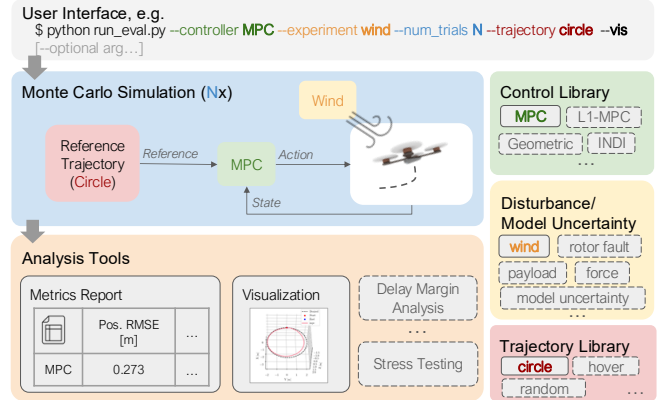


Fig. 1. Overview of our simulation testbed. The framework consists of a dynamics simulator based on *RotorPy* [21], a modular control interface supporting both adaptive and non-adaptive controllers, and evaluation tools for trajectory generation, disturbance injection, performance metrics, and visualization. Users interact with the system through a unified configuration interface that specifies controller selection, disturbance type, trajectory, and evaluation mode.

domain randomization [10], [11] and residual learning [12], [13].

To validate these robust adaptive control methods, simulation plays a critical role. It enables Monte Carlo analysis across randomized scenarios, and scalable experiments beyond what is feasible on hardware. Despite significant progress, simulation-based evaluation remains fragmented. For instance, [14] assesses MPC augmented by \mathcal{L}_1 adaptive control under model mismatch conditions using *RotorS* [15], an open-source simulator based on Gazebo and ROS. Similarly, [16] compares nonlinear MPC with differential-flatness-based control under model uncertainties using the simulation in *Agilicious* [17], a full-stack framework designed for agile flight. While these studies provide valuable insights, the need to repeatedly develop simulation environments for similar tasks can lead to overlapping efforts. In addition, existing simulators are frequently tied to specific application such as hardware integration ([17], [15]), reinforcement learning training ([18], [19]) or agricultural inspection [20]. As a result, evaluation tools are typically embedded within broader simulation frameworks, making it challenging for researchers to decouple them from the rest of the software stack. This separation between evaluation tools and simulation environments demands a high level of expertise, hindering accessibility for both newcomers and experienced researchers aiming for benchmarking. The closest work in this direction is *Safe Control Gym* [22], which provides a standardized evaluation environment for

¹The authors are with the High Performance Robotics Lab, Dept. of Mechanical Engineering, UC Berkeley. Contact at {dingqi, mwm}@berkeley.edu

²The authors are with Advanced Controls Research Laboratory, Dept. of Mechanical Engineering, University of Illinois Urbana-Champaign, Contact at {rant3, chengs, nhovakim}@illinois.edu

safe and robust control. However, its underlying dynamics are simplified (e.g., a 2D quadcopter model), limiting its ability to capture dynamics such as the aerodynamic and inertial effects crucial for realistic evaluation.

To address these challenges, we introduce a modular simulation framework that enables reproducible and systematic benchmarking of quadcopter controllers. Built on top of the lightweight and open-source simulator *RotorPy* [21], the testbed supports the evaluation of controllers across diverse tasks and includes a modular control library featuring representative robust adaptive control methods for quadcopters.

In summary, this paper provides:

- 1) A simulation benchmarking framework for quadcopter control that supports systematic, reproducible evaluation under disturbances and model uncertainties.
- 2) An open-source control library containing representative robust adaptive control methods for quadcopters.
- 3) A configurable suite of disturbance and uncertainty models, including wind, payload shifts, rotor fault, and control latency.
- 4) A set of evaluation tools for performance and robustness analysis.

The framework is designed to support reproducible, structured benchmarking of quadcopter control strategies under diverse and challenging conditions, without favoring a specific controller or architecture.

II. FRAMEWORK OVERVIEW

In this section, we detail the setup of our simulation testbed, covering quadcopter dynamics, the control pipeline, evaluation tasks, and associated performance metrics. Additionally, we describe the implementation of external disturbances and model uncertainties introduced into the plant. An overview of the system is provided in Figure 1. The testbed is designed to support consistent evaluation of diverse control methods, enabling users to conduct their own comparisons across disturbances and uncertainties.

A. Quadcopter Dynamics

While our testbed builds on *RotorPy*'s base dynamics, we summarize the model here to clarify the modifications introduced for adaptive control evaluation. In particular, we extend *RotorPy* by: (i) introducing external force and torque to the dynamics; (ii) incorporating a motor effectiveness vector to simulate rotor-level asymmetries; and (iii) adopting a planar drag formulation that has been shown effective on hardware experiments in [16], [10].

The translational dynamics are modeled in a world-fixed inertia frame, given by:

$$\dot{\mathbf{x}} = \mathbf{v}, \quad (1)$$

$$\dot{\mathbf{v}} = \frac{1}{m}(T\mathbf{z}_B + \mathbf{f}_a + \mathbf{f}_{\text{ext}}) + \mathbf{g}, \quad (2)$$

where \mathbf{x} and \mathbf{v} are the position and velocity vectors; T and m are the collective thrust and total mass, respectively; \mathbf{g} is the gravitational vector; \mathbf{f}_a indicates the aerodynamic drag force during flights, and \mathbf{f}_{ext} lumps the external forces as




Fig. 2. Illustration of the quadcopter's x -configuration motor layout in the body frame. The directions of rotation for each rotor (ω_1 to ω_4) and the geometric angle β between the rotor arms and the body axis x_B are indicated.

disturbances; the vector $\mathbf{z}_B \in \mathbb{R}^3$ is the unit vector along the body-frame z -axis.

The rotational dynamics are expressed as:

$$\dot{\mathbf{R}} = \mathbf{R}[\boldsymbol{\Omega}]_{\times}, \quad (3)$$

$$\dot{\boldsymbol{\Omega}} = \mathbf{J}^{-1}(-[\boldsymbol{\Omega}]_{\times}\mathbf{J}\boldsymbol{\Omega} + \boldsymbol{\tau} + \boldsymbol{\tau}_{\text{ext}} + \boldsymbol{\tau}_a), \quad (4)$$

where $\mathbf{R} \in \text{SO}(3)$ is the rotation matrix from the body frame to the inertial frame, and $\boldsymbol{\Omega}$ is the angular velocity expressed in the body frame. The operator $[\cdot]_{\times}$ denotes the skew-symmetric matrix associated with the cross product. We denote \mathbf{J} as the inertia matrix of the quadcopter; $\boldsymbol{\tau}$ is the torque generated by the rotors; $\boldsymbol{\tau}_a$ and $\boldsymbol{\tau}_{\text{ext}}$ are the aerodynamic torque and the disturbance torques, respectively.

The collective thrust T and torques $\boldsymbol{\tau}$ are functions of rotor speeds:

$$\begin{bmatrix} T \\ \boldsymbol{\tau} \end{bmatrix} = \mathbf{G}\mathbf{u}, \quad (5)$$

$$u_i = k_t \omega_i^2, \quad i \in \{1, 2, 3, 4\}, \quad (6)$$

where $\mathbf{u} \in \mathbb{R}^4$ is the vector of individual rotor thrusts, ω_i is the speed of rotor i , and k_t is the thrust coefficient. \mathbf{G} is a matrix defined as:

$$\mathbf{G} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ l \sin \beta & -l \sin \beta & -l \sin \beta & l \sin \beta \\ -l \cos \beta & -l \cos \beta & l \cos \beta & l \cos \beta \\ k_{\tau}/k_t & -k_{\tau}/k_t & k_{\tau}/k_t & -k_{\tau}/k_t \end{bmatrix}, \quad (7)$$

where l denotes the distance from the vehicle center to each rotor (arm length), and β is the angle between the x_B axis and the rotor arm direction, as illustrated in Figure 2; k_{τ} is the torque coefficient.

Quadcopter motors cannot instantaneously achieve commanded speeds due to rotor inertia and motor dynamics. Their transient behavior is commonly approximated using a first-order model, as adopted in prior work such as [23], [24]:

$$\dot{\omega} = \frac{1}{\tau_m}(\omega_{\text{des}} - \omega), \quad (8)$$

where ω_{des} is the desired motor speed; τ_m is the motor time constant which can be identified using static thrust stand testing.

The aerodynamic forces \mathbf{f}_a and torques $\boldsymbol{\tau}_a$ together form the aerodynamic wrench, representing the collection of forces and moments generated by the quadcopter’s motion through the air. These effects become particularly significant in high-speed flights or strong wind conditions. While *RotorPy* includes a parasite drag model identified on the Crazyflie, we adopt a planar drag formulation that has been validated across multiple quadcopter platforms and shown to transfer reliably to hardware [16], [10]. This choice favors generalizability for benchmarking, while we still retain *RotorPy*’s rotor drag and blade flapping terms. The planar drag model is expressed as:

$$\mathbf{f}_{a,p} = \begin{bmatrix} -k_{d,x}v_{a,x} \\ -k_{d,y}v_{a,y} \\ -k_{d,z}v_{a,z} + k_h(v_{a,x}^2 + v_{a,y}^2) \end{bmatrix}, \quad (9)$$

where $[v_{a,x}, v_{a,y}, v_{a,z}] = \mathbf{R}^T(\mathbf{v} - \mathbf{v}_w)$ is the relative body airspeed with respect to the wind velocity \mathbf{v}_w ; $k_{d,[x,y,z]}$ and k_h are aerodynamic parameters configured within the simulator to reflect nominal drag and lift effects.

In addition, we adopt the rotor drag and blade flapping models from *RotorPy*, which account for aerodynamic effects due to rotor-induced flow and blade motion (see Eq. (8-9) in [21]).

B. Control Library

We provide a library of adaptive control methods to facilitate rapid prototyping and benchmarking for researchers. Additionally, we include non-adaptive controllers as baselines to quantify performance improvements in the presence of external disturbances and model uncertainties.

The baseline controllers include:

- `geo`: a geometric controller [3];
- `mpc`: a nonlinear MPC [16].

The adaptive controllers include:

- `geo-a`: a geometric adaptive controller [8];
- `llgeo`: a geometric controller augmented with \mathcal{L}_1 adaptive control [7];
- `llmpc`: a nonlinear MPC augmented with \mathcal{L}_1 adaptive control [25];
- `indi-a`: an incremental nonlinear dynamic inversion controller [26], which incorporates onboard adaptive parameter estimation to update control effectiveness [9];
- `xadap`: a learning-based adaptive controller [10].

These controllers operate at different command levels: some, such as `geo-a` and `llgeo`, generate control commands at the collective thrust and moment level, while others, such as `xadap` and `indi-a`, directly output motor speed commands. To standardize execution across different control levels, we compute the required motor speed for those not commanding it using (5) and (6). Note the model parameters such as k_t and k_τ used in these computations are not the ground-truth values but reference model parameters that may contain uncertainties. We discuss these uncertainties in detail in Section II-E.

In addition, among these methods, `xadap` and `indi-a` are proposed as inner-loop controllers, which take intermediate control inputs such as commanded collective thrust and angular velocity from an outer-loop controller. To complete the control pipeline, we concatenate it with a geometric controller as the outer-loop controller.

C. Evaluation Task and Metrics

We evaluate tracking performance along a desired trajectory as the primary task in our testbed. For large-scale batch evaluations, we generate trajectories using randomized motion primitives following the method of [27]. Specifically, each trajectory is composed of a sequence of short-duration polynomial segments, where the motion parameters such as velocity, acceleration, and segment duration are randomly sampled within dynamically feasible bounds. This approach ensures physical feasibility, supports rapid generation, and produces a diverse set of trajectories that challenge different aspects of controller performance. However, users can customize the testbed by specifying alternative trajectories tailored to their specific tasks, such as hovering or circular flight. To quantify tracking accuracy, we use the root mean square error (RMSE) of position and the absolute error of heading as the primary precision metrics. Additionally, we assess robustness using the success rate, where a flight is considered failed if its position \mathbf{x} violates the following spatial constraint at an arbitrary instant:

$$\|\mathbf{x}_{\text{des}} - \mathbf{x}\|_2 \leq e_{x,\text{max}}. \quad (10)$$

We select $e_{x,\text{max}} = 5\text{m}$ in our testbed.

D. External Disturbances

To ensure a comprehensive evaluation, we implement a variety of disturbances commonly tested in previous studies such as [14], [10], [16].

1) *Gusty Wind*: Gusty wind varies rapidly, making it an effective test case for evaluating fast adaptation. To model this effect, we use the Dryden wind turbulence model [28], which characterizes turbulence based on the mean wind speed along each axis and its variance, defining the turbulence intensity. The resulting disturbance acts as an aerodynamic wrench, as defined in II-A, impacting both forces and moments on the quadcopter.

2) *Off-center Payload*: Off-center payloads are common in aerial logistics and introduce coupled effects on mass, center of mass (COM), and inertia. We model the payload as a rigidly attached point mass and account for its impact as follows:

- 1) The total mass is updated:

$$m_{\text{total}} = m_{\text{vehicle}} + m_{\text{payload}}, \quad (11)$$

- 2) The center of mass (COM) is shifted to:

$$\mathbf{x}_{\text{COM}} = \frac{m_{\text{vehicle}} \mathbf{x}_{\text{vehicle}} + m_{\text{payload}} \mathbf{x}_{\text{payload}}}{m_{\text{total}}}, \quad (12)$$

where $\mathbf{x}_{\text{vehicle}}$ is the nominal COM and the origin of the body frame defined in Figure 2. The shifted COM

changes the rotor moment arms, so we also update the control allocation matrix \mathbf{G} (defined in (7)) to reflect the new geometry.

- 3) The inertia tensor is recomputed using the parallel axis theorem [29], accounting for the shifted mass and added point mass.

The gravitational force on the payload is naturally captured by the increased mass. However, the offset between the payload and the new COM introduces a torque due to gravity, computed as:

$$\boldsymbol{\tau}_{\text{ext}} = [\mathbf{x}_{\text{payload}} - \mathbf{x}_{\text{COM}}]_{\times} (m_{\text{payload}} \mathbf{g}). \quad (13)$$

This torque is explicitly included as an external disturbance. To introduce variability, both the payload mass m_{payload} and its attachment location $\mathbf{x}_{\text{payload}}$ are randomized across trials. To simulate pickup and drop-off events, the payload is attached or removed at a random time during flight.

3) *External Forces and Torques:* External force and torque disturbances are implemented as independent perturbations, applied separately (unlike inherently coupled forces and moments in the off-center payload case), allowing us to analyze their individual effects on the system. Moreover, in contrast to the gusty wind case, these disturbances remain constant over time, enabling a controlled evaluation of their isolated impact without temporal variations.

4) *Loss of Rotor Effectiveness:* In real-world conditions, rotor effectiveness varies due to hardware differences, wear, or battery depletion. To model this effect, we introduce a positive effectiveness factor k_{eff} , which scales the contribution of each motor and modifies the system dynamics as follows:

$$k_{\text{eff},i} = 1 + \mathcal{U}(-\delta, \delta), \quad (14)$$

$$u_i = k_{\text{eff},i} k_t \omega_i^2, \quad i \in \{1, 2, 3, 4\}, \quad (15)$$

where $\mathcal{U}(-\delta, \delta)$ is a uniform distribution over $[-\delta, \delta]$; $\delta \in [0, 1)$ defines the variation range of this motor's effectiveness.

It is important to note that setting $k_{\text{eff},i}$ to zero does not fully capture the dynamics of a complete rotor failure, as explored in [30]. In the event of motor failure, the vehicle loses yaw authority and experiences asymmetric drag. While these effects can be modeled with moderate extensions, such as including rotational drag terms, we do not implement them in this work.

5) *Latency:* Latency refers to time delays between sensor measurements, controller decisions, and actuator execution. In quadcopters, such delays are typically due to communication overhead, processing constraints, or software pipeline structure. When present in feedback loops, latency can act as a form of disturbance that degrades control performance and may lead to instability.

To evaluate robustness to latency, we introduce a fixed control execution delay. Control inputs are buffered and applied after a specified delay, while sensor updates remain undelayed. This setup isolates the effect of actuation-side delay, resembling latency between a state estimator and controller or between ground station and vehicle. The robustness

of a controller to such delay can be characterized by its delay margin, defined as the maximum allowable latency before the system becomes unstable. While frequency-domain tools such as Bode or Nyquist plots can estimate delay margins for linear systems [31], they are less effective for nonlinear systems like quadcopters.

Instead, we define a delay margin as the maximum control delay under which the controller can complete a trajectory with the position error contained within a user-defined threshold at any time:

$$\|\mathbf{x}_{\text{des}} - \mathbf{x}\|_2 \leq e_{\text{th}}, \quad (16)$$

where e_{th} is a user-defined position error threshold. This value can be tuned depending on the scale and precision requirements of the application; for example, smaller thresholds (e.g., 0.5m) for micro aerial vehicles and larger ones (e.g., 2–5m) for full-scale platforms. We estimate this margin by incrementally increasing the control delay and recording the point at which the system violates the threshold across different trajectory types.

E. Model Uncertainties

Beyond external disturbances, a quadcopter's performance is influenced by uncertainties in the reference model, such as parameter offsets or errors. We consider two types of model uncertainty: *Uniform* and *Scaled*.

1) *Uniform Uncertainties:* In this type of uncertainties, each parameter is perturbed independently around its nominal value using a uniform distribution within a fixed range, which models unbiased parameter errors.

2) *Scaled Uncertainties:* When seeking a significantly different reference model, uniform perturbations may result in physically inconsistent configurations. For example, sampling mass and inertia independently from a uniform distribution over $\pm 50\%$ may yield combinations where the mass is reduced while the inertia is increased, violating physical relationships between size, mass distribution, and dynamics. To address this issue, we introduce a scaling factor c that proportionally adjusts the arm length:

$$c = \mathcal{U}(-c_{\text{max}}, c_{\text{max}}), \quad c_{\text{max}} \in [0, 1) \quad (17)$$

$$l_{\text{ref}} = (1 + c) l_{\text{true}}, \quad (18)$$

where l_{ref} is the arm length of the reference model; l_{true} is the ground truth value of the arm length. The range of the scaling factor c guarantees that the resulting arm length remains strictly positive.

Assuming constant density and isotropic scaling, we scale mass with the cube of the arm length, inertia with the fifth power, and drag with the square. To reflect the relationship between quadcopter size and motor strength, we choose to exponentially scale the thrust coefficient with c , ensuring that larger quadcopters have proportionally stronger motors. Other parameters, including maximum motor speed and propeller constants, are scaled linearly. This structured scaling preserves physical consistency, preventing unrealistic configurations such as small drones with excessively powerful motors. The single scaling factor c quantifies the deviation

TABLE I
QUADROTOR PARAMETERS

Parameter		Value
m	[kg]	0.75
l	[m]	0.166
β	[deg]	45
\mathbf{J}	[kg·m ²]	diag(0.0047, 0.005, 0.0074)
$(k_{d,x}, k_{d,y}, k_{d,z})$	[kg/s]	(0.62, 0.62, 0.62)
k_h	[kg/m]	0.0
k_t	[N/(rad/s) ²]	$7.64e^{-6}$
k_τ	[N·m/(rad/s) ²]	$1.07e^{-7}$
τ_m	[s]	0.01
ω_{\max}	[rad/s]	1000

of the reference model from the ground truth, with $\max |c|$ representing the maximum allowable deviation.

Finally, to account for deviations from idealized scaling laws, we introduce uniformly distributed noise in the range [-20%, 20%] across all parameters after applying the scaling factor c . This enhances model diversity while maintaining plausibility.

III. EVALUATION EXAMPLES

In this section, we present examples of evaluations conducted using our testbed, which supports automated generation of statistical tables and plots. To demonstrate the framework’s versatility, we show a subset of controllers for each disturbance type. This selection highlights the range of controllers supported by the testbed while ensuring that each example remains focused and interpretable. For all results presented in this section, we adopt the *large quadrotor* model from [10], using the parameters specified in Table I.

A. Uncertainty-free Condition

First, we present the performance of each controller in the absence of disturbances and model uncertainties. We fly the vehicle using a reference model that best represents our knowledge of the system dynamics. In this case, model parameters such as mass, inertia, and rotor arm lengths are assumed to be perfectly known, which means that the parameters used in model-based controllers match exactly with those in the simulation. These experiments help establish the best possible performance each method can achieve under the uncertainty/disturbance-free conditions. The only exception is the learning-based controller `xadap`, which does not rely on a reference model due to its design.

```
$ python run_eval.py --controller all --
  experiment no
```

This command runs the evaluation task using all available controllers under this uncertainty-free conditions and reports task statistics, such as position RMSE, heading error, and success rate, based on the metrics defined in the previous section. Table II shows baseline performance under idealized conditions. All controllers, except the learning-based `xadap`, depend on parameter tuning and can achieve comparable tracking accuracy and success rates in the absence of disturbances and model uncertainties.

TABLE II

PERFORMANCE OF EACH CONTROLLER IN THE ABSENCE OF DISTURBANCES AND MODEL UNCERTAINTIES. WITH APPROPRIATE TUNING, ALL CONTROLLERS CAN ACHIEVE SIMILAR PERFORMANCE.

Method	Position RMSE [m]	Heading Error [deg]	Success Rate
geo	0.074 ± 0.016	0.523 ± 0.137	100%
geo-a	0.065 ± 0.016	0.565 ± 0.148	100%
llgeo	0.048 ± 0.015	0.223 ± 0.156	100%
indi-a	0.030 ± 0.007	0.254 ± 0.157	100%
llmpc	0.035 ± 0.012	0.224 ± 0.123	100%
mpc	0.040 ± 0.012	0.890 ± 0.285	100%
xadap	0.046 ± 0.009	14.044 ± 0.648	100%




Fig. 3. Position tracking of `llgeo` and `geo-a` along a circular trajectory in the presence of gusty wind, generated using the Dryden turbulence model.

B. Gusty Wind

We now demonstrate how the framework can be used to evaluate controller performance under wind disturbances. In this example, we introduce a gusty wind field modeled using the Dryden turbulence model, which varies continuously in time and direction. Rather than comparing methods via aggregated tables, we focus on visualizing performance along multiple dimensions, such as position error, control effort, and system response, on a single trajectory. To visualize the performance of selected methods on a single rollout, we can enable the `--vis` flag in the evaluation script:

```
$ python run_eval.py --controller llgeo geo-a
  --experiment wind --trajectory circle --
  vis
```

This example, along with Figure 3 generated, shows the position tracking performance of the `llgeo` and `geo-a` controllers along a circular trajectory under gusty wind conditions. Similarly, external force and torque disturbances can be evaluated by modifying the command line argument accordingly, where the disturbance plot will update to reflect the corresponding disturbance type.

C. Off-center Payload

```
$ python run_eval.py --controller geo llgeo
  xadap --vis --experiment payload --
  trajectory circle
```

This command generates Figure 4, which visualizes the position tracking performance of `geo`, `llgeo`, and

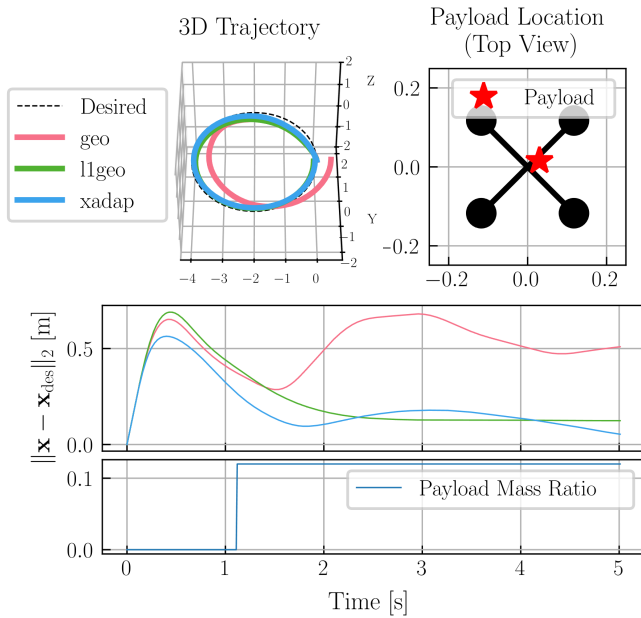


Fig. 4. Position tracking of `geo`, `llgeo`, and `xadap` under a circular trajectory with an off-center payload randomly attached mid-flight. The payload introduces coupled effects on center-of-mass, inertia, and external torque. The deviation from the desired trajectory before payload attachment arises from the transient controller response as the vehicle transitions from hovering to trajectory tracking.

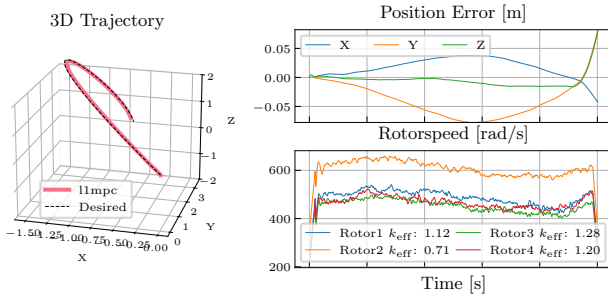


Fig. 5. Position tracking and motor speed response of the `llmpc` controller on a randomized trajectory under rotor effectiveness variations. Each rotor’s effectiveness factor is perturbed, simulating actuator-level faults.

`indi-a`. The figure also includes the payload’s location and the ratio of the payload mass to the vehicle’s mass over time.

D. Loss of Rotor Effectiveness

Figure 5 illustrates the motor speed response and tracking performance of the `llmpc` controller when following a randomized motion primitive trajectory under rotor effectiveness variations. Additionally, the perturbed effectiveness coefficient for each rotor is plotted, providing insight into how motor effectiveness influences system behavior. This visualization is generated by running:

```
$ python run_eval.py --controller llmpc --vis
--experiment rotoreff
```

E. Latency

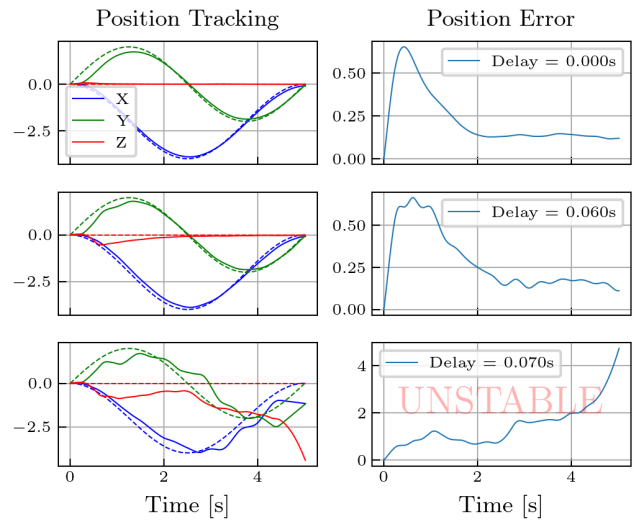


Fig. 6. Position tracking of the `geo` controller on a circular trajectory under increasing control latency. Three cases are shown: no delay, the maximum stable delay, and the first unstable delay instance.

```
$ python run_eval.py --controller geo --
delay_margin
```

This command generates Figure 6, which illustrates how control latency affects tracking performance under a circular trajectory using the `geo` controller. The plots show position tracking and position error under three latency settings: 0.00s (no delay), 0.06s (the final delay before instability), and 0.07s (the smallest delay at which instability occurs). The observed failure point is specific to the chosen controller, trajectory, and error threshold. Different controllers or trajectories may tolerate different amounts of latency before instability emerges. In addition, the delay margin is defined relative to a user-specified error threshold (Section II-D.5); varying this threshold will also affect the measured margin.

F. Model Uncertainty

```
$ python run_eval.py --controller xadap --vis
--experiment uncertainty
```

This example (Figure 7) visualizes both the tracking performance of `xadap` and the impact of model uncertainties. The model uncertainties are visualized by a polar chart illustrating the discrepancy between the reference model parameters used by the controller and the actual ground-truth parameters, providing insight into the magnitude of modeling errors across key parameters.

IV. EXTREME ADAPTATION EVALUATION

The previous section demonstrated how our testbed can simulate a variety of disturbance types and model uncertainties, enabling qualitative inspection of controller behavior under each condition. In this section, we shift focus toward evaluating the limits of controller robustness by progressively increasing the intensity of a disturbance to identify

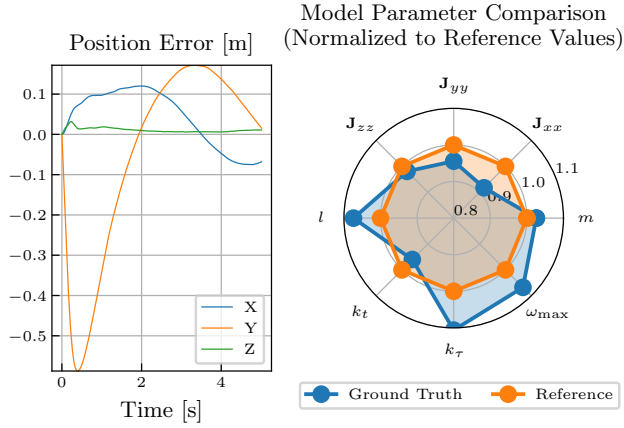


Fig. 7. Left: Position tracking of `xadap` under structured model parameter perturbations. Right: Polar chart comparing reference model parameters to ground truth, normalized to nominal values.

failure points. This type of evaluation has been explored in previous works [10], [14], [32], where the disturbance magnitude or uncertainty range is systematically increased (e.g., increasing the maximum trajectory speed or expanding the model uncertainty bounds) to determine the threshold at which performance degrades below an acceptable level. Such evaluations serve as valuable tools for assessing robustness and improving controller design. Additionally, they are particularly useful for analyzing the out-of-distribution performance of learning-based methods.

To facilitate this analysis, we implement an automatic stress-testing feature in our testbed. This feature progressively increases the range of randomization until either the success rate of the selected method drops to a selected threshold such as 10% or the maximum disturbance intensity is reached. The type of disturbance increment depends on the specific experiment setting:

- **Wind:** Increase in gust magnitude
- **Payload:** Increase in mass (but not arm length, as arbitrarily large offsets are unrealistic)
- **External Force:** Increase in force magnitude
- **External Torque:** Increase in torque magnitude
- **Loss of Rotor Effectiveness:** Increase in the maximum deviation of the motor effectiveness vector from a nominal, identical-motor setup, as quantified by δ in (14).
- **Model Uncertainty:** Increase in the maximum deviation from the ground-truth model, quantified by $\max |c|$ in Section II-E.2.

For example, we can evaluate the performance of `geo`, `mpc`, and `l1mpc` over 100 trials with gusty wind speed increasing from 0 to 5 m/s in increments of 1 m/s using the following command:

```
$ python run_eval.py --controller geo mpc
  l1mpc --experiment wind --when2fail --
  max_intensity 5 --intensity_step 1
```

The results, shown in Figure 8, indicate that while all methods maintain a 100% success rate, position tracking

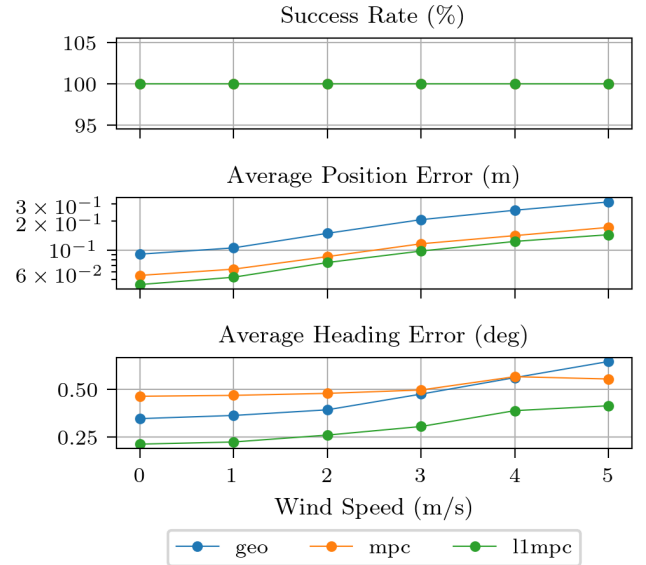


Fig. 8. Tracking error and heading error as a function of wind speed for `geo`, `mpc`, and `l1mpc`, under increasing gust intensity using automated stress testing.

error and heading error increase as wind speed intensifies.

Another example (Figure 9) examines how increasing the payload mass ratio impacts controller performance. The following command runs the evaluation:

```
$ python run_eval.py --controller geo indi-a
  xadap --experiment payload --when2fail --
  max_intensity 2 --intensity_step 0.5
```

This command increases the maximum possible payload mass ratio (relative to the nominal vehicle mass) from 0% to 200% in increments of 50%. The results show that while the non-adaptive `geo` controller degrades significantly as payload mass increases, both `indi-a` and `xadap` maintain high success rates and better tracking performance under high payload disturbances.

V. CONCLUSION

This paper introduces a simulation framework for benchmarking quadcopter controllers under external disturbances and dynamic uncertainties, including wind, payload, rotor faults, and latency. Built on *RotorPy*, the framework enables reproducible testing across different robust adaptive controllers and provides a set of evaluation tools for performance and robustness analysis. It supports repeatable and extensible experimentation, contributing to ongoing efforts in reproducible robotics research.

The framework is designed to be extensible with additional controllers and disturbance types. Future work may include incorporating more diverse wind models, simulating complete motor failures, and supporting additional real-world uncertainties. These extensions would further enhance the framework's ability to evaluate adaptive controllers under challenging and realistic conditions.

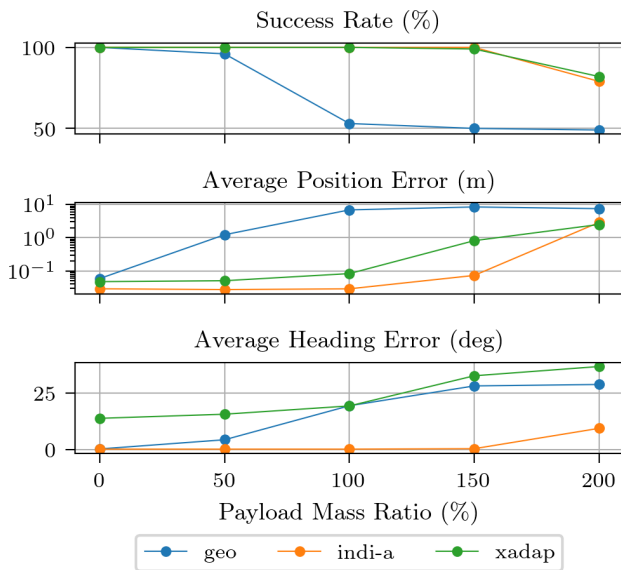


Fig. 9. Tracking performance of `geo`, `indi-a`, and `xadap` under increasing payload mass ratios using automated stress testing.

VI. ACKNOWLEDGEMENT

This work was supported by the Hong Kong Center for Logistics Robotics, NASA Cooperative Agreement (80NSSC20M0229), NASA ULI (80NSSC22M0070), and AFOSR (FA955021-1-0411). The experimental testbed at the HiPeRLab is the result of contributions of many people, a full list of which can be found at hiperlab.berkeley.edu/members/.

REFERENCES

- [1] A. Saviolo and G. Loianno, "Learning quadrotor dynamics for precise, safe, and agile flight control," *Annual Reviews in Control*, vol. 55, pp. 45–60, 2023.
- [2] M. A. Johnson and M. H. Moradi, *PID control*. Springer, 2005.
- [3] T. Lee, M. Leok, and N. Mcclamroch, "Geometric tracking control of a quadrotor uav on $se(3)$," 01 2011, pp. 5420 – 5425.
- [4] E. F. Camacho and C. B. Alba, *Model predictive control*. Springer Science & Business Media, 2013.
- [5] J. B. Rawlings, D. Q. Mayne, and M. M. Diehl, *Model Predictive Control: Theory, Computation, and Design, 2nd Ed.* Nob Hill Publishing, 2020.
- [6] N. Hovakimyan and C. Cao, *L1 adaptive control theory: Guaranteed robustness with fast adaptation*. SIAM, 2010.
- [7] Z. Wu, S. Cheng, P. Zhao, A. Gahlawat, K. A. Ackerman, A. Lakshmanan, C. Yang, J. Yu, and N. Hovakimyan, "L1quad: L1 adaptive augmentation of geometric control for agile quadrotors with performance guarantees," *IEEE Transactions on Control Systems Technology*, vol. 33, no. 2, pp. 597–612, 2025.
- [8] F. A. Goodarzi, D. Lee, and T. Lee, "Geometric adaptive tracking control of a quadrotor unmanned aerial vehicle on $se(3)$ for agile maneuvers," *Journal of Dynamic Systems, Measurement, and Control*, vol. 137, no. 9, Jun. 2015.
- [9] E. Smeur, Q. Chu, and G. Croon, "Adaptive incremental nonlinear dynamic inversion for attitude control of micro air vehicles," *Journal of Guidance, Control, and Dynamics*, vol. 39, pp. 1–12, 12 2015.
- [10] D. Zhang, A. Loquercio, J. Tang, T.-H. Wang, J. Malik, and M. W. Mueller, "A learning-based quadcopter controller with extreme adaptation," *IEEE Transactions on Robotics*, vol. 41, p. 3948–3964, 2025. [Online]. Available: <http://dx.doi.org/10.1109/TRO.2025.3577037>

- [11] E. Kaufmann, L. Bauersfeld, and D. Scaramuzza, "A benchmark comparison of learned control policies for agile quadrotor flight," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 10 504–10 510.
- [12] L. Bauersfeld*, E. Kaufmann*, P. Foehn, S. Sun, and D. Scaramuzza, "Neurobem: Hybrid aerodynamic quadrotor model," in *Robotics: Science and Systems XVII*, ser. RSS2021. Robotics: Science and Systems Foundation, Jul. 2021.
- [13] R. Zhang, D. Zhang, and M. W. Mueller, "Proxflly: Robust control for close proximity quadcopter flight via residual reinforcement learning," 2024. [Online]. Available: <https://arxiv.org/abs/2409.13193>
- [14] D. Hanover, P. Foehn, S. Sun, E. Kaufmann, and D. Scaramuzza, "Performance, precision, and payloads: Adaptive nonlinear mpc for quadrotors," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 690–697, 2022.
- [15] F. Furrer, M. Burri, M. Achtelik, and R. Siegwart, *RotorS—A Modular Gazebo MAV Simulator Framework*. Cham: Springer International Publishing, 2016, pp. 595–625.
- [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," 2024. [Online]. Available: <https://arxiv.org/abs/2109.01365>
- [17] P. Foehn, E. Kaufmann, A. Romero, R. Penicka, S. Sun, L. Bauersfeld, T. Laengle, G. Cioffi, Y. Song, A. Loquercio, and D. Scaramuzza, "Agilicious: Open-source and open-hardware agile quadrotor for vision-based flight," *AAAS Science Robotics*, 2022.
- [18] Y. Song, S. Naji, E. Kaufmann, A. Loquercio, and D. Scaramuzza, "Flightmare: A flexible quadrotor simulator," in *Conference on Robot Learning*, 2020.
- [19] J. Panerati, H. Zheng, S. Zhou, J. Xu, A. Prorok, and A. P. Schoellig, "Learning to fly—a gym environment with pybullet physics for reinforcement learning of multi-agent quadcopter control," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 7512–7519.
- [20] J. Zha, T. Yang, and M. W. Mueller, "Agri-fly: Simulator for uncrewed aerial vehicle flight in agricultural environments," *IEEE Access*, vol. 12, pp. 140 900–140 907, 2024.
- [21] S. Folk, J. Paulos, and V. Kumar, "RotorPy: A python-based multirotor simulator with aerodynamics for education and research," *arXiv preprint arXiv:2306.04485*, 2023.
- [22] Z. Yuan, A. W. Hall, S. Zhou, L. Brunke, M. Greeff, J. Panerati, and A. P. Schoellig, "safe-control-gym: a unified benchmark suite for safe learning-based control and reinforcement learning in robotics," 2022. [Online]. Available: <https://arxiv.org/abs/2109.06325>
- [23] J. Eschmann, D. Albani, and G. Loianno, "Data-driven system identification of quadrotors subject to motor delays," 2024. [Online]. Available: <https://arxiv.org/abs/2404.07837>
- [24] M. W. Mueller, *Dynamics and Control of Autonomous Flight*, 1st ed., ser. Mathematical Engineering. Springer Cham, 2025. [Online]. Available: <https://doi.org/10.1007/978-3-031-91871-1>
- [25] R. Tao, P. Zhao, I. Kolmanovsky, and N. Hovakimyan, "Robust adaptive mpc using uncertainty compensation," in *2024 American Control Conference (ACC)*, 2024, pp. 1873–1878.
- [26] 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*, vol. 29, no. 3, pp. 1203–1218, 2021.
- [27] M. W. Mueller, M. Hehn, and R. D'Andrea, "A computationally efficient motion primitive for quadcopter trajectory generation," *IEEE Transactions on Robotics*, vol. 31, no. 6, pp. 1294–1310, 2015.
- [28] B. Etkin, *Dynamics of Atmospheric Flight*. Dover Publications, 2005.
- [29] J. L. Meriam and L. G. Kraige, *Engineering Mechanics: Dynamics*, 8th ed. Hoboken, NJ: Wiley, 2015.
- [30] M. W. Mueller and R. D'Andrea, "Relaxed hover solutions for multicopters: Application to algorithmic redundancy and novel vehicles," *The International Journal of Robotics Research*, vol. 35, no. 8, pp. 873–889, 2016.
- [31] G. F. Franklin, J. D. Powell, and A. Emami-Naeini, *Feedback Control of Dynamic Systems*, 7th ed. Pearson, 2015.
- [32] A. Loquercio, E. Kaufmann, R. Ranftl, M. Müller, V. Koltun, and D. Scaramuzza, "Learning high-speed flight in the wild," *Science Robotics*, vol. 6, no. 59, p. eabg5810, 2021.