

# PyFlyt—UAV Simulation Environments for Reinforcement Learning Research

Jun Jet Tai, Jim Wong, Mauro S. Innocente, Nadjim Horri, James Brusey *Member, IEEE*,  
Swee King Phang *Senior Member, IEEE*

**Abstract**—Unmanned aerial vehicles (UAVs) have numerous applications, but their efficient and optimal flight can be a challenge. Reinforcement Learning (RL) has emerged as a promising approach to address this challenge, yet there is no standardized library for testing and benchmarking RL algorithms on UAVs. In this paper, we introduce PyFlyt, a platform built on the Bullet physics engine with native Gymnasium API support. PyFlyt provides modular implementations of simple components, such as motors and lifting surfaces, allowing for the implementation of UAVs of arbitrary configurations. Additionally, PyFlyt includes various task definitions and multiple reward function settings for each vehicle type. We demonstrate the effectiveness of PyFlyt by training various RL agents for two UAV models: quadrotor and fixed-wing. Our findings highlight the effectiveness of RL in UAV control and planning, and further show that it is possible to train agents in sparse reward settings for UAVs. PyFlyt fills a gap in existing literature by providing a flexible and standardised platform for testing RL algorithms on UAVs. We believe that this will inspire more standardised research in this direction.

**Index Terms**—PyFlyt, Reinforcement Learning, UAV Flight, Optimal Control

## I. INTRODUCTION

REINFORCEMENT Learning has rapidly advanced in the past decade, achieving various impressive feats [1]–[5]. Fundamentally, its allure boils down to its theoretical simplicity and generalisability, with growth fueled by open and accessible frameworks and software. Most notably, various environment API libraries [6], [7] and learning frameworks [8]–[10] have made standardising RL across a large range of tasks possible.

RL draws inspiration from control theory, evidenced by several RL tasks mimicking those in control theory. This also holds true for UAV flight, which has long held the interest of the controls engineering community. The classical approach to optimising UAV control involves testing in simulation. Popular open-source UAV simulation tools include Gazebo [11], FlightGear [12], X-Plane [13], AirSim [14] and UE4Sim [15]. However, their limited portability severely restricts their use in machine learning workflows, which usually run on

J. J. Tai, J. Wong, M. Innocente and J. Brusey are with Coventry University, Priory St, Coventry CV1 5FB, United Kingdom (emails: taij@uni.coventry.ac.uk, wongjim8@gmail.com, mauro.s.innocente@coventry.ac.uk, james.brusey@coventry.ac.uk)

N. Horri is with the School of Engineering, University of Leicester, University Road, Leicester LE1 7RH, United Kingdom (email: nmhorri@hotmail.com)

S. K. Phang is with the School of Engineering, Taylor’s University, Subang Jaya 47500, Malaysia (email: sweeking.phang@taylors.edu.my)

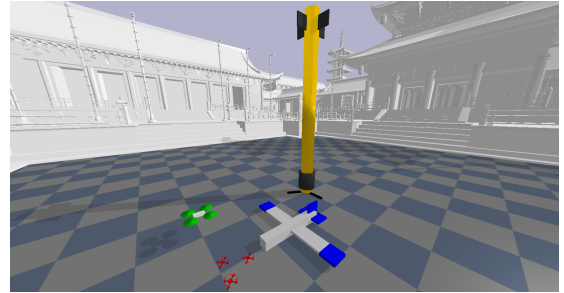


Fig. 1. The PyFlyt simulator, showing some prebuilt UAVs - one Rocket, one Fixedwing, one generic QuadX, and a cluster of three Crazyflie QuadXs.

highly parallel remote computing clusters. Moreover, with the exception of fully fledged robotics simulators like Gazebo, simulation is limited to one or two variants of UAVs, reducing their adaptability for RL environments using uncommon aircraft architectures like the quadplane [16]. Additionally, simulators like UE4Sim lack a physics model, rendering it impossible to modify physical UAV attributes such as motor thrust profile or moments of inertia.

In this work, we present an extensible and open source UAV design and simulation library named *PyFlyt* (pronounced pa-flaIt) that fills the void between complex UAV simulators and user-friendly machine learning libraries described in Section II. More formally, PyFlyt offers the following key features:

- *Ease of Use*: automatic handling of various physics, control and agent loop rates, collisions, multiple UAVs in the same environment, plus implementation of various classical control algorithms for low level flight control and extensible dynamics via the Bullet Physics Engine (Section III-B).
- *Configurability*: flexible construction of UAVs using modular implementations of basic components such as motors, boosters, lifting surfaces, gimbals and cameras (Section III-D).
- *Standardisation*: pre-implemented designs for various UAVs (Section III-E) along with corresponding environments compatible with the Gymnasium API (Section III-F), tested extensively with various RL algorithms (Section IV).
- *Quality of Life Improvements*: deterministic simulation depending on initialisation seed, Continuous Integration (CI) testing on the master repository and a packaging and release cycle to The Python Package Index (PyPI).

## II. RELATED WORKS

RL environments and UAV simulators are not a new; many already exist with various functionality and capabilities. In this section, we briefly cover the landscape of RL environment libraries and UAV simulation tools, including those catered for RL. At the end, we draw similarities between these tools and PyFlyt, and highlight the gap we aim to fill.

### A. Reinforcement Learning Environments

OpenAI’s release of the Gym library in 2016 [6] standardized benchmarking and interfacing for RL. Subsequently, various RL environment libraries built on the Gym API have emerged, including those based on video games [17], [18] or classic robotics problems [19], [20]. The original OpenAI Gym paper has been cited over 5000 times, and hundreds of well-documented and widely used environment libraries have since emerged, despite the existence of other APIs [7].

Recently, volunteer contributors to Gym have moved to *Gymnasium*, an open-source fork of Gym that will receive all further maintenance and development moving forward [21]. *Gymnasium* offers more official support for the `termination` and `truncation` API<sup>1</sup> and has seen popular libraries like *Gymnasium Robotics* [22], *ViZDoom* [17], *Minigrid* [23], *MiniWoB++* [24], and *MicroRTS-Py* [25] make the transition to its platform. In addition, a wrapper library by the authors of *Gymnasium* called *Shimmy* [26] allows for the conversion of environments in other APIs such as *DM-Env* and *DM-Control* into a *Gymnasium*-compatible one. The active development of *Gymnasium* and its constant adoption by other libraries make it an easy choice for *PyFlyt* to be built around the its API.

### B. UAV Simulators

When selecting a UAV simulator for a project, several factors need to be considered, including visual realism, physics requirements like turbulence or downwash, and *Hardware-In-The-Loop* (HITL) capabilities, among others. Commercially available simulators like *XPlane* and *FlightGear* can simulate a wide range of fixed-wing aircraft with pseudo-accurate world scenery. These simulators offer plugin systems that enable users to create extension modules, resulting in various research tools, including those that can be used with *Matlab/Simulink* for developing flight controllers [27]–[31]. Some of these tools have also been adapted for OpenAI’s *Gym* API interface [32], [33]. However, despite their capabilities, these simulators have complicated installation procedures, minimal Linux support, and require graphical user interfaces, making them less suitable for most machine learning workflows.

Multicopter simulators like *JMavSim* and Microsoft’s *AirSim* [14] are tools designed specifically for multicopter UAV simulation. Both software packages offer *Software-In-The-Loop* (SITL) and HITL simulation with the *PX4* [34] flight control firmware and *Robot Operating System* (ROS). *AirSim*, being the most popular, offers photorealistic rendering through

*Unreal Engine 4* at the expense of elevated computational requirements. However, creating custom multicopter UAV models for these simulators is challenging, making them less ideal when designing environments for RL algorithms revolving around UAV’s with different parameters, such as thrust profile or moment of inertia.

The most closely related work to ours is *gym-pybullet-drones* [35] and *PyFly* [36] (*PyFly*’s name was chosen before we knew of *PyFly*). Like *PyFlyt*, *Gym-pybullet-drones* is a *Bullet Physics Engine* based simulator, but it uses the older OpenAI *Gym* API and supports simulation for only the *Crazyflie 2.0* UAV [37]. *PyFly*, on the other hand, is a Python implementation of a 6 degree-of-freedom aerodynamic model for fixed-wing aircraft. Like *gym-pybullet-drones*, *PyFly* uses the older OpenAI *Gym* API, but lacks rendering capabilities and only has RL environments for attitude control. In comparison, *PyFlyt* allows for the construction of *both* multicopter and fixedwing UAVs, in arbitrary configurations such as the trirotor, hexarotor, tube-and-wing, flying-wing or biplane.

Overall, existing simulators provide various capabilities, but they also have limitations that make them less suitable for certain applications. *PyFlyt* aims to address these limitations by providing a more flexible and user-friendly platform for running RL algorithms on a range of configurable UAV platforms, all within a single library<sup>2</sup>.

## III. DESIGN OF PYFLYT

In this section, we break down the design methodology of *PyFlyt* by first describing its main goals in Section III-A before going into the overall architecture in Section III-B. The detail the modular UAV components in Section III-D and how they form the basic UAV platforms in Section III-E. Finally, we describe the default *Gymnasium* compatible environments in Section III-F and describe potential extensions, possibilities, and improvements in Section V.

### A. Main Goals

*PyFlyt* is a library primarily designed for researchers to develop and test complex RL algorithms for UAVs, but the development of classical control algorithms is also well within scope. It treats single agent simulations as multiagent simulations with a single agent. This allows easy extension to the multiagent setting. *PyFlyt* also automatically tracks collisions between different entities, allowing the testing and evaluation of collision avoidance algorithms either with obstacles or with other UAVs.

Furthermore, *PyFlyt* supports the construction of arbitrary UAV configurations by combining various basic components. This includes the addition of onboard cameras — gimballed or fixed — with RGB, depth maps, and segmentation maps, allowing for the development of vision-based RL algorithms.

Moreover, there is built-in support for developing generic onboard flight control. This simplifies the process of developing and testing RL algorithms by allowing researchers to map

<sup>1</sup><https://github.com/openai/gym/issues/2510>

<sup>2</sup>Available for install with a single command from: <https://pypi.org/project/PyFlyt/>

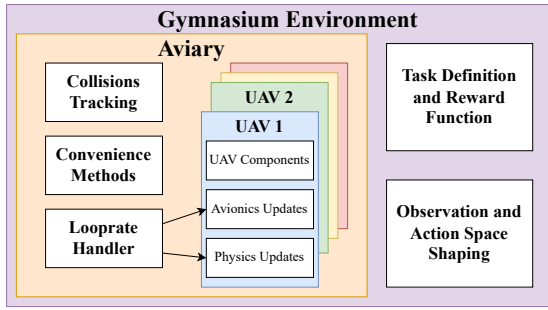


Fig. 2. Top level overview of the PyFlyt architecture, showing the composition of classes that make up core UAV flight engine and how it integrates into a Gymnasium environment. Coloured boxes describe a class definition, while white boxes describe various functionality of the classes.

RL agent actions to setpoints in the flight controller, rather than directly controlling raw actuator outputs. In addition to that, the library supports running RL, control, and physics at different frequencies. This leads to more realistic simulation, as the RL loop is unlikely to always be running at the fastest looprate. This capability is useful when doing research on a more macro level, such as when studying swarming behaviours [38].

Finally, the creation of Gymnasium compatible environments using PyFlyt is done as an auxiliary module instead of being integrated into the core architecture of PyFlyt. This is in contrast to the implementation in gym-pybullet drones, allowing researchers to utilize PyFlyt to construct arbitrary environments using any desired RL API. Several interesting and well tested environments have also been implemented for various vehicles. Their goal is to serve as a good baseline for upcoming research, as well as to serve as templates for users aiming to create their own environments.

### B. Overall Architecture

PyFlyt utilizes the Bullet Physics Engine, chosen for its open-source nature, fast C++ backend, flexible rendering options, and time-discrete steppable physics, Python bindings, and support for various robot file definitions. It was designed with extensibility in mind, allowing users to customize the simulation environment without relying on inheritance-based approaches [39].

The library's architecture is composed of various classes, shown in Figure 2. At the core lies the Aviary, serving as a domain for all UAVs. The Aviary handles collision tracking between entities and performs step scheduling according to the various looprates. It also includes convenience functions for setting setpoints, arming status, and flight modes for all UAVs.

The Aviary can accommodate any number of UAVs, each defined by a combination of a Unified Robot Description Format (URDF) file and a Python class. The URDF file describes the geometrical and physical properties of the UAV, including its mass, inertia, as well as visual and collision meshes, which can be imported as Digital Asset Exchange (DAE) files from standard CAD software. The Python class

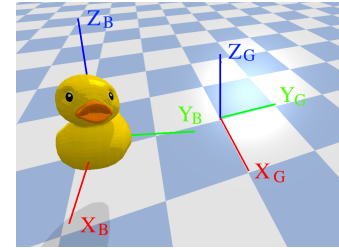


Fig. 3. An illustration of body (subscript  $B$ ) and ground (subscript  $G$ ) frame axes used within PyFlyt.

associated with each UAV inherits from a common Drone Class, which defines only the core functionality shared by all UAVs, necessary for interfacing with the Aviary. Each UAV defines its own base control laws, state spaces, actuator mappings, and physical dynamics interactions. Users can further store arbitrary parameters in a YAML Ain't Markup Language (YAML) file for easy configuration. The use of Python classes allows for a high degree of flexibility in defining complex interactions between the UAVs and their environment.

### C. Frame Conventions

Frame conventions are essential for describing the orientation and movement of objects in space. We use two reference frames, the ground frame and the body frame, shown in Figure 3. The ground frame defines three axes relative to the local horizontal plane and the direction of gravity. The body frame defines the axes relative to the body of the UAV being controlled [40]. We utilize the East-North-Up (ENU) frame convention, wherein for the ground frame, the X-axis points East, Y-axis points North, and Z-axis points Up (East and North are just cardinal references). On the body frame, the ENU convention defines the X, Y, and Z axes to point out the front, left, and upward direction from the UAV respectively.

To transform a vector  $\mathbf{v}_G \in \mathbb{R}^3$  from the ground frame to the body frame  $\mathbf{v}_B \in \mathbb{R}^3$ , we can use the following equation:

$$\mathbf{v}_{XYZ} = \mathbf{R} \times \mathbf{v}_{ENU} \quad (1)$$

Where  $\mathbf{R}$  denotes the rotation matrix from ground to body frame:

$$\mathbf{R} = \begin{bmatrix} c_\theta c_\psi & s_\phi c_\theta & -s_\theta \\ -s_\psi c_\phi + c_\psi s_\theta s_\phi & c_\psi c_\phi + s_\psi s_\theta s_\phi & c_\theta s_\phi \\ s_\psi s_\phi + c_\psi s_\theta c_\phi & -c_\psi s_\phi + s_\psi s_\theta c_\phi & c_\theta c_\phi \end{bmatrix} \quad (2)$$

Where  $\phi$ ,  $\theta$ , and  $\psi$  are the Euler angles describing the rotation between the two frames,  $c_*$  and  $s_*$  are shorts for  $\cos(*)$  and  $\sin(*)$  respectively.

### D. UAV Component Descriptions

The field of engineering often shares common mathematical models across similar components. For example, the performance difference between two DC electric motors with varying stator heights are typically limited to parameters such as motor ramp time constant, maximum rotational speed, and torque output [41]. To provide a clear abstraction, PyFlyt offers configurable implementations of common components used

in UAV construction. These components include vectorized implementations for brushless DC motor driven propellers, electric gimbals, fueled boosters, plus non-vectorized implementations for flapped lifting surfaces and cameras. In this section, we describe the modelling of each component and their configurations.

1) *Brushless DC Motor Driven Propellers*: A brushless DC motor driven propeller, simply termed motor here, is a propulsion unit consisting of a fixed pitch propeller paired to a brushless DC electric motor, commonly seen on various UAVs as a means of propulsion. We model the motor's RPM  $\omega$  as a function of the throttle input  $\zeta \in [-1, 1]$ , using a first-order transfer function plus noise depending on the motor time constant  $\tau_m$ , maximum RPM  $k_\omega$  and motor fluctuation standard deviation  $\sigma_m$ :

$$\omega_{t+1} = \left( \frac{1}{\tau_m r_p} (\zeta k_\omega - \omega_t) + \omega_t \right) (1 + \sigma_m \omega_t \eta) \quad (3)$$

Where the subscript  $t$  denotes the physics timestep,  $r_p$  denotes the physics looprate in Hz, and  $\eta$  is noise sampled from a standard Normal. Depending on the thrust coefficient  $k_F$  and torque coefficient  $k_T$ , this gets turned into a torque and force on the motor:

$$T_m = k_T \omega^2 \quad (4)$$

$$F_m = k_F \omega^2 \quad (5)$$

2) *Fueled Boosters*: Unlike electric motors, fueled boosters are propulsion units that produce no meaningful torque around their thrust axis and have limited throttleability. More crucially, they depend on a fuel source that depletes with usage, changing the mass and inertia properties of the UAV. Additionally, some boosters, typically of the solid fuel variety [42], cannot be extinguished and reignited, a property we call reignitability.

In PyFlyt, the booster is parameterized with total fuel mass  $m_{b_{\max}}$ , maximum fuel mass flow rate  $\dot{m}_{b_{\max}}$ , the diagonal elements of the moment of inertia matrix at full fuel load  $J_{xx_{\max}}$ ,  $J_{yy_{\max}}$ ,  $J_{zz_{\max}}$ , minimum thrust  $T_{b_{\min}}$ , maximum thrust  $T_{b_{\max}}$ , a boolean for reignitability  $\mathbb{I}_b$ , the booster ramp time constant  $\tau_b$  and a thrust output fluctuation standard deviation  $\sigma_b$ . The booster component also allows placement of the fuel tank at an arbitrary location on the UAV, useful for simulating jet fuel powered turbojet UAVs which have fuel stored in the main body and wings.

Fueled boosters have complex mathematical models that drastically change depending on altitude, temperature, total fuel remaining, and throttle setting [43]–[45]. In favour of simplicity, we adopt a simpler model for the booster in PyFlyt, useful as a baseline for developing more complex boosters should the user see necessary. The control inputs to the booster include a boolean ignition state  $\iota$  and a normalized throttle setting  $\beta \in [0, 1]$ . The minimum duty cycle is defined as the ratio of minimum thrust to maximum thrust:

$$\lambda_{\min} = \frac{k_{t_{\min}}}{k_{t_{\max}}} \quad (6)$$

Then, the current duty cycle of the booster  $\lambda \in [0, 1]$  changes according to the throttle setting through a first order transfer

function, plus some noise modelled similarly to that of the motor:

$$\lambda_{t+1} = \mathbb{I}_\lambda \mathbb{I}_{m_b > 0} \left( \frac{1}{\tau_b r_p} (\beta(1 - \lambda_{\min}) - \lambda_t) + \lambda_t \right) (1 + \sigma_b \lambda_t \eta) \quad (7)$$

Where  $\mathbb{I}_\lambda$  is an indicator on whether the booster is lit while  $\mathbb{I}_{m_b > 0}$  indicates the availability of remaining fuel.

The fuel consumption  $\dot{m}_b$  of the booster depends on the duty cycle:

$$\dot{m}_b = -\lambda k_{f_r} \mathbb{I}_\lambda \mathbb{I}_{m_b > 0} \quad (8)$$

To prevent reignition of un-reignitable boosters, the booster lit indicator is recomputed at every timestep:

$$\mathbb{I}_{\lambda, t+1} = (\neg \mathbb{I}_b \wedge \mathbb{I}_{\lambda, t}) \vee \iota_t \quad (9)$$

In short, a booster that is not reignitable cannot be extinguished, while one that is reignitable can be extinguished and reignited at will.

From here, the thrust output depends on the duty cycle:

$$T_b = \lambda T_{b_{\max}} \quad (10)$$

And inertia properties of the fuel tank depend on the remaining fuel balance:

$$J_i = \frac{m_b}{m_{b_{\max}}} J_{i_{\max}}; i \in \{xx, yy, zz\} \quad (11)$$

There are several deficiencies with this model. First, it does not model the sloshing of liquid fuel within the fuel tank. Second, it does not model engine deterioration with the number of reignitions. Third, solid rocket fuel boosters can have non-linear thrust profiles depending on the amount of fuel remaining [46], this is not modelled in the default booster model. Lastly, there is presently no limit on the number of times a booster can be reignited, unlike most liquid-fueled rocket engines [47]. That said, it is very difficult to have a parameterised booster model that suits the wide range of possible boosters in literature. For that reason, the booster component is designed with generalisability and simplicity in mind; interested users are encouraged to modify the model depending on their needs.

3) *Electric Gimbals*: An electric gimbal is a powered joint that rotates around two axes. The first servo motor rotates the joint about the first axis,  $\mathbf{c}_a$ , referenced to the body frame. Similarly, the second servo motor rotates the joint about the second axis,  $\mathbf{c}_b$ , also referenced to the body frame. Typically, the two rotational axes are orthogonal to each other, although this is not a necessity. There are several uses for such gimbaling, but the most common application is for thrust vectoring of propulsion systems [48]–[50], especially in scaled model rocketry [51], [52].

The gimbal is controlled using a normalized input  $\xi_i \in [-1, 1]$ ;  $i \in a, b$  and produces three dimensional rotation matrix about the UAV's body frame. Configurable parameters include the unit vectors for the two rotational axes  $\mathbf{c}_a, \mathbf{c}_b$ , maximum gimbaling angles  $k_{g-a}, k_{g-b}$ , and actuation time constant  $\tau_g$ . We utilize the same first order transfer function when converting input setpoint to actual gimbal angles  $\epsilon$ :

$$\epsilon_{i, t+1} = \frac{1}{\tau_g r_p} (\xi_i k_{g-i} - \epsilon_{i, t}) + \epsilon_{i, t} \quad (12)$$

Where  $i \in \{a, b\}$  is an axis index. To obtain the final rotation matrix, Rodrigues' rotation formula is used [53]. Let the unit vector for an axis of rotation be  $\mathbf{c}_i = \langle u_x, u_y, u_z \rangle$ . We then define  $\mathbf{W}$  as:

$$\mathbf{W}_i = \begin{bmatrix} 0 & -u_z & u_y \\ u_z & 0 & -u_x \\ -u_y & u_x & 0 \end{bmatrix} \quad (13)$$

The Rodrigues' rotation matrix is then constructed as:

$$\mathbf{R}_i = \mathbf{I} + \sin \epsilon_i \mathbf{W}_i + \left(2 \sin^2 \frac{\epsilon_i}{2}\right) \mathbf{W}_i^2 \quad (14)$$

Where  $\mathbf{I}$  is the identity matrix. The final rotation matrix is thus:

$$\mathbf{R} = \mathbf{R}_a \times \mathbf{R}_b \quad (15)$$

The term multiplying  $\mathbf{W}_i^2$  is typically written as  $1 - \cos(\epsilon_i)$ . However, the sine variant is more numerically stable for values near  $2n\pi; n \in \mathbb{Z}$ . Since  $\mathbf{W}_i$  and  $\mathbf{W}_i^2$  only depend on the axis of rotation which is predetermined at initialization, they can be precomputed and stored, resulting in a very fast computation for the rotation matrix around each axis.

4) *Flapped Lifting Surface*: Lifting surfaces are effectively flat plates with an aerofoil cross section that travel through the air. In doing so, they experience a net force resulting in a lifting force  $F_{s_{\text{lift}}}$ , a drag force  $F_{s_{\text{drag}}}$  and a net moment  $T_s$  about the aerodynamic center.

The forces and moments that act on them depend on a number of factors, most notably the angle-of-attack  $\alpha$  and freestream velocity  $V_\infty$  according to the equation:

$$\begin{bmatrix} F_{s_{\text{lift}}} \\ F_{s_{\text{drag}}} \\ T_s \end{bmatrix} = \begin{bmatrix} C_L \\ C_D \\ C_M \end{bmatrix} \frac{1}{2} \rho A_s V_\infty^2 \quad (16)$$

Where  $A_s$  is the area of the lifting surface and  $C_L, C_D$  and  $C_M$  represent dimensionless lift, drag, and pitching moment coefficients, forming what we term here as a coefficient vector.

The coefficient vector can be captured in a lift, drag, and polar curve — graphical representations of how the values vary with Angle of Attack (AoA). These curves, also known as a performance profile, are unique to each aerofoil and are often obtained through extensive Computational Fluid Dynamics (CFD) simulation or wind tunnel testing.

In our approach, we adopt the method proposed by [54] to obtain the performance profile. Specifically, a lifting surface can be parameterized by its zero-lift AoA  $\alpha_0$ , positive stall AoA  $\alpha_{\text{stall}_+}$ , negative stall AoA  $\alpha_{\text{stall}_-}$ , zero-lift drag coefficient  $C_{D_0}$ , lift coefficient slope  $C_{L_\alpha}$ , and a dimensionless viscosity correction factor  $k_{s_\eta}$ <sup>3</sup>. The exact equations for curve generation are quite involved, interested readers are instead directed to the exact equations in [54] or to implementations in code<sup>4</sup>.

The lifting surfaces in PyFlyt can have actuated flaps, hence the term flapped lifting surface. This flap is parameterized with a flap-to-chord ratio — a ratio of how much of the lifting surface is movable, a flap deflection limit  $\delta_{\text{max}}$ , and

an actuation time constant  $\tau_s$ . The flap allows the lifting surface to change its aerofoil geometry, consequently altering its performance profile curves. We employ an approximation by [56] to model this modification and utilize the familiar first order transfer function to model the flap deflection  $\delta$  as a function of time and normalized actuation input  $\nu \in [-1, 1]$ :

$$\delta_{t+1} = \frac{1}{\tau_s r_p} (\nu \delta_{\text{max}} - \delta_t) + \delta_t \quad (17)$$

Our approximation is generally valid for normal operating conditions, but may not accurately model the aerodynamic behavior of the lifting surface in extreme conditions, such as when experiencing major side slip, highly turbulent air, or supersonic flows.

5) *Camera*: The camera component allows the development of UAVs with vision-based capabilities. To further expand on this idea, there are no limits to the number of cameras a UAV can have, and no restrictions on their locations relative to the UAV. This flexibility allows simulating stereo vision for depth perception [57], multiple viewpoints for collision avoidance [58], or multi-camera resolution [59].

The camera in PyFlyt is parameterized with an image pixel height  $h$  and width  $w$ , field of view in degrees, camera orientation and pose relative to the UAV's body frame, and whether the camera is gimbal stabilized. The option for gimbal stabilization forces the camera to utilize rotation about the ground frame X and Y axes, providing a locked horizon view with the viewpoint following the body frame yaw.

When queried, each camera returns a set of 3 images readily obtained from the Bullet Physics Engine — an RGB image in  $\mathbb{R}^{w \times h \times 4}$ , a depth map in  $\mathbb{R}_+^{w \times h \times 1}$ , and a segmentation map in  $\mathbb{Z}^{w \times h \times 1}$  with values indicating the ID of the entity from which the pixel came from.

### E. UAV Model Description

Several default UAV models built using the components discussed in Section III-D are shipped with PyFlyt. Furthermore, it is possible to spawn multiples of each UAV in the same environment and control them simultaneously with dynamic interactions. In this section, we detail the design of three particular models that come with PyFlyt - a quadrotor, a fixedwing, and a rocket.

1) *QuadX*: The *QuadX* UAV describes a multirotor UAV in the Quad-X configuration as described by ArduPilot<sup>5</sup> and PX4<sup>6</sup>. In general, all quadrotor UAVs share a common mathematical model, albeit with varying physical and geometric parameters, such as mass, inertia, and size, as well as corresponding control parameters. PyFlyt provides a generic implementation of the QuadX model utilizing the vectorized motor component implementation. To construct arbitrary QuadX UAVs, a URDF and YAML file pair is simply passed to the underlying QuadX class to implement all mass, geometric, and control parameters of a QuadX UAV. As a result, users are able to construct QuadX UAVs without needing to modify the underlying Python code.

<sup>3</sup>Page 112 of [55].

<sup>4</sup>[https://github.com/jjshoots/PyFlyt/blob/master/PyFlyt/core/abstractions/lifting\\_surfaces.py#L74](https://github.com/jjshoots/PyFlyt/blob/master/PyFlyt/core/abstractions/lifting_surfaces.py#L74)

<sup>5</sup><https://ardupilot.org/copter/docs/connect-escs-and-motors.html#quadcopter>

<sup>6</sup>[https://dev.px4.io/v1.10/en/airframes/airframe\\_reference.html#quadrotor-x](https://dev.px4.io/v1.10/en/airframes/airframe_reference.html#quadrotor-x)

In addition, the PyFlyt implementation of the QuadX UAV comes with various default control modes that researchers can build off using the native `register_controller` method. These control modes utilize a cascaded Proportional Integral Derivative (PID) architecture [60] to achieve increasingly higher levels of attitude control. The implemented controllers allows for the following setpoint definitions:

- Raw motor outputs
- Angular velocity control and normalized thrust output
- Angular position control and climb rate
- Linear velocity control and yaw rate
- Linear velocity control and yaw angle
- Position control and yaw rate
- Position control and yaw angle

As a reference, we provide two different QuadX models in PyFlyt. The first is the model of a Bitcraze Crazyflie 2.x, leveraging the extensive amounts of system identification work available [61]–[63]. However, instead of adhering to a true 1:1 representation of the Crazyflie quadrotor, we modify the model to have an 8:1 thrust to weight ratio, allowing for more interesting and dynamic performance characteristics in downstream RL tasks. The second model is that of a generic 1 kg quadrotor UAV with parameters mimicking a generic F450 quadrotor UAV<sup>7</sup>, meant to serve as a starting point for users implementing models of their own without using a DAE file.

2) *Fixedwing*: The *Fixedwing* UAV is a tube-and-wing aircraft consisting of five lifting surface components—main wing, left aileron, right aileron, horizontal tail, vertical tail—and a motor. All of the lifting surfaces except the main wing are flapped, and they have parameters of an ideal 2D aerofoil. The ailerons have a maximum flap deflection of 30°, while the rudder and elevators are limited to 20°. It has a wingspan of 2.4 m, a nose to tail length of 1.55 m, with a 0.8 thrust to weight ratio. The total mass of the UAV is 2.35 kg, with 1 kg uniformly distributed in the fuselage, 0.5 kg in the main wing and the remaining mass distributed among the other lifting surface components.

3) *Rocket*: The *Rocket* UAV is a 1:10th scale rocket closely modeled after the design of SpaceX’s Falcon 9<sup>8</sup> v1.2 first stage and interstage. Mass and geometry properties were extracted from a datasheet by Space Launch Report<sup>9</sup>. It consists of booster, gimbal, and lifting surface components. The SpaceX Falcon 9 features grid fins [64] as supplementary control actuators. However, at the time of writing, limited studies on the aerodynamic properties of the Falcon 9’s grid fins were made available to the public [65]. As a consequence to this shortcoming, the Rocket’s control fin properties were simply approximated. In addition, the Bullet Physics Engine approximating the world as a flat 2D surface, it is not feasible to simulate the orbital mechanics of the Rocket within the PyFlyt framework.

<sup>7</sup><https://www-v1.dji.com/flame-wheel-arf/feature.html>

<sup>8</sup><https://www.spacex.com/vehicles/falcon-9/>

<sup>9</sup><https://web.archive.org/web/20170825204357/spacelaunchreport.com/falcon9ft.html#f9stglog>

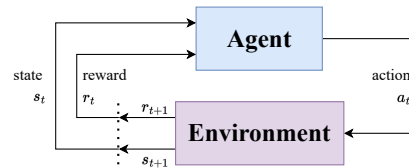


Fig. 4. A classic model of an MDP represented as a cyclic graph.

## F. Gymnasium Environments

The primary goal of PyFlyt is to facilitate RL research for UAVs. To that end, we provide several default RL environments utilizing the Gymnasium API for a variety of UAV models with various tasks and reward functions. This section briefly covers the mathematical model of a Markov Decision Process (MDP) and the Gymnasium API, before providing details about the various environments provided.

1) *Markov Decision Processes*: The Markov Decision Process (MDP) is a mathematical framework used to model decision-making in situations where outcomes are uncertain. It consists of a set of states  $s \in \mathbf{S}$ , actions  $a \in \mathbf{A}$ , and rewards  $r \in \mathbb{R}$ . At each time step, the agent chooses an action based on the current state, which leads to a new state and a reward. The goal is to find a policy that maximizes the expected cumulative reward over time. Figure 4 illustrates the MDP as a cyclic graph.

2) *Gymnasium API*: Gymnasium is a standard API for defining RL environments. It is the maintained fork of the OpenAI Gym library, with a simple Pythonic interface. The Gymnasium API allows for the universal definition of state and action spaces through a set of default constructs. It also defines environment termination and truncation behaviour, convenient environment wrappers, as well as a plethora of native environments from which users can use as benchmarks or use as reference.

Generally, libraries that utilize the Gymnasium API can be segregated into one of two parties — learning libraries which implement a suite of well-studied and effective learning algorithms, and environment libraries which allow for the benchmarking and construction of RL problems. The existence of such a standardized API allows for the application of general purpose RL algorithms to arbitrary problems. In this context, PyFlyt is a means of standardizing the construction of UAVs for RL problems, with which we provide several default environments built to match the Gymnasium API.

### 3) *Base Environment Parameters*:

a) *Observation Spaces*: All currently implemented environments within PyFlyt utilize either the `Box` space, or a `Dict` space consisting of a `Box` space and a `Sequence` space. The `Box` space defines a space in  $\mathbb{R}^{n \times \dots}$ , while a `Sequence` space defines a space of a sequence of elements in  $\mathbb{R}^{n \times \dots}$ . Meanwhile, the `Dict` space is a means of combining various spaces together to form a composite space.

In Pyflyt, the UAV’s attitude, previous action, and component statuses are represented as a `Box` space. Waypoint reaching environments have targets defined as a `Sequence` space, which is combined with the aforementioned `Box` state

through a `Dict`. More concisely, environments without waypoints have the following state space:

$$\mathbf{s}_{\text{attitude}} = [\dot{\theta}_B, \theta_G, \dot{\mathbf{p}}_B, \mathbf{p}_G, \mathbf{a}_{t-1}, \mathbf{s}_{\text{aux}}] \quad (18)$$

Where  $\theta$  represents the angular position of the UAV and  $p$  represents the linear position, both referenced to either body (subscript  $B$ ) or ground (subscript  $G$ ) frames. In addition,  $\mathbf{s}_{\text{aux}}$  refers to auxiliary information, such as normalized lifting surface deflection angle, motor duty cycle, or booster fuel and ignition statuses.

For environments with waypoints, the state space has the following form:

$$\mathbf{s} = \{\mathbf{s}_{\text{attitude}}, \mathbf{s}_{\text{target\_deltas}}\} \quad (19)$$

Where  $\mathbf{s}_{\text{target\_deltas}}$  represent the sequence of waypoints remaining in the environment:

$$\mathbf{s}_{\text{target\_deltas}} = [\Delta \mathbf{p}_1^T, \Delta \mathbf{p}_2^T, \dots, \Delta \mathbf{p}_n^T] \quad (20)$$

Where  $\Delta \mathbf{p}_i$  represent the vector to the  $i$ -th remaining waypoint defined in the UAV's body frame.

*b) Action Spaces:* All action spaces for currently implemented environments use a `Box` action space. For QuadX environments, the action is a four wide vector representing roll, pitch and yaw rates along with a normalized thrust command. For Fixedwing environments, the action is a four wide vector representing normalized aileron, horizontal and vertical stabiliser deflection angles along with a normalized thrust command. For Rocket environments, the action is a seven wide vector representing grid fin deflections in the X plane, Y plane, and cyclical deflections, along with the booster ignition, throttle settings and two values for booster gimbaling about the body frame XY axes.

*c) Looprates:* All environments operate with a physics looprate of 240 Hz while controllers run at 120 Hz—allowing UAV models to operate lower level controllers under the hood. At the environment level, all interactions between an agent and the environment happen at a reconfigurable rate of 30 Hz. This configuration allows the RL agents to send setpoints to the UAVs instead of raw actuator commands, a scheme similar to human pilots for racing quadrotor UAVs [66].

*4) Reward Functions:* In RL, the reward function directly influences the learning dynamics and final policy obtained from a particular task. A dense reward function is one that provides consistent feedback to the agent about its actions. This means that the reward signal is provided at almost every time step, and the agent can easily understand the impact of its actions on the environment.

On the other hand, a sparse reward is one that is given infrequently and provides limited information to the agent about its actions. This means that the reward signal is only provided occasionally or at certain milestones, making it harder for the agent to understand the impact of its actions on the environment. Sparse rewards are commonly encountered in real-world tasks where the reward function is not immediately obvious and may require additional exploration in different parts of the environment to discover it.

In general, dense rewards can encourage an agent to learn quickly and converge to an optimal solution. However, sparse

rewards often lead to a more robust and generalised solution as the agent is forced to learn more about the environment through additional exploration. In addition, sparse rewards represent a more holistic view of the goal of the underlying task, leading to agents that are more focused on achieving the primary objective without getting distracted by irrelevant factors [67]. All of the default environments in PyFlyt come with a dense and sparse reward variants, available through the `sparse_reward` argument on environment initialization.

*5) Provided Environments:* At the time of writing, there are four pre-built environments in PyFlyt, described in this section.

*a) PyFlyt/QuadX-Hover-v0:* The task of this environment is for an agent to hover the QuadX UAV indefinitely. In the dense setting, the reward of the environment is described as follows:

$$r = \begin{cases} -100 & \text{if crash,} \\ -\|\theta, \phi\| - \|[x, y, z - 1]\| & \text{otherwise} \end{cases} \quad (21)$$

Where  $\theta$  and  $\phi$  are the pitch and roll angles in radians,  $x$ ,  $y$  and  $z$  describe the linear position of the UAV.

In the sparse reward variant, the reward is simply -0.1 for each timestep and -100 for crashing.

*b) PyFlyt/QuadX-Waypoints-v0:* This environment describes the task of getting a QuadX UAV to reach a series of randomly generated waypoints in 3D space in a specific order. By default, there are five waypoints in total, but this number is reconfigurable. The dense reward function of the environment is as follows:

$$r = \begin{cases} -100 & \text{if crash,} \\ c_a \delta^{-1} - c_b \dot{\delta} & \text{if } \dot{\delta} < 0, \\ 100 & \text{if waypoint reached,} \\ \delta^{-1} & \text{otherwise} \end{cases} \quad (22)$$

Where  $\delta$  is the displacement from the UAV to the next target, and  $c_a$  and  $c_b$  are scaling coefficients, set at 0.1 and 3 respectively. To avoid reward function blowup caused by the  $\delta^{-1}$  component, a waypoint is considered reached when the UAV's center of mass is some distance  $d$  away from the waypoint. This distance is a reconfigurable parameter that is set at 0.2 m by default. In the sparse reward setting, the reward is simply -100 for crashing and 100 for reaching a waypoint.

*c) PyFlyt/Fixedwing-Waypoints-v0:* This environment is similar to PyFlyt/QuadX-Waypoints-v0 but is modified to use the Fixedwing UAV. Due to the increase in size of the UAV model,  $d$  is set at 2,  $c_a$  is set at 1 and  $c_b$  is set at 3. In addition, the waypoints are also generated much further apart, at a scale of 20 times that of the QuadX environment.

*d) PyFlyt/Rocket-Landing-v0:* The objective in this environment is to safely land the Rocket UAV model on a 4-meter diameter pad from an initial drop height of 450 meters—about  $100\times$  the height of the Rocket. The landing is considered successful if the rocket contacts the pad with a velocity of less than 1 m/s and comes to an upright halt.

In the dense reward setting, the rocket is penalized according to the absolute distance between the center of the rocket and the landing pad,  $d_{\text{pad}}$ :

$$r_d = -0.2 \cdot d_{\text{pad}} \quad (23)$$

Note that  $d_{\text{pad}} > 0$  as it is the distance from the center of mass of the rocket to the landing pad. We further define a penalty for having the rocket favour lower absolute velocities when near the landing pad:

$$r_v = -d_{\text{pad}}^{-2} \cdot v \quad (24)$$

Finally, a termination reward was computed based on the landing condition:

$$r_{\text{pad}} = \begin{cases} -20 & \text{if fatal,} \\ 100 & \text{if safe} \end{cases} \quad (25)$$

Combined, this leads to a reward function of:

$$r = r_d + r_v + \mathbb{I}_{\text{pad}} r_{\text{pad}} - 100 \mathbb{I}_{\text{ground}} \quad (26)$$

Where  $\mathbb{I}_{\text{pad}}$  indicates whether the rocket is on the landing pad and  $\mathbb{I}_{\text{ground}}$  is an indicator for ground contact. More concisely, we penalize the rocket heavily for landing outside of the landing pad, while awarding policies that favour low velocities near the landing pad leading to safe touchdowns. To increase environment difficulty and prevent the execution of otherwise irrelevant manoeuvres, the Rocket starts with only 1% of total fuel—just enough for one landing attempt.

#### IV. REINFORCEMENT LEARNING RESULTS

Preliminary results for training RL algorithms in `PyFlyt/QuadX-Waypoints-v0` and `PyFlyt/Fixedwing-Waypoints-v0` are provided here, as we think these represent the an interesting class of problems for UAV flight. The environments were tested using Soft Actor Critic (SAC) [68] in the dense reward setting, while the sparse reward variants were tested using Critic Confidence Guided Exploration (CCGE) [69] and Advantage Weighted Actor Critic (AWAC) [70] by bootstrapping off a policy partially trained using SAC. We obtain 50 learning curves for each environment-algorithm combination over different random seeds. We then plot the interquartile means of the learning curves along with their bootstrapped confidence intervals using the RLiabile library [71]. Our results, shown in Figure 5, are meant to show the feasibility of training agents in these environments across multiple different algorithms. Examples of flight trajectories for each environment are also shown in Figure 6 <sup>10</sup>.

#### V. POSSIBLE EXTENSIONS AND FUTURE WORK

PyFlyt v1.0 has been released and is ready for use. That said, further improvements have been considered as part of PyFlyt’s development roadmap. First, we plan to implement default creation of multi-agent environments for multi-agent

<sup>10</sup>Videos of trajectories are available at <https://github.com/jjshoots/PyFlyt/blob/master/readme.md>

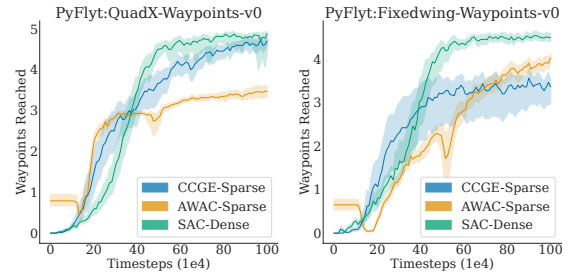


Fig. 5. Results of training RL agents in PyFlyt on the different environments.

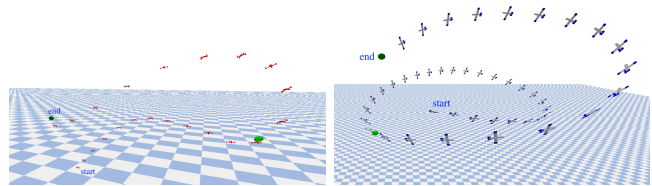


Fig. 6. Example trajectories of trained agents in `PyFlyt/QuadX-Waypoints-v0` and `PyFlyt/Fixedwing-Waypoints-v0`, both with only two waypoints for easier legibility.

RL, with PettingZoo support as our primary focus [72]. Various additional UAV types are also currently being modelled, such as quadplane, flying wing, and hexacopter UAVs. We also aim to introduce a number of aerodynamics associated simulation attributes, such as Dryden turbulence modeling, down wash, and aerodynamic shielding. Finally, user experience and documentation are vital components of any library, we therefore aim to provide clear, concise documentation, examples, and tutorials for PyFlyt, with the intention of reducing the barrier of entry for new users and increasing its adoption within the community.

#### VI. CONCLUSION

This paper introduced PyFlyt, a novel open-source simulation tool designed for machine learning and robotics workflows in UAV research. With its Pythonic interface and modular implementation of UAV components, PyFlyt enables the configuration of arbitrary UAV types, and offers prebuilt UAV implementations. A number of default Gymnasium environments were also introduced with both dense and sparse reward variants. We further present successful completion of the default environment by three distinct RL algorithms within the PyFlyt environment, underscoring its potential as a valuable tool for future UAV research. We are confident that PyFlyt will facilitate innovative research in robotics and autonomous UAV flight, and we look forward to see its utilization by the community.

#### REFERENCES

[1] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, *et al.*, “Mastering the game of go without human knowledge,” *nature*, vol. 550, no. 7676, pp. 354–359, 2017.



- [2] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, *et al.*, “A general reinforcement learning algorithm that masters chess, shogi, and go through self-play,” *Science*, vol. 362, no. 6419, pp. 1140–1144, 2018.
- [3] C. Berner, G. Brockman, B. Chan, V. Cheung, P. D. Keabiak, C. Dennison, D. Farhi, Q. Fischer, S. Hashme, C. Hesse, *et al.*, “Dota 2 with large scale deep reinforcement learning,” *arXiv preprint arXiv:1912.06680*, 2019.
- [4] O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi, R. Powell, T. Ewalds, P. Georgiev, *et al.*, “Grandmaster level in starcraft ii using multi-agent reinforcement learning,” *Nature*, vol. 575, no. 7782, pp. 350–354, 2019.
- [5] J. Degraeve, F. Felici, J. Buchli, M. Neunert, B. Tracey, F. Carpanese, T. Ewalds, R. Hafner, A. Abdolmaleki, D. de Las Casas, *et al.*, “Magnetic control of tokamak plasmas through deep reinforcement learning,” *Nature*, vol. 602, no. 7897, pp. 414–419, 2022.
- [6] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, “Openai gym,” *arXiv preprint arXiv:1606.01540*, 2016.
- [7] A. Muldal, Y. Doron, J. Aslanides, T. Harley, T. Ward, and S. Liu, “dm\_env: A python interface for reinforcement learning environments,” 2019.
- [8] A. Raffin, A. Hill, M. Ernestus, A. Gleave, A. Kanervisto, and N. Dornmann, “Stable baselines3,” 2019.
- [9] E. Liang, R. Liaw, R. Nishihara, P. Moritz, R. Fox, K. Goldberg, J. Gonzalez, M. Jordan, and I. Stoica, “Rllib: Abstractions for distributed reinforcement learning,” in *International Conference on Machine Learning*, pp. 3053–3062, PMLR, 2018.
- [10] S. Huang, R. F. J. Dossa, C. Ye, J. Braga, D. Chakraborty, K. Mehta, and J. Araújo, “Cleanrl: High-quality single-file implementations of deep reinforcement learning algorithms,” 2021.
- [11] N. Koenig and A. Howard, “Design and use paradigms for gazebo, an open-source multi-robot simulator,” in *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)(IEEE Cat. No. 04CH37566)*, vol. 3, pp. 2149–2154, IEEE, 2004.
- [12] A. R. Perry, “The flightgear flight simulator,” in *Proceedings of the USENIX Annual Technical Conference*, vol. 686, pp. 1–12, 2004.
- [13] R. Garcia and L. Barnes, “Multi-uav simulator utilizing x-plane,” in *Selected papers from the 2nd International Symposium on UAVs, Reno, Nevada, USA June 8–10, 2009*, pp. 393–406, Springer, 2010.
- [14] S. Shah, D. Dey, C. Lovett, and A. Kapoor, “Airsim: High-fidelity visual and physical simulation for autonomous vehicles,” in *Field and Service Robotics: Results of the 11th International Conference*, pp. 621–635, Springer, 2018.
- [15] M. Mueller, V. Casser, J. Lahoud, N. Smith, and B. Ghanem, “Ue4sim: A photo-realistic simulator for computer vision applications,” 2017.
- [16] A. Mathur and E. M. Atkins, “Design, modeling and hybrid control of a quadplane,” in *AIAA Scitech 2021 Forum*, p. 0374, 2021.
- [17] M. Kempka, M. Wydmuch, G. Runc, J. Toczek, and W. Jaśkowski, “Vizdoom: A doom-based ai research platform for visual reinforcement learning,” in *2016 IEEE conference on computational intelligence and games (CIG)*, pp. 1–8, IEEE, 2016.
- [18] S. Milani, N. Topin, B. Houghton, W. H. Guss, S. P. Mohanty, K. Nakata, O. Vinyals, and N. S. Kuno, “Retrospective analysis of the 2019 minerl competition on sample efficient reinforcement learning,” in *NeurIPS 2019 Competition and Demonstration Track*, pp. 203–214, PMLR, 2020.
- [19] W. Koch, R. Mancuso, R. West, and A. Bestavros, “Reinforcement learning for uav attitude control,” *ACM Transactions on Cyber-Physical Systems*, vol. 3, no. 2, pp. 1–21, 2019.
- [20] N. G. Lopez, Y. L. E. Nuin, E. B. Moral, L. U. S. Juan, A. S. Rueda, V. M. Vilches, and R. Kojcev, “gym-gazebo2, a toolkit for reinforcement learning using ros 2 and gazebo,” *arXiv preprint arXiv:1903.06278*, 2019.
- [21] Farama Foundation, “Announcing the farama foundation.” Website, October 2022.
- [22] R. de Lazcano, K. Andreas, J. J. Tai, S. R. Lee, and J. Terry, “Gymnasium robotics,” 2023.
- [23] M. Chevalier-Boisvert, L. Willems, and S. Pal, “Minimalistic gridworld environment for gymnasium,” 2018.
- [24] T. Shi, A. Karpathy, L. Fan, J. Hernandez, and P. Liang, “World of bits: An open-domain platform for web-based agents,” in *International Conference on Machine Learning*, pp. 3135–3144, PMLR, 2017.
- [25] S. Huang, S. Ontañón, C. Bamford, and L. Grela, “Gym-μrts: toward affordable full game real-time strategy games research with deep reinforcement learning,” in *2021 IEEE Conference on Games (CoG)*, pp. 1–8, IEEE, 2021.
- [26] Jun Jet Tai, Mark Towers, E. Tower, and J. Terry, “Shimmy: Gymnasium and pettingzoo wrappers for commonly used environments.”
- [27] L. R. Ribeiro and N. M. F. Oliveira, “Uav autopilot controllers test platform using matlab/simulink and x-plane,” in *2010 IEEE Frontiers in Education Conference (FIE)*, pp. S2H–1, IEEE, 2010.
- [28] A. Bittar, H. V. Figueiredo, P. A. Guimaraes, and A. C. Mendes, “Guidance software-in-the-loop simulation using x-plane and simulink for uavs,” in *2014 International Conference on Unmanned Aircraft Systems (ICUAS)*, pp. 993–1002, IEEE, 2014.
- [29] H. V. Figueiredo and O. Saotome, “Simulation platform for quadricopter: Using matlab/simulink and x-plane,” in *2012 Brazilian Robotics Symposium and Latin American Robotics Symposium*, pp. 51–55, IEEE, 2012.
- [30] E. Sorton and S. Hammaker, “Simulated flight testing of an autonomous unmanned aerial vehicle using flightgear,” in *Infotech@ Aerospace*, p. 7083, 2005.
- [31] J. Ying, H. Luc, J. Dai, and H. Pan, “Visual flight simulation system based on matlab/flightgear,” in *2017 IEEE 2nd Advanced Information Technology, Electronic and Automation Control Conference (IAEAC)*, pp. 2360–2363, IEEE, 2017.
- [32] D. J. Richter and R. A. Calix, “Qplane: an open-source reinforcement learning toolkit for autonomous fixed wing aircraft simulation,” in *Proceedings of the 12th ACM Multimedia Systems Conference*, pp. 261–266, 2021.
- [33] S. Chin, A. Wood, A. Sidney, B. Tarpa, and R. Ross, “Gymfg: a framework with a gym interface for flightgear,” 2020.
- [34] L. Meier, D. Honegger, and M. Pollefeys, “Px4: A node-based multithreaded open source robotics framework for deeply embedded platforms,” in *2015 IEEE international conference on robotics and automation (ICRA)*, pp. 6235–6240, IEEE, 2015.
- [35] 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)*, pp. 7512–7519, IEEE, 2021.
- [36] E. Bøhn, E. M. Coates, S. Moe, and T. A. Johansen, “Deep reinforcement learning attitude control of fixed-wing uavs using proximal policy optimization,” in *2019 International Conference on Unmanned Aircraft Systems (ICUAS)*, pp. 523–533, IEEE, 2019.
- [37] W. Giernacki, M. Skwierczyński, W. Witwicki, P. Wroński, and P. Kozierski, “Crazyflie 2.0 quadrotor as a platform for research and education in robotics and control engineering,” in *2017 22nd International Conference on Methods and Models in Automation and Robotics (MMAR)*, pp. 37–42, IEEE, 2017.
- [38] M. S. Innocente and P. Grasso, “Self-organising swarms of firefighting drones: Harnessing the power of collective intelligence in decentralised multi-robot systems,” *Journal of Computational Science*, vol. 34, pp. 80–101, 2019.
- [39] O. Alam and J. Kienzle, “Designing with inheritance and composition,” in *Proceedings of the 3rd international workshop on Variability & Composition*, pp. 19–24, 2012.
- [40] S. K. Phang, K. Li, B. M. Chen, T. H. Lee, K. P. Valavanis, and G. J. Vachtsevanos, “Systematic design methodology and construction of micro aerial quadrotor vehicles,” *Handbook of Unmanned Aerial Vehicles*, pp. 181–206, 2014.
- [41] N. Hemati and M.-C. Leu, “A complete model characterization of brushless dc motors,” *IEEE Transactions on Industry Applications*, vol. 28, no. 1, pp. 172–180, 1992.
- [42] A. Abdelraouf, O. Mahmoud, and M. Al-Sanabawy, “Thrust termination of solid rocket motor,” in *Journal of Physics: Conference Series*, vol. 2299, p. 012018, IOP Publishing, 2022.
- [43] L. d’Agostino, L. Biagioni, and G. Lamberti, “An ignition transient model for solid propellant rocket motors,” in *37th Joint Propulsion Conference and Exhibit*, p. 3449, 2001.
- [44] F. Barato, N. Bellomo, M. Faenza, M. Lazzarin, A. Bettella, and D. Pavarin, “Numerical model to analyze transient behavior and instabilities on hybrid rocket motors,” *Journal of Propulsion and Power*, vol. 31, no. 2, pp. 643–653, 2015.
- [45] D. Greatrix, “Acceleration-based combustion augmentation modelling for noncylindrical grain solid rocket motors,” in *31st Joint Propulsion Conference and Exhibit*, p. 2876, 1995.
- [46] G. Püskülcü and A. Ulas, “3-d grain burnback analysis of solid propellant rocket motors: Part 1—ballistic motor tests,” *Aerospace Science and Technology*, vol. 12, no. 8, pp. 579–584, 2008.
- [47] L. Ma, K. Wang, Z. Xu, Z. Shao, and Z. Song, “Trajectory optimization for powered descent and landing of reusable rockets with restartable engines,” in *Proceedings of the 69th International Astronautical Congress*, 2018.

- [48] Z. Cen, T. Smith, P. Stewart, and J. Stewart, "Integrated flight/thrust vectoring control for jet-powered unmanned aerial vehicles with a cheone propulsion," *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, vol. 229, no. 6, pp. 1057–1075, 2015.
- [49] R. Kumar, M. Bhargavapuri, A. M. Deshpande, S. Sridhar, K. Cohen, and M. Kumar, "Quaternion feedback based autonomous control of a quadcopter uav with thrust vectoring rotors," in *2020 american control conference (acc)*, pp. 3828–3833, IEEE, 2020.
- [50] C. Papachristos, K. Alexis, and A. Tzes, "Efficient force exertion for aerial robotic manipulation: Exploiting the thrust-vectoring authority of a tri-tiltrotor uav," in *2014 IEEE international conference on robotics and automation (ICRA)*, pp. 4500–4505, IEEE, 2014.
- [51] U. Jain, H. Shukla, S. Kapoor, A. Pandey, and H. Nirwal, "Design and analysis of 2-axis rocket motor stand for thrust vectoring," in *AIAA propulsion and energy 2020 forum*, p. 3920, 2020.
- [52] A. G. Kamath, *Robust thrust vector control for precision rocket-landing*. University of California, Davis, 2021.
- [53] A. G. Valdenebro, "Visualizing rotations and composition of rotations with the rodrigues vector," *European Journal of Physics*, vol. 37, no. 6, p. 065001, 2016.
- [54] W. Khan and M. Nahon, "Real-time modeling of agile fixed-wing uav aerodynamics," in *2015 International Conference on Unmanned Aircraft Systems (ICUAS)*, pp. 1188–1195, 2015.
- [55] B. W. McCormick, *Aerodynamics, aeronautics, and flight mechanics*. John Wiley & Sons, 1994.
- [56] W. Khan and M. Nahon, "Real-time modeling of agile fixed-wing uav aerodynamics," in *2015 international conference on unmanned aircraft systems (ICUAS)*, pp. 1188–1195, IEEE, 2015.
- [57] R. Iyengar, J. Zhu, and B. P. Tripp, "Simulator for the oreo robotic head," *Journal of Computational Vision and Imaging Systems*, vol. 7, no. 1, pp. 13–15, 2021.
- [58] I. Akinola, Z. Wang, and P. Allen, "Clamgen: Closed-loop arm motion generation via multi-view vision-based rl," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2376–2382, IEEE, 2021.
- [59] J. Lee, M. Lee, S. Cho, and S. Lee, "Reference-based video super-resolution using multi-camera video triplets," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17824–17833, 2022.
- [60] S. K. Phang, C. Cai, B. M. Chen, and T. H. Lee, "Design and mathematical modeling of a 4-standard-propeller (4sp) quadrotor," in *Proceedings of the 10th world congress on intelligent control and automation*, pp. 3270–3275, IEEE, 2012.
- [61] B. Landry *et al.*, *Planning and control for quadrotor flight through cluttered environments*. PhD thesis, Massachusetts Institute of Technology, 2015.
- [62] C. Luis and J. L. Ny, "Design of a trajectory tracking controller for a nanoquadcopter," *arXiv preprint arXiv:1608.05786*, 2016.
- [63] J. Förster, "System identification of the crazyflie 2.0 nano quadcopter," B.S. thesis, ETH Zurich, 2015.
- [64] W. WASHINGTON and M. MILLER, "Grid fins-a new concept for missile stability and control," in *31st Aerospace Sciences Meeting*, p. 35, 1993.
- [65] H. Lee, W. Cho, S. Ko, and Y. Lee, "Aerodynamic characteristics of the grid fins on spacex falcon 9," *Journal of the Korean Society for Aeronautical & Space Sciences*, vol. 48, no. 10, pp. 745–752, 2020.
- [66] J. J. Tai, S. K. Phang, and C. L. Hoo, "Application of steady-state integral proportional integral controller for inner dynamics control loop of multi-rotor uavs," in *2018 Fourth International Conference on Advances in Computing, Communication & Automation (ICACCA)*, pp. 1–6, IEEE, 2018.
- [67] J. Eschmann, "Reward function design in reinforcement learning," *Reinforcement Learning Algorithms: Analysis and Applications*, pp. 25–33, 2021.
- [68] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, *et al.*, "Soft actor-critic algorithms and applications," *arXiv preprint arXiv:1812.05905*, 2018.
- [69] J. J. Tai, J. K. Terry, M. S. Innocente, J. Brusey, and N. Horri, "Some supervision required: Incorporating oracle policies in reinforcement learning via epistemic uncertainty metrics," 2022.
- [70] A. Nair, A. Gupta, M. Dalal, and S. Levine, "Awac: Accelerating online reinforcement learning with offline datasets," *arXiv preprint arXiv:2006.09359*, 2020.
- [71] R. Agarwal, M. Schwarzer, P. S. Castro, A. Courville, and M. G. Bellemare, "Deep reinforcement learning at the edge of the statistical precipice," *Advances in Neural Information Processing Systems*, 2021.
- [72] J. Terry, B. Black, N. Grammel, M. Jayakumar, A. Hari, R. Sullivan, L. S. Santos, C. Dieffendahl, C. Horsch, R. Perez-Vicente, *et al.*, "Pettingzoo: Gym for multi-agent reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 34, pp. 15032–15043, 2021.

## VII. BIOGRAPHY SECTION

**Jun Jet Tai** is a PhD student at Coventry University working in the field of reinforcement learning and unmanned aerial robotics. Previously, he obtained his undergraduate's degree in Mechanical Engineering from Taylor's University. While there, he worked as a research assistant under the Taylor's Unmanned Aerial Vehicles research group under Dr. Swee King Phang, where he designed novel control and obstacle avoidance algorithms for quadrotor UAVs.

**Jim Wong** is a final year Aerospace Engineering student at Coventry University. His final year research project focuses on Path Planning and Path Following control algorithms for a fixed wing UAV, with emphasis on optimal control and flight dynamics modelling. He is also interested in the practical applications of artificial intelligence and machine learning in his research in the development of autonomous unmanned aerial vehicles.

**Mauro S. Innocente** is an Assistant Professor in Aerospace Engineering, a member of the Centre for Future Transport and Cities, and the founder of the Autonomous Vehicles & Artificial Intelligence Laboratory at Coventry University. He received his engineering degree from the National University of the Northeast (2003), his MSc degree in numerical methods from the Polytechnic University of Catalonia (2007), and his PhD degree from Swansea University (2011). His research focuses on mathematical modelling, computational intelligence, swarm intelligence, swarm robotics, reinforcement learning, and optimisation within a comprehensive spectrum of applications.

**Nadjim Horri** received his PhD degree from the University of Surrey in 2004. He is currently a lecturer in aerospace control at the University of Leicester. He worked as a research fellow at the Surrey Space Centre and as an assistant professor in aerospace engineering at Coventry University. His research interests include the state and parameter estimation and control of manned and unmanned aircraft and spacecraft, including fault tolerance and the handling of flight constraints and perturbations.

**James Brusey** (M'13) received the B.Ap.Sc. and Ph.D. degrees from RMIT University in 1996 and 2003, respectively. He is currently a Professor of Computer Science and co-Director of the Centre for Computer Science and Mathematical Methods at Coventry University, leading on AI for Cyberphysical Systems. His research interests include Reinforcement Learning applied to real-world, human-centric control systems.

**Swee King Phang** (SM'22) received the B.Eng. and Ph.D. degrees from the National University of Singapore, Singapore, in 2010 and 2014. He is currently a Senior Lecturer of School of Engineering, and hub leader for Automation and Robotics Hub at Taylor's University, Malaysia. His research interests are mainly on unmanned aerial systems, including modeling and flight controller design, indoor navigation and trajectory optimization.