Reinforcement Learning for UAV Attitude Control

WILLIAM KOCH, RENATO MANCUSO, RICHARD WEST, and AZER BESTAVROS,
Boston University, USA

Autopilot systems are typically composed of an “inner loop” providing stability and control, whereas an “outer loop” is responsible for mission-level objectives, such as way-point navigation. Autopilot systems for unmanned aerial vehicles are predominately implemented using Proportional-Integral-Derivative (PID) control systems, which have demonstrated exceptional performance in stable environments. However, more sophisticated control is required to operate in unpredictable and harsh environments. Intelligent flight control systems is an active area of research addressing limitations of PID control most recently through the use of reinforcement learning (RL), which has had success in other applications, such as robotics. Yet previous work has focused primarily on using RL at the mission-level controller. In this work, we investigate the performance and accuracy of the inner control loop providing attitude control when using intelligent flight control systems trained with state-of-the-art RL algorithms—Deep Deterministic Policy Gradient, Trust Region Policy Optimization, and Proximal Policy Optimization. To investigate these unknowns, we first developed an open source high-fidelity simulation environment to train a flight controller attitude control of a quadrotor through RL. We then used our environment to compare their performance to that of a PID controller to identify if using RL is appropriate in high-precision, time-critical flight control.

CCS Concepts: • **Computing methodologies** → **Reinforcement learning**; *Control methods*; *Machine learning*; • **Computer systems organization** → **Embedded systems**;

Additional Key Words and Phrases: Attitude control, UAV, reinforcement learning, quadcopter, autopilot, machine learning, PID, intelligent control, adaptive control

ACM Reference format:

William Koch, Renato Mancuso, Richard West, and Azer Bestavros. 2019. Reinforcement Learning for UAV Attitude Control. *ACM Trans. Cyber-Phys. Syst.* 3, 2, Article 22 (February 2019), 21 pages.
<https://doi.org/10.1145/3301273>

1 INTRODUCTION

Over the past decade, there has been an uptrend in the popularity of Unmanned Aerial Vehicles (UAVs). In particular, quadrotors have received significant attention in the research community, where a significant number of seminal results and applications has been proposed and experimented. This recent growth is primarily attributed to the drop in cost of onboard sensors, actuators, and small-scale embedded computing platforms. Despite the significant progress, flight control is still considered an open research topic. On the one hand, flight control inherently implies

This work was partially supported by a grant from the National Science Foundation under awards #1430145, #1414119, and #1718135.

Authors' addresses: W. Koch, R. Mancuso, R. West, and A. Bestavros, Department of Computer Science, Boston University, 111 Cummington Mall, Boston, MA 02215; emails: {wfkoch, rmancuso, richwest, best}@bu.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2378-962X/2019/02-ART22 \$15.00

<https://doi.org/10.1145/3301273>

the ability to perform highly time-sensitive sensory data acquisition, processing, and computation of forces to apply to the aircraft actuators. On the other hand, it is desirable that UAV flight controllers are able to tolerate faults, adapt to changes in the payload and/or the environment, and optimize flight trajectory, to name a few.

Autopilot systems for UAVs are typically composed of an “inner loop” responsible for aircraft stabilization and control, and an “outer loop” to provide mission-level objectives (e.g., way-point navigation). Flight control systems for UAVs are predominately implemented using Proportional-Integral-Derivative (PID) control systems. PIDs have demonstrated exceptional performance in many circumstances, including in the context of drone racing, where precision and agility are key. In stable environments, a PID controller exhibits close to ideal performance. When exposed to unknown dynamics (e.g., wind, variable payloads, voltage sag), however, a PID controller can be far from optimal [29]. For next-generation flight control systems to be intelligent, a way needs to be devised to incorporate adaptability to mutable dynamics and environment.

The development of intelligent flight control systems is an active area of research [32], specifically through the use of artificial neural networks, which are an attractive option given that they are universal approximators and resistant to noise [30].

Online learning methods (e.g., [14]) have the advantage of learning the aircraft dynamics in real time. The main limitation with online learning is that the flight control system is only knowledgeable of its past experiences. It follows that its performances are limited when exposed to a new event. Training models offline using *supervised learning* is problematic, as data is expensive to obtain and derived from inaccurate representations of the underlying aircraft dynamics (e.g., flight data from a similar aircraft using PID control), which can lead to suboptimal control policies [9, 35, 40]. To construct high-performance intelligent flight control systems, it is necessary to use a hybrid approach. First, accurate offline models are used to construct a baseline controller, whereas online learning provides fine tuning and real-time adaptation.

An alternative to supervised learning for creating offline models is known as *reinforcement learning* (RL). In RL, an agent is given a reward for every action it takes in an environment, with the objective to maximize the rewards over time. Using RL, it is possible to develop optimal control policies for a UAV without making any assumptions about the aircraft dynamics. Recent work has shown RL to be effective for UAV autopilots, providing adequate path tracking [18]. Nonetheless, previous work on intelligent flight control systems has primarily focused on guidance and navigation.

Open challenges in RL for attitude control. RL is currently being applied to a wide range of applications, each with its own set of challenges. Attitude control for UAVs is a particularly interesting RL problem for several reasons. We have highlighted three areas we find important:

- C1: *Precision and Accuracy.*** Many RL tasks can be solved in a variety of ways. For example, to win a game, there may be several sequential moves that will lead to the same outcome. In the case of optimal attitude control, there is little tolerance and flexibility as to the sequence of control signals that will achieve the desired attitude (e.g., angular rate) of the aircraft. Even the slightest deviations can lead to instabilities. It remains unclear what level of control accuracy can be achieved when using intelligent control trained with RL for time-sensitive attitude control (i.e., the “inner loop”). Therefore, determining the achievable level of accuracy is critical in establishing if RL is suitable for attitude flight control.
- C2: *Robustness and Adaptation.*** In the context of control, robustness refers to the controller’s performance in the presence of uncertainty when control parameters are fixed, whereas adaptiveness refers to the controller’s performance to adapt to the uncertainties by adjusting the control parameters [37]. It is assumed that the neural network trained with

RL will face uncertainties when transferred to physical hardware due to the gap between the RL environment and the real world. However, it remains unknown in what range of uncertainty the controller can operate safely before adaptation is necessary. Characterizing the controller's robustness will provide valuable insight into the design of the intelligent flight control system architecture. For instance, what will be the necessary adaptation rate, and what sensor data can be collected from the real world to update the RL environment?

C3: *Reward Engineering*: In RL, reward engineering is the process of designing a reward system to provide the agent a signal showing that they are doing the right thing [12]. In the context of attitude control, the reward must encapsulate the agent's performance in achieving the desired attitude goals. As goals become more complex and demanding (e.g., minimizing energy consumption or stability in presence of damage), identifying which performance metrics are most expressive will be necessary to push the performance of intelligent control systems trained with RL.

Our contributions. In this article, we study C1 in depth the accuracy and precision of attitude control provided by intelligent flight controllers trained using RL. Although we specifically focus on the creation of controllers for the Iris quadcopter [4], the methods developed apply to a wide range of multirotor UAVs and can also be extended to fixed-wing aircraft. We develop a novel training environment called GymFC with the use of a high-fidelity physics simulator for the agent to learn attitude control. GymFC is an OpenAI Environment [11] providing a common interface for researchers to develop intelligent flight control systems. The simulated environment consists of an Iris quadcopter digital replica or *digital twin* [16] with the intention of eventually be used to transfer the trained controller to physical hardware. Controllers are trained using state-of-the-art RL algorithms: Deep Deterministic Policy Gradient (DDPG), Trust Region Policy Optimization (TRPO), and Proximal Policy Optimization (PPO). We then compare the performance of our synthesized controllers with that of a PID controller. Our evaluation finds that controllers trained using PPO outperform PID control and are capable of exceptional performance. To summarize, this article makes the following contributions:

- GymFC, an open source [22] environment for developing an intelligent attitude flight controller, providing the research community a tool to progress performance.
- A learning architecture for attitude control utilizing digital twinning concepts for minimal effort when transferring trained controllers into hardware.
- An evaluation for state-of-the-art RL algorithms, such as DDPG, TRPO, and PPO, learning policies for aircraft attitude control. As a first work in this direction, our evaluation also establishes a baseline for future work.
- An analysis of intelligent flight control performance developed with RL compared to traditional PID control.

The remainder of this article is organized as follows. In Section 2, we provide an overview of the quadcopter flight dynamics and RL. Next, in Section 3, we briefly survey existing literature on intelligent flight control. In Section 4, we present our training environment and then use this environment to evaluate RL performance for flight control in Section 5. Finally, Section 6 concludes the article and provides several future research directions.

2 BACKGROUND

In this section, we provide an overview of quadcopter flight dynamics required to understand this work and an introduction to developing flight control systems with RL.

Fig. 1. Quadcopter rotational movement.

2.1 Quadcopter Flight Dynamics

A quadcopter is an aircraft with six degrees of freedom (DOF), three rotational and three translational. With four control inputs (one to each motor), this results in an underactuated system that requires an onboard computer to compute motor signals to provide stable flight. We indicate with $\omega_i, i \in 1, \dots, M$ the rotation speed of each rotor where $M = 4$ is the total number of motors for a quadcopter. These have a direct impact on the resulting Euler angles ϕ, θ, ψ , for instance, roll, pitch, and yaw, respectively, which provide rotation in $D = 3$ dimensions. Moreover, they produce a certain amount of upward thrust, indicated with f .

The aerodynamic effect that each ω_i produces depends on the configuration of the motors. The most popular configuration is an “X” configuration, depicted in Figure 1(a), which has the motors mounted in an “X” formation relative to what is considered the front of the aircraft. This configuration provides more stability compared to a “+” configuration, which in contrast has its motor configuration rotated an additional 45° along the z -axis. This is due to the differences in torque generated along each axis of rotation in respect to the distance of the motor from the axis. The aerodynamic affect u that each rotor speed ω_i has on thrust and Euler angles is given by

$$u_f = b(\omega_1^2 + \omega_2^2 + \omega_3^2 + \omega_4^2) \quad (1)$$

$$u_\phi = b(\omega_1^2 + \omega_2^2 - \omega_3^2 - \omega_4^2) \quad (2)$$

$$u_\theta = b(\omega_1^2 - \omega_2^2 + \omega_3^2 - \omega_4^2) \quad (3)$$

$$u_\psi = b(\omega_1^2 - \omega_2^2 - \omega_3^2 + \omega_4^2). \quad (4)$$

where $u_f, u_\phi, u_\theta, u_\psi$ is the thrust, roll, pitch, and yaw effect, respectively, whereas b is a thrust factor that captures propeller geometry and frame characteristics. For further details about the mathematical models of quadcopter dynamics, please refer to Bouabdallah et al. [10].

To perform a rotational movement, the velocity of each rotor is manipulated according to the relationship expressed in Equation (4) and as illustrated in Figure 1(b) through (d). For example, to roll right (Figure 1(b)), more thrust is delivered to motors 3 and 4. Yaw (Figure 1(d)) is not achieved directly through difference in thrust generated by the rotor as roll and pitch are, but

Fig. 2. RL architecture using the GymFC environment for training intelligent attitude flight controllers.

instead through a difference in torque in the rotation speed of rotors spinning in opposite directions. For example, as shown in Figure 1(d), higher rotation speed for rotors 1 and 4 allow the aircraft to yaw clockwise. Because a net positive torque counterclockwise causes the aircraft to rotate clockwise due to Newton's second law of motion.

Attitude, in respect to orientation of a quadcopter, can be expressed by its angular velocities of each axis $\Omega = [\Omega_\phi, \Omega_\theta, \Omega_\psi]$. The objective of attitude control is to compute the required motor signals to achieve some desired attitude Ω^* .

In autopilot systems, attitude control is executed as an inner control loop and is time sensitive. Once the desired attitude is achieved, translational movement (in the X, Y, Z direction) is accomplished by applying thrust proportional to each motor.

In commercially available quadcopters, the vast, if not all, use PID attitude control. A PID controller is a linear feedback controller expressed mathematically as

$$u(t) = K_p e(t) + K_i \int_0^t e(\tau) d\tau + K_d \frac{de(t)}{dt}, \quad (5)$$

where K_p, K_i, K_d are configurable constant gains and $u(t)$ is the control signal. The effect of each term can be thought of as the P term considers the current error, the I term considers the history of errors, and the D term estimates the future error. For attitude control in a quadcopter aircraft, there is PID control for each roll, pitch, and yaw axis. At each cycle in the inner loop, each PID sum is computed for each axis and then these values are translated into the amount of power to deliver to each motor through a process called *mixing*. Mixing uses a table consisting of constants describing the geometry of the frame to determine how the axis control signals are summed based on the torques that will be generated by the length of each arm (recall differences between “X” and “+” frames). The control signal for each motor y_i is loosely defined as

$$y_i = f(m_{(i,\phi)}u_\phi + m_{(i,\theta)}u_\theta + m_{(i,\psi)}u_\psi), \quad (6)$$

where $m_{(i,\phi)}, m_{(i,\theta)}, m_{(i,\psi)}$ are the mixer values for motor i and f is the throttle coefficient.

2.2 Reinforcement Learning

In this work, we consider an RL architecture (depicted in Figure 2) consisting of a neural network flight controller as an agent interacting with an Iris quadcopter [4] in a high-fidelity physics simulated environment $\mathcal{E}$, more specifically using the Gazebo simulator [23].

At each discrete timestep t , the agent receives an observation x_t from the environments consisting of the angular velocity error of each axis $e = \Omega^* - \Omega$ and the angular velocity of each rotor ω_i , which are obtained from the quadcopter's inertial measurement unit (IMU) and electronic speed controller (ESC) sensors, respectively. These observations are in the continuous observation spaces $x_t \in \mathbb{R}^{(M+D)}$. Once the observation is received, the agent executes an action a_t within $\mathcal{E}$. In return, the agent receives a single numerical reward r_t indicating the performance of this action. The action is also in a continuous action space $a_t \in \mathbb{R}^M$ and corresponds to the four control signals $u(t)$ sent to each ESC driving the attached motor $\mathbb{M}$. Because the agent is only receiving this sensor data, it is unaware of the physical environment and the aircraft dynamics, and therefore $\mathcal{E}$ is only partially observed by the agent. Motivated by Minh et al. [31], we consider the state to be a sequence of the past observations and actions $s_t = x_i, a_i, \dots, a_{t-1}, x_t$.

The interaction between the agent and $\mathcal{E}$ is formally defined as a Markov decision process (MDP) where the state transitions are defined as the probability of transitioning to state s' given the current state and action respectively are s and a , $Pr\{s_{t+1} = s' | s_t = s, a_t = a\}$. The behavior of the agent is defined by its policy π , which is essentially a mapping of what action should be taken for a particular state. The objective of the agent is to maximize the returned reward overtime to develop an optimal policy. We welcome the reader to refer to Sutton and Barto [36] for further details on RL.

Until recently, control in a continuous action space was considered difficult for RL. Significant progress has been made combining the power of neural networks with RL. In this work, we elected to use DDPG [27] and TRPO [33] due to the recent use of these algorithms for quadcopter navigation control [18]. DDPG provides improvement to the Deep Q-Network (DQN) [31] for the continuous action domain. It employs an actor-critic architecture using two neural networks for each actor and critic. It is also a model-free algorithm, meaning that it can learn the policy without having to first generate a model. TRPO is similar to natural gradient policy methods; however, this method guarantees monotonic improvements. We additionally include a third algorithm for our analysis: PPO [34]. PPO is known to outperform other state-of-the-art methods in challenging environments. PPO is also a policy gradient method and has similarities to TRPO while being easier to implement and tune.

3 RELATED WORK

Aviation has a rich history in flight control dating back to the 1960s. During this time, supersonic aircraft were being developed, which demanded more sophisticated dynamic flight control than what a static linear controller could provide. Gain scheduling [26] was developed, allowing multiple linear controllers of different configurations to be used in designated operating regions. This, however, was inflexible and insufficient for handling the nonlinear dynamics at high speeds but paved the way for adaptive control. For a period of time, many experimental adaptive controllers were being tested but were unstable. Later advances were made to increase stability with model reference adaptive control (MRAC) [39] and L_1 [17], which provided reference models during adaptation. Additionally numerous other nonlinear control algorithms have been developed and applied to flight control, including feedback linearization [25], sliding mode control [25], and backstepping [28], in an attempt to model the nonlinear dynamics of the aircraft.

As the cost of small-scale embedded computing platforms dropped, intelligent flight control options became realistic and have been actively researched over the past decade to design flight control solutions that are able to *adapt* and also *learn*. The ability to learn is an important distinction that differentiates intelligent control from the control algorithms discussed previously. Intelligent control architectures are capable of planning for future events such as system failure,

damage, and emergencies—tasks that would otherwise be difficult or impossible for other control algorithms [24].

As performance demands for UAVs continue to increase, we are beginning to see signs of flight control history repeating itself. The popular high-performance drone racing firmware Betaflight [2] has recently added a gain scheduler to adjust PID gains depending on throttle and voltage levels. Intelligent PID flight control [15] methods have been proposed in which PID gains are dynamically updated online, providing adaptive control as the environment changes. However, these solutions still inherit disadvantages associated with PID control, such as integral windup, need for mixing, and, most significantly, they are feedback controllers and therefore inherently *reactive*. Yet feedforward control (or predictive control) is *proactive* and allows the controller to output control signals before an error occurs. For feedforward control, a model of the system must exist. Learning-based intelligent control has been proposed to develop models of the aircraft for predictive control using artificial neural networks.

Notable work by Dierks and Jagannathan [14] proposes an intelligent flight control system constructed with neural networks to learn the quadcopter dynamics, online, to navigate along a specified path. This method allows the aircraft to adapt in real time to external disturbances and unmodeled dynamics. Matlab simulations demonstrate that their approach outperforms a PID controller in the presence of unknown dynamics, specifically in regard to control effort required to track the desired trajectory. Nonetheless, the proposed approach does require prior knowledge of the aircraft mass and moments of inertia to estimate velocities. Online learning is an essential component to constructing a complete intelligent flight control system. It is fundamental, however, to develop accurate offline models to account for uncertainties encountered during online learning [32]. To build offline models, previous work has used supervised learning to train intelligent flight control systems using a variety of data sources, such as test trajectories [9], and PID step responses [35]. The limitation of this approach is that training data may not accurately reflect the underlying dynamics. In general, supervised learning on its own is not ideal for interactive problems such as control [36].

RL has similar goals to adaptive control in which a policy improves over time when interacting with its environment. RL has been applied to autonomous helicopters to learn how to track trajectories (guidance), specifically how to hover in place and perform various maneuvers [6, 7, 21]. Kim et al. [21] and Abbeel et al. [6] demonstrated their trained helicopter's capabilities in helicopter competitions requiring the aircraft to perform advanced aerobatic maneuvers. Performance was compared to trained pilots; nevertheless, it is unknown how their controllers compare to PID control for tracking trajectories. In contrast to their work, we investigate the use and accuracy of using RL for low-level manipulation of the aircraft actuators to maintain a desired attitude, not for guidance control. Furthermore our goal is to compare controllers taught with RL to PID control to determine what applications, if any, would be more appropriate. The first use of RL in quadcopter control was presented by Waslander et al. [38] for altitude control. The authors developed a model-based RL algorithm to search for an optimal control policy. The controller was rewarded for accurate tracking and damping. Their design provided significant improvements in stabilization in comparison to a linear control system. More recently, Hwangbo et al. [18] used RL for quadcopter control, particularly for navigation control. They developed a novel deterministic on-policy learning algorithm that outperformed TRPO [33] and DDPG [27] in regard to training time. Furthermore, the authors validated their results in the real world, transferring their simulated model to a physical quadcopter. Path tracking turned out to be adequate. Notably, the authors discovered major differences in transferring from simulation to the real world. Known as the *reality gap*, transferring from simulation to the real world has been researched extensively as being problematic without taking additional steps to increase realism in the simulator [19, 30].

Most prior work has focused on performance of navigation and guidance. There is limited and insufficient data justifying the accuracy and precision of neural network-based intelligent attitude flight control and none to our knowledge for controllers trained using RL. Furthermore, this work uses physics simulations in contrast to mathematical models of the aircraft and environments used in aforementioned prior work for increased realism. The goal of this work is to provide a platform for training attitude controllers with RL, and to provide performance baselines in regard to attitude controller accuracy.

4 ENVIRONMENT

In this section, we describe our learning environment, GYMFC, for developing intelligent flight control systems using RL. The goal of the proposed environment is to allow the agent to learn attitude control of an aircraft with only the knowledge of the number of actuators. GYMFC includes both an *episodic task* and a *continuous task*. In an episodic task, the agent is required to learn a policy for responding to individual angular velocity commands. This allows the agents to learn the step response from rest for a given command, allowing its performance to be accurately measured. Episodic tasks, however, are not reflective of realistic flight conditions. For this reason, in a continuous task, pulses with random widths and amplitudes are continuously generated and correspond to angular velocity setpoints. The agent must respond accordingly and track the desired target over time. In Section 5, we evaluate our synthesized controllers via episodic tasks, but we have strong experimental evidence that training via episodic tasks produces controllers that behave correctly in continuous tasks as well (Appendix A).

GYMFC has a multilayer hierarchical architecture composed of three layers: (i) a digital twin layer, (ii) a communication layer, and (iii) an agent-environment interface layer. This design decision was made to clearly establish roles and allow layer implementations to change (e.g., to use a different simulator) without affecting other layers as long as the layer-to-layer interfaces remain intact. A high-level overview of the environment architecture is illustrated in Figure 3. We will now discuss in greater detail each layer with a bottom-up approach.

4.1 Digital Twin Layer

At the heart of the learning environment is a high-fidelity physics simulator that provides functionality and realism that is hard to achieve with an abstract mathematical model of the aircraft and environment. One of the primary design goals of GYMFC is to minimize the effort required to transfer a controller from the learning environment into the final platform. For this reason, the simulated environment exposes identical interfaces to actuators and sensors as they would exist in the physical world. In the ideal case, the agent should not be able to distinguish between interaction with the simulated world (i.e., its digital twin) and its hardware counterpart. In this work, we use the Gazebo simulator [23] in light of its maturity, flexibility, extensive documentation, and active community.

In a nutshell, the *digital twin layer* is defined by (i) the simulated world and (ii) its interfaces to the above communication layer (see Figure 3).

Simulated world. The simulated world is constructed specifically for UAV attitude control in mind. The technique we developed allows attitude control to be accomplished independently of guidance and/or navigation control. This is achieved by fixing the center of mass of the aircraft to a ball joint in the world, allowing it to rotate freely in any direction, which would be impractical, if not impossible, to achieve in the real world due to gimbal lock and friction of such an apparatus. In this work, the aircraft to be controlled in the environment is modeled off of the Iris quadcopter [4] with a weight of 1.5Kg and 550mm motor-to-motor distance. An illustration of the quadcopter in the environment is displayed in Figure 4. Note during training that Gazebo runs in headless mode

Fig. 3. Overview of environment architecture, GymFC. Blue blocks with dashed borders are implementations developed for this work.

Fig. 4. The Iris quadcopter in Gazebo 1m above the ground. The body is transparent to show where the center of mass is linked as a ball joint to the world. Arrows represent the various joints used in the model.

without this user interface to increase simulation speed. This architecture, however, can be used with any multicopter as long as a digital twin can be constructed. Helicopters and multicopters represent excellent candidates for our setup because they can achieve a full range of rotations along all three axes. This is typically not the case with fixed-wing aircraft. Our design can, however, be expanded to support fixed-wing aircraft by simulating airflow over the control surfaces for attitude control. Gazebo already integrates a set of tools to perform airflow simulation.

Interface. The digital twin layer provides two command interfaces to the communication layer: simulation reset and motor update. Simulation reset commands are supported by Gazebo's API and are not part of our implementation. Motor updates are provided by a UDP server. We hereby discuss our approach to developing this interface.

To keep synchronicity between the simulated world and the controller of the digital twin, the pace at which the simulation should progress is directly enforced. This is possible by controlling

the simulator step by step. In our initial approach, Gazebo's Google Protobuf [5] API was used, with a specific message to progress by a single simulation step. By subscribing to status messages (which include the current simulation step), it is possible to determine when a step has completed and to ensure synchronization. However, as we attempted to increase the rate of advertising step messages, we discovered that the rate of status messages is capped at 5Hz. Such a limitation introduces a consistent bottleneck in the simulation/learning pipeline. Furthermore, it was found that Gazebo silently drops messages that it cannot process.

A set of important modifications were made to increase experiment throughput. The key idea was to allow motor update commands to directly drive the simulation clock. By default, Gazebo comes preinstalled with an ArduPilot ArduCopter [1] plugin to receive motor updates through a UDP server. These motor updates are in the form of pulse width modulation (PWM) signals. At the same time, sensor readings from the IMU on board the aircraft is sent over a second UDP channel. ArduCopter is an open source multicopter firmware, and its plugin was developed to support software in the loop (SITL).

We derived our Aircraft Plugin from the ArduCopter plugin with the following modifications (as well as those discussed in Section 4.2). On receiving a motor command, the motor forces are updated as normal but then a simulation step is executed. Sensor data is read and then sent back as a response to the client over the same UDP channel. In addition to the IMU sensor data, we also simulate sensor data obtained from the ESC. The ESC provides the angular velocities of each rotor, which are relayed to the client as well. Implementing our Aircraft Plugin with this approach successfully allowed us to work around the limitations of the Google Protobuf API and increased step throughput by more than 200 times.

4.2 Communication Layer

The communication layer is positioned in between the digital twin and the agent-environment interface. This layer manages the low-level communication channel to the aircraft and simulation control. The primary function of this layer is to export a high-level synchronized API to the higher layers for interacting with the digital twin that uses asynchronous communication protocols. This layer provides the commands `pwm_write` and `reset` to the agent-environment interface layer.

The function call `pwm_write` takes as input a vector of PWM values for each actuator, corresponding to the control input $u(t)$. These PWM values correspond to the same values that would be sent to an ESC on a physical UAV. The PWM values are translated to a normalized format expected by the Aircraft Plugin and then packed into a UDP packet for transmission to the Aircraft Plugin UDP server. The communication layer blocks until a response is received from the Aircraft Plugin, forcing synchronized writes for the above layers. The UDP reply is unpacked and returned in response.

During the learning process, the simulated environment must be reset at the beginning of each learning episode. Ideally, one could use the `gz` command line utility included with the Gazebo installation, which is lightweight and does not require additional dependencies. Unfortunately, there is a known socket handle leak [3] that causes Gazebo to crash if the command is issued more than the maximum number of open files allowed by the operating system. Given that we are running thousands episodes during training, this was not an option for us. Instead, we opted to use the Google Protobuf interface, so we did not have to deploy a patched version of the utility on our test servers. Because resets only occur at the beginning of a training session and are not in the critical processing loop using Google Protobufs here is acceptable.

On start of the communication layer, a connection is established with the Google Protobuf API server, and we subscribe to world statistics messages that include the current simulation iteration. To reset the simulator, a world control message is advertised, instructing the simulator to reset the

simulation time. The communication layer blocks until it receives a world statistics message indicating that the simulator has been reset and then returns back control to the agent-environment interface layer. Note that the world control message is only resetting the simulation time, not the entire simulator (i.e., models and sensors). This is because we found that in some cases when a world control message was issued to perform a full reset, the sensor data took a few additional iterations for reset. To ensure proper reset to the above layers, this time reset message acts as a signaling mechanism to the Aircraft Plugin. When the plugin detects that a time reset has occurred, it resets the whole simulator and, most importantly, steps the simulator until the sensor values have also reset ensuring above layers that when a new training session starts, reading sensor values accurately reflect the current state and not the previous state from stale values.

4.3 Environment Interface Layer

The topmost layer interfacing with the agent is the environment interface layer that implements the OpenAI Gym [11] environment API. Each OpenAI Gym environment defines an observation space and an action space. These are used to inform the agent of the bounds to expect for environment observations and what are legal bounds for the action input, respectively. As mentioned in Section 2.2, GymFC is in both the continuous observation space and action space domain. The state is of size $m \times (M + D)$, where m is the memory size indicating the number of past observations, $M = 4$ as we consider a four-motors configuration, and $D = 3$ since each measurement is taken in the three dimensions. Each observation value is in $[-\infty : \infty]$. The action space is of size M equivalent to the number of control actuators of the aircraft (i.e., four for a quadcopter), where each value is normalized between $[-1 : 1]$ to be compatible with most agents who squash their output using the tanh function.

GymFC implements two primary OpenAI functions, namely reset and step. The reset function is called at the start of an episode to reset the environment and returns the initial environment state. This is also when the desired target angular velocity Ω^* or setpoint is computed. The setpoint is randomly sampled from a uniform distribution between $[\Omega_{min}, \Omega_{max}]$. For the continuous task, this is also set at a random interval of time. Selection of these bounds may refer to the desired operating region of the aircraft. Although it is highly unlikely during normal operation that a quadcopter will be expected to reach the majority of these target angular velocities, the intention of these tasks are to push and stress the performance of the aircraft.

The step function executes a single simulation step with the specified actions and returns to the agent the new state vector, together with a reward indicating how well the given action was performed. Reward engineering can be challenging. If careful design is not performed, the derived policy may not reflect what was originally intended. Recall from Section 2.2 that the reward is ultimately what shapes the policy. For this work, with the goal of establishing a baseline of accuracy, we develop a reward to reflect the current angular velocity error (i.e., $e = \Omega^* - \Omega$). In the future, GymFC will be expanded to include additional environments aiding in the development of more complex policies' particularly to showcase the advantages of using RL to adapt and learn. We translate the current error e_t at time t into a derived reward r_t normalized between $[-1, 0]$ as follows:

$$r_t = -clip(sum(|\Omega_t^* - \Omega_t|)/3\Omega_{max}, 0, 1), \quad (7)$$

where the *sum* function sums the absolute value of the error of each axis, and the *clip* function clips the result between $[0, 1]$ in cases where there is an overflow in the error. Since the reward is negative, it signifies a penalty, and the agent maximizes the rewards (thus minimizing error) over time to track the target as accurately as possible. Rewards are normalized to provide standardization and stabilization during training [20].

Additionally, we experimented with a variety of other rewards. We found sparse binary rewards¹ to give poor performance. We believe that this is due to complexity of quadcopter control and the limitations of the RL algorithms we tested. In the early stages of learning, the agent explores its environment. However, the event of randomly reaching the target angular velocity within some threshold was rare and thus did not provide the agent with enough information to converge. Conversely, we found that signaling at each timestep was best.

5 EVALUATION

In this section, we present our evaluation on the accuracy of studied neural network-based attitude flight controllers trained with RL. Due to space limitations, we present evaluation and results only for episodic tasks, as they are directly comparable to our baseline (PID). Nonetheless, we have obtained strong experimental evidence that agents trained using episodic tasks perform well in continuous tasks (Appendix A). To our knowledge, this is the first RL baseline conducted for quadcopter attitude control.

5.1 Setup

We evaluate the RL algorithms DDPG, TRPO, and PPO using the implementations in the OpenAI Baselines project [13]. The goal of the OpenAI Baselines project is to establish a reference implementation of RL algorithms, providing baselines for researchers to compare approaches and build on. Every algorithm is run with defaults except for the number of simulations steps, which we increased to 10 million.

The episodic task parameters were configured to run each episode for a maximum of 1 second of simulated time, allowing enough time for the controller to respond to the command and additional time to identify if a steady state has been reached. The bounds the target angular velocity is sampled from is set to $\Omega_{min} = -5.24$ rad/s, $\Omega_{max} = 5.24$ rad/s (± 300 deg/s). These limits were constructed by examining PID's performance to make sure we expressed physically feasible constraints. The max step size of the Gazebo simulator, which specifies that the duration of each physics update step was set to 1 ms to develop highly accurate simulations. In other words, our physical world "evolved" at 1kHz. Training and evaluations were run on Ubuntu 16.04 with an eight-core i7-7700 CPU and an NVIDIA GeForce GT 730 graphics card.

For our PID controller, we ported the mixing and SITL implementation from Betaflight [2] to Python to be compatible with GymFC. The PID controller was first tuned using the classical Ziegler-Nichols method [41] and then manually adjusted to improve performance of the step response sampled around the midpoint $\pm \Omega_{max}/2$. We obtained the following gains for each axis of rotation: $K_\phi = [2, 10, 0.005]$, $K_\theta = [10, 10, 0.005]$, $K_\psi = [4, 50, 0.0]$, where each vector contains the $[K_p, K_i, K_d]$ (proportional, integrative, derivative) gains, respectively. Next, we measured the distances between the arms of the quadcopter to calculate the mixer values for each motor m_i , $i \in \{1, \dots, 4\}$. Each vector m_i is of the form $m_i = [m_{(i,\phi)}, m_{(i,\theta)}, m_{(i,\psi)}]$ —for instance, roll, pitch, and yaw (see Section 2.1). The final values were $m_1 = [-1.0, 0.598, -1.0]$, $m_2 = [-0.927, -0.598, 1.0]$, $m_3 = [1.0, 0.598, 1.0]$, and last, $m_4 = [0.927, -0.598, -1.0]$. The mix values and PID sums are then used to compute each motor signal y_i according to Equation (6), where $f = 1$ for no additional throttle.

To evaluate and compare the accuracy of the different algorithms, we used a set of metrics. First, we define "initial error" as the distance between the rest velocities and the current setpoint. A notion of progress toward the setpoint from rest can then be expressed as the percentage of the initial error that has been "corrected." Correcting 0% of the initial error means that no progress

¹A reward structured so that $r_t = 0$ if $\text{sum}(|e_t|) < \text{threshold}$; otherwise, $r_t = -1$.

Fig. 5. Average normalized rewards shown in magenta received during training of 10,000 episodes (10 million steps) for each RL algorithm and memory m sizes 1, 2, and 3. Plots share a common y and x axis. Additionally, yellow represents the 95% confidence interval and the black line is a 2-degree polynomial added to illustrate the trend of the rewards over time.

has been made, whereas 100% indicates that the setpoint has been reached. Each metric value is independently computed for each axis. We hereby list our metrics. *Success* captures the number of experiments (in percentage) in which the controller eventually settles in a band within 90% and 110% of the initial error (i.e., $\pm 10\%$ from the setpoint). *Failure* captures the average percentage error relative to the initial error after $t = 500\text{ms}$ for those experiments that do not make it in the $\pm 10\%$ error band. The latter metric quantifies the magnitude of unacceptable controller performance. The delay in the measurement ($t > 500\text{ms}$) is to exclude the rise regime. The underlying assumption is that a steady state is reached before 500ms. *Rise* is the average time in milliseconds it takes the controller to go from 10% to 90% of the initial error. *Peak* is the max achieved angular velocity represented as a percentage relative to the initial error. Values greater than 100% indicate overshoot, whereas values less than 100% represent undershoot. *Error* is the average sum of the absolute value error of each episode in radians per second. This provides a generic metric for performance. Our last metric is *Stability*, which captures how stable the response is halfway through the simulation (i.e., at $t > 500\text{ms}$). Stability is calculated by taking the linear regression of the angular velocities and reporting the slope of the calculated line. Systems that are unstable have a nonzero slope.

5.2 Results

Each learning agent was trained with an RL algorithm for a total of 10 million simulation steps, equivalent to 10,000 episodes or about 2.7 simulation hours. The agents configuration is defined as the RL algorithm used for training and its memory size m . Training for DDPG took approximately 33 hours, whereas PPO and TRPO took approximately 9 hours and 13 hours, respectively. The average sum of rewards for each episode is normalized between $[-1, 0]$ and displayed in Figure 5. This computed average in magenta is from three independently trained agents with the

Table 1. RL Performance Evaluation Averages From 3,000 Command Inputs
Per Configuration With 95% Confidence

m	Rise (ms)				Peak (%)			Error (rad/s)			Stability		
	ϕ	θ	ψ		ϕ	θ	ψ	ϕ	θ	ψ	ϕ	θ	ψ
P	1	65.9±2.4	94.1±4.3	73.4±2.7	113.8±2.2	107.7±2.2	128.1±4.3	309.9±7.9	440.6±13.4	215.7±6.7	0.0±0.0	0.0±0.0	0.0±0.0
	2	58.6±2.5	125.4±6.0	105.0±5.0	116.9±2.5	103.0±2.7	126.8±3.7	305.2±7.9	674.5±19.1	261.3±7.6	0.0±0.0	0.0±0.0	0.0±0.0
	3	101.5±5.0	128.8±5.8	79.2±3.3	108.9±2.2	94.2±5.3	119.8±2.7	405.9±10.9	1,403.8±58.4	274.4±5.3	0.0±0.0	0.0±0.0	0.0±0.0
T	1	103.9±6.2	150.2±6.7	109.7±8.0	125.1±9.3	110.4±3.9	139.6±6.8	1,644.5±52.1	929.0±25.6	1,374.3±51.5	-0.4±0.1	-0.2±0.0	-0.1±0.0
	2	161.3±6.9	162.7±7.0	108.4±9.6	100.1±5.1	144.2±13.8	101.7±5.4	1,432.9±47.5	2,375.6±84.0	1,475.6±46.4	0.1±0.0	0.4±0.0	-0.1±0.0
	3	130.4±7.1	150.8±7.8	129.1±8.9	141.3±7.2	141.2±8.1	147.1±6.8	1,120.1±36.4	1,200.7±34.3	824.0±30.1	0.1±0.0	-0.1±0.1	-0.1±0.0
D	1	68.2±3.7	100.0±5.4	79.0±5.4	133.1±7.8	116.6±7.9	146.4±7.5	1,201.4±42.4	1,397.0±62.4	992.9±45.1	0.0±0.0	-0.1±0.0	0.1±0.0
	2	49.2±1.5	99.1±4.9	40.7±1.8	42.0±5.5	46.7±8.0	71.4±7.0	2,388.0±63.9	2,607.5±72.2	1,953.4±58.3	-0.1±0.0	-0.1±0.0	-0.0±0.0
	3	85.3±5.9	124.3±7.2	105.1±8.6	101.0±8.2	158.6±21.0	120.5±7.0	1,984.3±59.3	3,280.8±98.7	1,364.2±54.9	0.0±0.1	0.2±0.1	0.0±0.0

Note: P = PPO, T = TRPO, and D = DDPG.

same configuration, whereas the 95% confidence is shown in yellow. Additionally, we have added a 2-degree polynomial in black fit to the data to illustrate the reward trend over time. Training results show clearly that PPO converges consistently compared to TRPO and DDPG, and overall PPO accumulates higher rewards. What is also interesting and counterintuitive is that the larger memory size actually *decreases* convergence and stability among all trained algorithms. Recall from Section 2 that RL algorithms learn a policy to map states to action. A reason for the decrease in convergence could be attributed to the state space increasing, causing the RL algorithm to take longer to learn the mapping to the optimal action. As part of our future work, we plan to investigate using separate memory sizes for the error and rotor velocity to decrease the state space. Reward gains during training of TRPO and DDPG are quite inconsistent with large confidence intervals. Although performance for DDPG $m = 1$ looks promising, on further investigation into the large confidence interval, we found that this was due to the algorithm completely failing to respond to certain command inputs, thus questioning whether the algorithm has learned the underlying flight dynamics (this is emphasized later in Table 2).

In the future, we plan to investigate methods to decrease training times by addressing C2 and C3 in Section 1. Specific to C2 to support a large range of aircraft, we will explore whether we can construct a generic neural network taught general flight dynamics (Section 2.1), which will provide a baseline to extend training to create intelligent controllers unique to an aircraft (otherwise known as domain adaptation [8]). Additionally, considering C3, we will experiment with developing more expressive reward functions to decrease training times.

Each trained agent was then evaluated on 1,000 never before seen command inputs in an episodic task. Since there are three agents per configuration, each configuration was evaluated over a total of 3,000 episodes. The average performance metrics for Rise, Peak, Error, and Stability for the response to the 3,000 command inputs is reported in Table 1. Results show that the agent trained with PPO outperforms TRPO and DDPG in every measurement. In fact, PPO is the only one that is able to achieve stability (for every m), whereas all other agents have at least one axis where the Stability metric is nonzero.

Next, the best-performing agent for each algorithm and memory size is compared to the PID controller. The best agent was selected based on the lowest sum of errors of all three axes reported by the Error metric. The Success and Failure metrics are compared in Table 2. Results show that agents trained with PPO would be the only ones good enough for flight, with a success rate close to perfect, and where the roll failure of 0.2% is only off by about 0.1% from the setpoint. However, the

Table 2. Success and Failure Results for Considered Algorithms

	m	Success (%)			Failure (%)		
		ϕ	θ	ψ	ϕ	θ	ψ
P	1	99.8±0.3	100.0±0.0	100.0±0.0	0.1±0.1	0.0±0.0	0.0±0.0
	2	100.0±0.0	53.3±3.1	99.8±0.3	0.0±0.0	20.0±2.4	0.0±0.0
	3	98.7±0.7	74.7±2.7	99.3±0.5	0.4±0.2	5.4±0.7	0.2±0.2
T	1	32.8±2.9	59.0±3.0	87.4±2.1	72.5±10.6	17.4±3.7	9.4±2.6
	2	19.7±2.5	48.2±3.1	56.9±3.1	76.6±5.0	43.0±6.5	38.6±7.0
	3	96.8±1.1	60.8±3.0	73.2±2.7	1.5±0.8	20.6±4.1	20.6±3.4
D	1	84.1±2.3	52.5±3.1	90.4±1.8	11.1±2.2	41.1±5.5	4.6±1.0
	2	26.6±2.7	26.1±2.7	50.2±3.1	82.7±8.5	112.2±12.9	59.7±7.5
	3	39.2±3.0	44.8±3.1	60.7±3.0	52.0±6.4	101.8±13.0	33.9±3.4
PID		100.0±0.0	100.0±0.0	100.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0

Note: P = PPO, T = TRPO, and D = DDPG. The row highlighted in blue refers to our best-performing learning agent PPO, whereas the rows with gray highlight correspond to the best agents for the other two algorithms.

Table 3. RL Performance Evaluation Compared to PID of the Best-Performing Agent

m		Rise (ms)			Peak (%)			Error (rad/s)			Stability		
		ϕ	θ	ψ	ϕ	θ	ψ	ϕ	θ	ψ	ϕ	θ	ψ
P	1	66.6±3.2	70.8±3.6	72.9±3.7	112.6±3.0	109.4±2.4	127.0±6.2	317.0±11.0	326.3±13.2	217.5±9.1	0.0±0.0	0.0±0.0	0.0±0.0
	2	64.4±3.6	102.8±6.7	148.2±7.9	118.4±4.3	104.2±4.7	124.2±3.4	329.4±12.3	815.3±31.4	320.6±11.5	0.0±0.0	0.0±0.0	0.0±0.0
	3	97.9±5.5	121.9±7.2	79.5±3.7	111.4±3.4	111.1±4.2	120.8±4.2	396.7±14.7	540.6±22.6	237.1±8.0	0.0±0.0	0.0±0.0	0.0±0.0
T	1	119.9±8.8	149.0±10.6	103.9±9.8	103.0±11.0	117.4±5.8	142.8±6.5	1965.2±90.5	930.5±38.4	713.7±34.4	0.7±0.1	0.3±0.0	0.0±0.0
	2	108.0±8.3	157.1±9.9	47.3±6.5	69.4±7.4	117.7±9.2	126.5±7.2	2020.2±71.9	1316.2±49.0	964.0±31.2	0.1±0.1	0.5±0.1	0.0±0.0
	3	115.2±9.5	156.6±12.7	176.1±15.5	153.5±8.1	123.3±6.9	148.8±11.2	643.5±20.5	895.0±42.8	1108.9±44.5	0.1±0.0	0.0±0.0	0.0±0.0
D	1	64.7±5.2	118.9±8.5	51.0±4.8	165.6±11.6	135.4±12.8	150.8±6.2	929.1±39.9	1490.3±83.0	485.3±25.4	0.1±0.1	-0.2±0.1	0.1±0.0
	2	49.2±2.1	99.1±6.9	40.7±2.5	84.0±10.4	93.5±15.4	142.7±12.5	2074.1±86.4	2498.8±109.8	1336.9±50.1	-0.1±0.0	-0.2±0.1	-0.0±0.0
	3	73.7±8.4	172.9±12.0	141.5±14.5	103.7±11.5	126.5±17.8	119.6±8.2	1585.4±81.4	2401.3±109.8	1199.0±74.0	-0.1±0.1	-0.2±0.1	0.1±0.0
PID		79.0±3.5	99.8±5.0	67.7±2.3	136.9±4.8	112.7±1.6	135.1±3.3	416.1±20.4	269.6±11.9	245.1±11.5	0.0±0.0	0.0±0.0	0.0±0.0

Note: Values reported are the average of 1,000 command inputs with 95% confidence where P = PPO, T = TRPO, and D = DDPG. PPO $m = 1$ (highlighted in blue) outperforms all other agents, including PID control. Metrics highlighted in red for PID control are outperformed by the PPO agent.

best-trained agents for TRPO and DDPG are often significantly far away from the desired angular velocity. For example, TRPO's best agent, 39.2% (60.8% success, see Table 2) of the time, does not reach the desired pitch target with upward of a 20% error from the setpoint.

Next, we provide our thorough analysis comparing the best agents in Table 3. We have found that RL agents trained with PPO using $m = 1$ provide performance and accuracy exceeding that of our PID controller in regard to rise time, peak velocities achieved, and total error. What is interesting is that usually a fast rise time could cause overshoot; however, the PPO agent has, on average, a faster rise time and less overshoot. Both PPO and PID reach a stable state measured halfway through the simulation.

To illustrate the performance of each of the best agents, a random simulation is sampled and the step response for each attitude command is displayed in Figure 6, along with the target angular velocity to achieve Ω^* . All algorithms reach some steady state; however, only PPO and PID

Fig. 6. Step response of best-trained RL agents compared to PID. Target angular velocity is $\Omega^* = [2.20, -5.14, -1.81]$ rad/s, shown by a dashed black line. Error bars $\pm 10\%$ of the initial error from Ω^* are shown in a dashed red line.

Fig. 7. Step response and PWM motor signals in microseconds (μs) of the best-trained PPO agent compared to PID. Target angular velocity is $\Omega^* = [2.11, -1.26, 5.00]$ rad/s, shown by a dashed black line. Error bars $\pm 10\%$ of initial error from Ω^* are shown in a dashed red line.

do so within the error band indicated by the dashed red lines. TRPO and DDPG have extreme oscillations in both the roll and yaw axes, which would cause instability during flight. In this particular example, we can observe PID to perform better with a 19% decrease in error compared to PPO, most visibly in yaw control. However, globally speaking, in terms of error, PPO has shown to be a more accurate attitude controller.

To highlight the performance and accuracy of the PPO agent, we sample another simulation and show the step response and also the PWM control signals generated by each controller in Figure 7. In this figure, we can see that the PPO agent has exceptional tracking capabilities of the

desired attitude. Compared to PID, the PPO controller has a 44% decrease in error. The PPO agent has a 2.25 times faster rise time on the roll axis, is 2.5 times faster on the pitch axis, and is 1.15 times faster on the yaw axis. Furthermore, the PID controller experiences slight overshoot in both the roll and yaw axes, whereas the PPO agent does not. In regard to the control output, the PID controller exerts more power to motor 3 but then motor values eventually level off, whereas the PPO control signal oscillates comparably more.

6 FUTURE WORK AND CONCLUSION

In this article, we presented our RL training environment, GymFC, for developing intelligent attitude controllers for UAVs and addressed C1: Precision and Accuracy in depth, which identifies if neural networks trained with RL can produce accurate attitude controllers. We placed an emphasis on digital twinning concepts to allow transferability to real hardware. We used GymFC to evaluate the performance of state-of-the-art RL algorithms PPO, TRPO, and DDPG to identify if they are appropriate to synthesize high-precision attitude flight controllers. Our results highlight that (i) RL can train accurate attitude controllers and (ii) that those trained with PPO outperformed a fully tuned PID controller on almost every metric. Although we base our evaluation on results obtained in episodic tasks, we found that trained agents were able to perform exceptionally well also in continuous tasks without retraining (Appendix A). This suggests that training using episodic tasks is sufficient for developing intelligent attitude controllers. The results presented in this work can be considered as a first milestone and a good motivation to further inspect the boundaries of RL for intelligent control. With this premise, we plan to develop our future work along three main avenues. On the one hand, we plan to investigate C2: Robustness and Adaptation and C3: Reward Engineering (Section 1) to harness the true power of RL's ability to adapt and learn in environments with dynamic properties (e.g., wind, variable payload, system damage and failure). On the other hand, we intend to transfer our trained agents onto a real aircraft to evaluate their live performance, including timing and memory analysis of the neural network. This will allow us to define the minimum hardware specifications required to use neural network attitude control. Furthermore, we plan to expand GymFC to support other aircraft, such as fixed-wing aircraft, while continuing to increase the realism of the simulated environment by improving the accuracy of our digital twins.

APPENDIX

A CONTINUOUS TASK EVALUATION

In this section, we briefly expand on our findings that show that even if agents are trained through episodic tasks, their performance transfers to continuous tasks without the need for additional training. Figure 8 shows that an agent trained with Proximal Policy Optimization (PPO) using episodic tasks has exceptional performance when evaluated in a continuous task. Figure 9 is a closeup of another continuous task sample showing the details of the tracking and corresponding motor output. These results are quite remarkable, as they suggest that training with episodic tasks is sufficient for developing intelligent attitude flight controller systems capable of operating in a continuous environment. In Figure 10, another continuous task is sampled and the PPO agent is compared to a PID agent. The performance evaluation shows the PPO agent to have a 22% decrease in overall error compared to the PID agent.

Fig. 8. Performance of a PPO agent trained with episodic tasks but evaluated using a continuous task for a duration of 60 seconds. The time in seconds at which a new command is issued is randomly sampled from the interval $[0.1, 1]$, and each issued command is maintained for a random duration also sampled from $[0.1, 1]$. Desired angular velocity is specified by the black line, whereas the red line is the attitude tracked by the agent.

Fig. 9. Closeup of continuous task results for a PPO agent with PWM values.

Fig. 10. Response comparison of a PID and PPO agent evaluated in a continuous task environment. The PPO agent, however, is only trained using episodic tasks.

ACKNOWLEDGMENT

We would like to thank the anonymous reviewers for their comments which helped us improve the quality of this manuscript.

REFERENCES

- [1] ArduPilot. 2018. ArduPilot Home Page. Retrieved January 20, 2019 from <http://ardupilot.org/>.
- [2] GitHub. 2018. BetaFlight. Retrieved January 20, 2019 from <https://github.com/betaflight/betaflight>.
- [3] Open Source Robotics Foundation. 2018. gzserver doesn't close disconnected sockets. Retrieved January 20, 2019 from <https://bitbucket.org/osrf/gazebo/issues/2397/gzserver-doesnt-close-disconnected-sockets>.
- [4] APM Copter. 2018. Iris QuadCopter. Retrieved January 20, 2019 from <http://www.ardupilot.co.uk/iris-quadcopter-uav.html>.
- [5] Google. 2018. Protocol Buffers. Retrieved January 20, 2019 from <https://developers.google.com/protocol-buffers/>.
- [6] Pieter Abbeel, Adam Coates, Morgan Quigley, and Andrew Y. Ng. 2007. An application of reinforcement learning to aerobatic helicopter flight. In *Advances in Neural Information Processing Systems*. 1–8.
- [7] J. Andrew Bagnell and Jeff G. Schneider. 2001. Autonomous helicopter control using reinforcement learning policy search methods. In *Proceedings of the 2001 IEEE International Conference on Robotics and Automation (ICRA'01)*, Vol. 2. IEEE, Los Alamitos, CA, 1615–1620.
- [8] John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman. 2008. Learning bounds for domain adaptation. In *Advances in Neural Information Processing Systems*. 129–136.
- [9] Alexey Bobtsov, Alexei Guirik, Marina Budko, and Mikhail Budko. 2016. Hybrid parallel neuro-controller for multirotor unmanned aerial vehicle. In *Proceedings of the 2016 8th International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT'16)*. IEEE, Los Alamitos, CA, 1–4.
- [10] Samir Bouabdallah, Pierpaolo Murrieri, and Roland Siegwart. 2004. Design and control of an indoor micro quadrotor. In *Proceedings of the 2004 IEEE International Conference on Robotics and Automation (ICRA'04)*, Vol. 5. IEEE, Los Alamitos, CA, 4393–4398.
- [11] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, et al. 2016. Openai gym. arXiv:1606.01540.
- [12] Daniel Dewey. 2014. Reinforcement learning and the reward engineering principle. In *Proceedings of the 2014 AAAI Spring Symposium Series*.
- [13] Prafulla Dhariwal, Christopher Hesse, Oleg Klimov, Alex Nichol, Matthias Plappert, Alec Radford, et al. 2017. OpenAI Baselines. *GitHub*. Retrieved January 20, 2019 from <https://github.com/openai/baselines>.

- [14] Travis Dierks and Sarangapani Jagannathan. 2010. Output feedback control of a quadrotor UAV using neural networks. *IEEE Transactions on Neural Networks* 21, 1 (2010), 50–66.
- [15] Mehdi Fatan, Bahram Lavi Sefidgari, and Ali Vatankeh Barenji. 2013. An adaptive neuro PID for controlling the altitude of quadcopter robot. In *Proceedings of the 2013 18th International Conference on Methods and Models in Automation and Robotics (MMAR'13)*. IEEE, Los Alamitos, CA, 662–665.
- [16] Thomas Gabor, Lenz Belzner, Marie Kiermeier, Michael Till Beck, and Alexander Neitz. 2016. A simulation-based architecture for smart cyber-physical systems. In *Proceedings of the 2016 IEEE International Conference on Autonomic Computing (ICAC'16)*. IEEE, Los Alamitos, CA, 374–379.
- [17] Naira Hovakimyan, Chengyu Cao, Evgeny Kharisov, Enric Xargay, and Irene M. Gregory. 2011. L1 adaptive control for safety-critical systems. *IEEE Control Systems* 31, 5 (2011), 54–104.
- [18] Jemin Hwangbo, Inkyu Sa, Roland Siegwart, and Marco Hutter. 2017. Control of a quadrotor with reinforcement learning. *IEEE Robotics and Automation Letters* 2, 4 (2017), 2096–2103.
- [19] Nick Jakobi, Phil Husbands, and Inman Harvey. 1995. Noise and the reality gap: The use of simulation in evolutionary robotics. In *Proceedings of the 3rd European Conference on Advances in Artificial Life*. 704–720.
- [20] Andrej Karpathy. 2018. Deep Reinforcement Learning: Pong from Pixels. Retrieved January 20, 2019 from <http://karpathy.github.io/2016/05/31/rl/>.
- [21] H. Jin Kim, Michael I. Jordan, Shankar Sastry, and Andrew Y. Ng. 2004. Autonomous helicopter flight via reinforcement learning. In *Advances in Neural Information Processing Systems*. 799–806.
- [22] William Koch. 2018. GymFC. *GitHub*. Retrieved January 20, 2019 from <https://github.com/wil3/gymfc>.
- [23] Nathan Koenig and Andrew Howard. [n.d.]. Design and use paradigms for Gazebo, an open-source multi-robot simulator. In *Proceedings of the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'04)*, Vol. 3. IEEE, Los Alamitos, CA, 2149–2154.
- [24] Kalmanje KrishnaKumar and Karen Gundy-Burlet. 2002. Intelligent control approaches for aircraft applications. In *Proceedings of the JANAFF Interagency Propulsion Committee Meeting*. Destin, FL.
- [25] Daewon Lee, H. Jin Kim, and Shankar Sastry. 2009. Feedback linearization vs. adaptive sliding mode control for a quadrotor helicopter. *International Journal of Control, Automation and Systems* 7, 3 (2009), 419–428.
- [26] Douglas J. Leith and William E. Leithead. 2000. Survey of gain-scheduling analysis and design. *International Journal of Control* 73, 11 (2000), 1001–1025.
- [27] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, et al. 2015. Continuous control with deep reinforcement learning. arXiv:1509.02971.
- [28] Tarek Madani and Abdelaziz Benallegue. 2006. Backstepping control for a quadrotor helicopter. In *Proceedings of the 2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 3255–3260.
- [29] K. Niki Maleki, Kaveh Ashenayi, Loyd R. Hook, Justin G. Fuller, and Nathan Hutchins. 2016. A reliable system design for nondeterministic adaptive controllers in small UAV autopilots. In *Proceedings of the 2016 IEEE/AIAA 35th Digital Avionics Systems Conference (DASC'16)*. IEEE, Los Alamitos, CA, 1–5.
- [30] Orazio Miglino, Henrik Hautop Lund, and Stefano Nolfi. 1995. Evolving mobile robots in simulated and real environments. *Artificial Life* 2, 4 (1995), 417–434.
- [31] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, et al. 2013. Playing Atari with deep reinforcement learning. arXiv:1312.5602.
- [32] Fendy Santoso, Matthew A. Garratt, and Sreenatha G. Anavatti. 2017. State-of-the-art intelligent flight control systems in unmanned aerial vehicles. *IEEE Transactions on Automation Science and Engineering* 15, 2, 613–627.
- [33] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. 2015. Trust region policy optimization. In *Proceedings of the International Conference on Machine Learning*. 1889–1897.
- [34] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. arXiv:1707.06347.
- [35] Jack F. Shepherd III and Kagan Tumer. 2010. Robust neuro-control for a micro quadrotor. In *Proceedings of the 12th Annual Conference on Genetic and Evolutionary Computation*. ACM, New York, NY, 1131–1138.
- [36] Richard S. Sutton and Andrew G. Barto. 1998. *Reinforcement Learning: An Introduction*. Vol. 1. MIT Press, Cambridge, MA.
- [37] Le Yi Wang and Ji-Feng Zhang. 2001. Fundamental limitations and differences of robust and adaptive control. In *Proceedings of the 2001 American Control Conference*, Vol. 6. IEEE, Los Alamitos, CA, 4802–4807.
- [38] Steven Lake Waslander, Gabriel M. Hoffmann, Jung Soon Jang, and Claire J. Tomlin. 2005. Multi-agent quadrotor testbed control design: Integral sliding mode vs. reinforcement learning. In *Proceedings of the 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'05)*. IEEE, Los Alamitos, CA, 3712–3717.
- [39] H. Philip Whitaker, Joseph Yamron, and Allen Kezer. 1958. *Design of Model-Reference Adaptive Control Systems for Aircraft*. MIT Instrumentation Laboratory, Cambridge, MA.

- [40] Peggy S. Williams-Hayes. 2005. Flight test implementation of a second generation intelligent flight control system. In *Proceedings of the Infotech@Aerospace Conference*.
- [41] John G. Ziegler and Nathaniel B. Nichols. 1942. Optimum settings for automatic controllers. *Transactions of the ASME* 64, 11 (Nov. 1942), 759–768.

Received May 2018; revised September 2018; accepted December 2018