Air Learning: a deep reinforcement learning gym for autonomous aerial robot visual navigation

Srivatsan Krishnan¹ · Behzad Boroujerdian² · William Fu¹ · Aleksandra Faust³ · Vijay Janapa Reddi^{1,2}

Received: 16 March 2020 / Revised: 2 January 2021 / Accepted: 21 May 2021 /
Published online: 7 July 2021
© The Author(s) 2021

Abstract

We introduce Air Learning, an open-source simulator, and a gym environment for deep reinforcement learning research on resource-constrained aerial robots. Equipped with domain randomization, Air Learning exposes a UAV agent to a diverse set of challenging scenarios. We seed the toolset with point-to-point obstacle avoidance tasks in three different environments and Deep Q Networks (DQN) and Proximal Policy Optimization (PPO) trainers. Air Learning assesses the policies' performance under various quality-of-flight (QoF) metrics, such as the energy consumed, endurance, and the average trajectory length, on resource-constrained embedded platforms like a Raspberry Pi. We find that the trajectories on an embedded Ras-Pi are vastly different from those predicted on a high-end desktop system, resulting in up to 40% longer trajectories in one of the environments. To understand the source of such discrepancies, we use Air Learning to artificially degrade high-end desktop performance to mimic what happens on a low-end embedded system. We then propose a mitigation technique that uses the hardware-in-the-loop to determine the latency distribution of running the policy on the target platform (onboard compute on aerial robot). A randomly sampled latency from the latency distribution is then added as an artificial delay within the training loop. Training the policy with artificial delays allows us to minimize the hardware gap (discrepancy in the flight time metric reduced from 37.73% to 0.5%). Thus, Air Learning with hardware-in-the-loop characterizes those differences and exposes how the onboard compute's choice affects the aerial robot's performance. We also conduct reliability studies to assess the effect of sensor failures on the learned policies. All put together, Air Learning enables a broad class of deep RL research on UAVs. The source code is available at: <https://github.com/harvard-edge/AirLearning>.

Keywords Deep reinforcement learning · Autonomous aerial robots · Resource-constrained deep RL · Robotics · Deep RL challenges · Sim2Real · Real life RL

Editors: Yuxi Li, Alborz Geramifard, Lihong Li, Csaba Szepesvari, Tao Wang.

✉ Srivatsan Krishnan
srivatsan@seas.harvard.edu

Extended author information available on the last page of the article

1 Introduction

Deep Reinforcement Learning (RL) has shown promising results in domains like sensori-motor control for cars (Bojarski et al., 2016), indoor robots (Chiang et al., 2019), as well as UAVs (Gandhi et al., 2017; Sadeghi & Levine, 2016). Deep RL's ability to adapt and learn with minimum apriori knowledge makes them attractive for use in complex systems (Kretchmar, 2000).

Unmanned Aerial Vehicles (UAVs) serve as a great platform for advancing state of the art for deep RL research. UAVs have practical applications, such as search and rescue (Waharte & Trigoni, 2010), package delivery (Faust et al., 2017; Goodchild & Toy, 2018), construction inspection (Peng et al., 2017). Compared to other robots such as self-driving cars, robotic arm, they are vastly cheap to prototype and build, which makes them truly scalable.¹ Also, UAVs have fairly diverse control requirements. Targeting low-level UAV control (e.g. attitude control) requires continuous control (e.g., angular velocities) whereas, targeting high-level tasks such as point-to-point navigation can use discrete control. Last but not least, at deployment time they must be a fully autonomous system, running onboard computationally- and energy-constrained computing hardware.

But despite the promise of deep RL, there are several practical challenges in adopting reinforcement learning for the UAV navigation task as shown in Fig. 1. Broadly, the problems can be grouped into four main categories: (1) environment simulator, (2) learning algorithms, (3) policy architecture, and (4) deployment on resource-constrained UAVs. To address these challenges, the boundaries between reinforcement learning algorithms, robotics control, and the underlying hardware must soften. The figure illustrates the cross-layer, and interdisciplinary nature of the field, spanning from environment modeling to the underlying system. Each layer, in isolation, has a complex design space that needs to be explored for optimization. In addition, there are interactions across the layers that are also important to consider (e.g., policy size on a power-constrained mobile or embedded computing system). Hence, there is a need for a platform that can aid interdisciplinary research. More specifically, we need a research platform that can benchmark each of the layers individually (for depth), as well as end-to-end execution for capturing the interactions across the layers (for breadth).

In this paper, we present Air Learning (Sect. 4)—an open source deep RL research simulation suite and benchmark for autonomous UAVs. As a simulation suite of tools, Air Learning provides a scalable and cost-effective applied reinforcement learning research. It augments existing frameworks such as AirSim (Shah et al., 2017) with capabilities that make it suitable for deep RL experimentation. *As a gym, Air Learning enables RL research for resource constrained systems.*

Air Learning addresses the simulator level challenge, by providing domain randomization. We develop a configurable environment generator with a range of knobs to generate different environments with varying difficulty levels. The knobs (randomly) tune the number of static and dynamic obstacles, their speed (if relevant), texture and color, arena size, etc. In the context of our benchmarking autonomous UAV navigation task, the knobs help the learning algorithms generalize well without overfitting to a specific instance of an environment.²

¹ For example, an FPV hobbyist drone can be built under \$100: <https://bit.ly/2TR3rMQ>.

² The environment generator can be applied to other challenges in aerial robots, such as detecting thin wires and coping with translucent objects.

Fig. 1 Aerial robotics is a cross-layer, interdisciplinary field. Designing an autonomous aerial robot to perform a task involves interactions between various boundaries, spanning from environment modeling down to the choice of hardware for the onboard compute

Air Learning addresses the learning challenges (RL algorithm, policy design, and reward optimization) by exposing the environment generator as an OpenAI gym (Brockman et al., 2016) interface and integrating it with Baselines (Hill et al., 2018), which has high-quality implementations of the state-of-the-art RL algorithms. We provide templates which the researchers can use for building multi-modal input policies based on Keras/Tensorflow. And as a DRL benchmark, the OpenAI gym interface enables easy addition of new deep RL algorithms. At the point of writing this paper, we provide two different reinforcement learning algorithms Deep Q-Networks (DQN) (Mnih et al., 2013) and Proximal Policy Optimization (PPO) (Schulman et al., 2017). DQN is an off-policy, discrete action RL algorithm, and PPO is an on-policy, continuous action control of UAVs. Both come ready with curriculum learning (Bengio et al., 2009) support.

To address the resource-constrained challenge early on in the design and development of deep RL algorithms and policies, Air Learning uses a “hardware-in-the-loop” (HIL) (Adiprawita et al., 2008) method to enable robust hardware evaluation without risking real UAV platform. Hardware in the loop, which requires plugging in the computing platform used in the UAV into the software simulation, is a form of real-time simulation that allows us to understand how the UAV responds to simulated stimuli on a target hardware platform.³ HIL simulation helps us quantify the real-time performance of reinforcement learning policies on various compute platforms, without risking experiments on real robot platforms before they are ready.

We use HIL simulation to understand how a policy performs on an embedded compute platform that might potentially be the onboard computer of the UAVs. To enable systematic HIL evaluation, we use a variety of Quality-of-Flight (QoF) metrics, such as the total energy consumed by the UAV, the average length of its trajectory and endurance, to compare the different reinforcement learning policies. To demonstrate that Air Learning’s HIL simulation is essential and that it can reveal interesting insights, we take the best performing policy from our policy exploration stage and evaluate the performance of the policy on a resource-constrained low-performance platform (Ras-Pi 4) and compare it with a high-performance desktop counterpart (Intel Core-i9). The difference between the Ras-Pi 4 and the Core-i9 based performance for the policy is startling. The Ras-Pi 4 sometimes takes trajectories that are nearly 40% longer in some environments. We investigate the reason for the difference in the performance of the policy on Ras-Pi 4 versus Intel Core-i9 and show that the choice of onboard compute platform directly affects the policy processing latency,

³ Demonstration of HIL: <https://bit.ly/2NDRjex>.

and hence the trajectory lengths. The discrepancy in the policy behavior from training to deployment hardware is a challenge that must be taken into account when designing the DRL algorithm for a resource-constrained robot. We define this behavior as ‘Hardware induced gap’ because of the performance gap in training machine versus deployment machine. We use a variety of metrics to quantify the hardware gap, such as percentage change between the QoF metrics that include flight time, success rate, the energy of flight, and trajectory distance.

In summary, we present an open-source gym environment and research platform for deep RL research for autonomous aerial vehicles. The contributions within this context include:

- We present an open source benchmark to develop and train different RL algorithms, policies, and reward optimizations using regular and curriculum learning.
- We present a UAV mapless navigation task benchmark for RL research on resource constrained systems.
- We present a random environment generator for domain randomization to enable RL generalization.
- We introduce and show ‘Hardware induced gap’ – that the policy’s behavior depends on a computing platform it is running on, and that the same policy can result in a very different behavior if the target deployment platform is very different from the training platform.
- We describe the significance of taking energy consumption and the platform’s processing capabilities into account when evaluating policy success rates.
- To alleviate the hardware-induced gap, we train a policy using HIL to match the target platform’s latencies. Using this mitigation technique, we minimized the hardware gap between the training platform and resource-constrained target platform from 38% to less than 0.5% on flight time, 16.03% to 1.37% on the trajectory length, and 15.49% to 0.1% on the energy of flight metric.

Air Learning will be of interest to both fundamental and applied RL research community. The point to point UAV navigation benchmark can yield to progress on fundamental RL algorithm development for resource-constrained systems where training and deployment platforms are different. From that point of view, Air Learning is another OpenAI Gym environment. For the applied RL researchers, interested in RL applications for UAV domains such as source seeking, search and reuse, etc., Air Learning serves as a simulation platform and toolset for full-stack research and development.

2 Real world challenges

We describe the real-world challenges associated with developing deep RL algorithms on resource-constrained UAVs. We consolidate the challenges into four categories, namely Environment simulator, challenges related to the learning algorithm, policy selection challenges, and hardware-level challenges.

Environment Simulator Challenges: The first challenge is that deep RL algorithms targeted for robotics need simulator. Collecting large amounts of real-world data is challenging because most commercial and off-the-shelf UAVs operate for less than 30 mins. To put this into perspective, creating a dataset as large as the latest “ImageNet” by

Tencent for ML Images (Wu et al., 2019) would take close to 8000 flights (assuming a standard 30 FPS camera), thus making it a logistically challenging issue. But perhaps an even more critical and difficult aspect of this data collection is that there is a need for negative experiences, such as obstacle collisions, which can severely drive up the cost and logistics of collecting data (Gandhi et al., 2017). More importantly, it has been shown the environment simulator having high fidelity and ability to perform domain randomization aids in the better generalization of reinforcement learning algorithms (Tobin et al., 2017). Hence, any infrastructure for deep RL must have features to address these challenges to deploy RL policies in real-world robotics applications.

Learning algorithm challenges: The second challenge is associated with reinforcement learning algorithms. Choosing the right variant of a reinforcement learning algorithm for a given task requires fairly exhaustive exploration. Furthermore, since the performance and efficiency of a particular reinforcement learning algorithm are greatly influenced by its reward function, to get good performance, there is a need to perform design exploration between the reinforcement learning algorithms and its reward function. Though these challenges are innate to the deep RL domain, having an environment simulator exposed as a simple interface (Brockman et al., 2016) can allow us to efficiently automate the RL algorithm selection, rewards shaping, hyperparameter tuning (Chiang et al., 2019).

Policy selection challenges: The third challenge is associated with the selection of policies for robot control. Choosing the right policy architecture is a fairly exhaustive task. Depending upon the available sensor suite on the robot, the policy can be uni-modal or multi-modal in nature. Also, for effective learning, the hyperparameters associated with the policy architecture have to be appropriately tuned. Hyperparameter tuning and policy architecture search is still an active area of research, which has led to techniques such as AutoML (Zoph et al., 2017) to determine the optimal neural network architecture. In the context of DRL policy selection, having a standard machine learning back-end tool such as Tensorflow/Keras (Abadi et al., 2015) can allow DRL researchers (or roboticist) to automate the policy architecture search.

Hardware-level challenges: The fourth challenge is regarding the deployment of deep RL policies on the resource-constrained UAVs. Since UAVs are mobile machines, they need to accomplish their tasks with a limited amount of onboard energy. Because onboard compute is a scarce resource and RL policies are computationally intensive, we need to carefully co-design the policies with the underlying hardware so that the compute platform can meet the real-time requirements under power constraints. As the UAV size decreases, the problem exacerbates because battery capacity (i.e., size) decreases, which reduces the total onboard energy (even though the level of intelligence required remains the same). For instance, a nano-UAV such as a CrazyFlie (2018) must have the same autonomous navigation capabilities as compared to its larger mini counterpart, e.g., DJI-Mavic Pro (2018) while the CrazyFlie's onboard energy is $\frac{1}{15}$ th that of the Mavic Pro. Typically in deep RL research for robotics, the system and onboard computers are based on commercial off the shelf hardware platforms. However, whether the selection of these compute platforms is optimal is mostly unknown. Hence, having the ability to characterize the onboard computing platform early on can lead to resource-friendly deep RL policies.

Air Learning is built with features to overcome the challenges listed above. Due to the interdisciplinary nature of the tool, it provides flexibility to researchers to focus on a given layer (e.g., policy architecture design) while also understanding its impact on the subsequent layer (e.g., hardware performance). In the next section, we describe the related work and list of features that Air Learning supports out of the box.

3 Related work

Related work in deep RL toolset and benchmarks can be divided into three categories. The first category of related work includes environments for designing and benchmarking new deep RL algorithms. The second category of related work includes tools used specifically for deep RL based aerial robots. In the third category, we include other learning-based toolsets that support features that are important for deep RL training. The feature list and comparison of related work to Air Learning are tabulated in Table 1.

Benchmarking environments: The first category of related work includes benchmarking environments such as OpenAI Gym (Brockman et al., 2016), Arcade Learning Environments (Bellemare et al., 2015), and MuJoCo (Todorov et al., 2012). These environments are simple by design and allow designing and benchmarking of new deep RL algorithms. However, using these environments for real-life applications such as robotics is challenging because they do not address the hardware-level challenges (Sect. 2) for transferring trained RL policies to real robots. Air Learning addresses these limitations by introducing Hardware-in-the-Loop (HIL), which allows end-user to benchmark and characterize the RL policy performance on a given onboard computing platform.

UAV specific deep RL benchmarks: The second category of related work includes benchmarks that focus on UAVs. For example, AirSim (Shah et al., 2017) provides a high-fidelity simulation and dynamics for the UAVs in the form of a plugin that can be imported in any UE4 (Unreal Engine 4) (Valcasara, 2015) project. However, there are three AirSim limitations that AirLearning addresses. First, the generation of the environment that includes domain randomization for the UAV task is left to the end-user to either develop or source it from the UE4 market place. The domain randomizations (Tobin et al., 2017) are very critical for generalization of the learning algorithm, and we address this limitation in AirSim using the Air Learning environment generator.

Second, AirSim does not model UAV energy consumption. Energy is a scarce resource in UAVs that affects overall mission capability. Hence, learning algorithms need to be evaluated for energy efficiency. Air Learning uses energy model (Boroujerdian et al., 2018) within AirSim to evaluate learned policies. Air Learning also allows studying the impact of the performance of the onboard compute platform on the overall energy of UAVs, allowing us to estimate in the simulation how many missions UAV can do, without running in the simulation.

Third, AirSim does not offer interfaces with OpenAI gym or other reinforcement learning framework such as stable baselines (Hill et al., 2018). We address this drawback by exposing the Air Learning random environment generator with OpenAI gym interfaces and integrate it with a high-quality implementation of reinforcement learning algorithms available in the framework such as baselines (Hill et al., 2018) and Keras-RL (Plappert, 2016). Using Air Learning, we can quickly explore and evaluate different RL algorithms for various UAV tasks.

Another related work that uses a simulator and OpenAI gym interface in the context of UAVs is GYMFC (Koch et al., 2018). GYMFC uses Gazebo (Koenig & Howard, 2004) simulator and OpenAI gym interfaces for training an attitude controller for UAVs using reinforcement learning. The work primarily focuses on replacing the conventional flight controller with a real-time controller based on a neural network. This is a highly specific, low-level task. We focus more on high-level tasks, such as point-to-point UAV navigation in an environment with static and dynamic obstacles, and we provide the necessary infrastructure to carry research to enable on-edge autonomous navigation in UAVs. Adapting

Table 1 Comparison of features commonly present in deep RL research infrastructures

Features	UAV specific			UAV agnostic			Robot-Grasping	ROBEL (Ahn et al., 2020)	SenseAct (Mahmood et al., 2018)
	Air Learning	AirSim	GymFC	CARLA	Gazebo	PyRobot			
Photorealism	✓	✓	×	✓	×	×	✓	×	×
Domain randomization	✓	×	×	✓	×	✓	✓	×	×
Open-AI gym interface	✓	×	✓	✓	✓	✓	✓	✓	✓
RL Algorithm exploration	✓	✓	✓	✓	✓	✓	✓	✓	✓
ML backend integration	✓	✓	✓	✓	✓	✓	✓	✓	✓
UAV physics	✓	✓	✓	×	✓	×	×	×	×
Energy modelling	✓	×	×	×	×	×	×	×	×
Compute benchmarking	✓	×	×	×	×	×	×	✓	✓
RL policy deployment on robot	✓	✓	✓	✓	✓	✓	✓	✓	✓

✓ Denotes that the feature exists. × Denotes missing feature or requires significant effort from end-user to enable that feature

this work to support a high-level task such as navigation will involve overcoming the limitations of Gazebo, specifically in the context of photorealism. One of the motivations of building AirSim is to overcome the limitations of Gazebo by using state-of-the-art rendering techniques for modeling the environment, which is achieved using robust game engines such as Unreal Engine 4 (Valcasara, 2015) and Unity (Menard & Wagstaff, 2015).

UAV agnostic deep RL benchmarks: The third category of related work includes deep RL benchmarks used for other robot tasks, such as grasping by a robotic arm or self-driving car. These related works are highly relevant to Air Learning because it contains essential features that improve the utility/performance of deep RL algorithms.

The most prominent work in learning-based approaches for self-driving cars is CARLA (Dosovitskiy et al., 2017). It supports a photorealistic environment built on top of a game engine. It also exposes the environment as an OpenAI gym interface, which allows researchers to experiment with different deep RL algorithms. The physics is based on the game engine, and they do not model energy or focus on the compute hardware performance. Since the CARLA was built explicitly for self-driving cars, porting these features to UAVs will require significant engineering effort.

For the robotic arm grasping/manipulation task, prior work (Ahn et al., 2020; Gu et al., 2016; Kalashnikov et al., 2018; Quillen et al., 2018) include infrastructure support to train and deploy deep RL algorithms on these robots. In Yahya et al. (2016), they introduce collective learning where they provide distributed infrastructure to collect large amounts of data with real platform experiments. They introduce an asynchronous variant of guided policy search to maximize the utilization (computer and synchronization between different agents), where each agent trains a local policy while a single global policy is trained based on the data collected from individual agents. However, these kinds of robots are fixed in a place; hence, they are not limited by energy or by onboard compute capability. So the inability to process or calculate the policy's outcome in real-time only slows down the grasping rate. It does not cause instability. In UAVs, which have a higher control loop rate, uncertainty due to slow processing latency can cause fatal crashes (Giusti et al., 2016; Hwangbo et al., 2017).

For mobile robots with/without grasping such as LocoBot (Locobot, 2018), PyRobot (Murali et al., 2019), ROBEL (Ahn et al., 2020) provides open-source tools and benchmarks for training and deploying deep RL policies on the LocoBot. The simulation infrastructure is based on Gazebo or MuJoCo, and hence it lacks photorealism in the environment and other domain randomization features. Similar to CARLA and robot grasping benchmarks, PyRobot does not model energy or focus on computing hardware performance.

In soft learning (Haarnoja et al., 2018), the authors apply a soft-actor critic algorithm for the quadrupedal robot. They use Nvidia TX2 on the robot for data collection and also running the policy. The data collected is then used to train the global policy, which is then periodically updated to the robot. In contrast, in our work, we show that training policy on a high-end machine can result in a discrepancy in performance for aerial robot platform. Aerial robots are much more complex to control and unstable compared to ground-based quadrupedal robots. Hence small differences in processing time can hinder its safety. We propose training a policy using the HIL technique with the target platform's latency distribution to mitigate the difference.

Effect of action time in RL agents: Prior works (Riedmiller, 2012; Travník et al., 2018) have studied the relationship between decision making time (i.e., time taken to decide an action) and task performance in RL agents. The authors propose reactive reinforcement learning algorithms propose a new “reactive SARSA” algorithm that orders computational

components without affecting the training convergence to make decision making faster. In Air Learning, we expose a similar effect where differences in training hardware (high-end CPU/GPU) and deployment hardware (embedded CPUs) can result in entirely different agent behavior. To that end, we propose a novel action scaling technique based on Hardware-in-the-loop that minimizes the differences between training and deployment of the agent on resource-constrained hardware. Unlike “reactive SARSA” (Travnik et al., 2018), we do not make any changes to the RL algorithm.

Another related work (Mahmood et al., 2018) studies the impact of delays in the action time in the robotic arm’s context. The authors use previously computed action until a new action is processed. We study the same problem in aerial robots, where we show that the differences in training and deployment hardware are another source of introducing processing delays and often overlooked. Since drones are deployed in a more dynamic environment, delayed action reduces the drones’ reactivity and can severely hinder their safety. To mitigate the performance gaps (hardware gap), we use the HIL methodology to model the target hardware delays and use them for training the policy.

In summary, Air Learning provides an open source toolset and benchmark loaded with the features to develop deep RL based applications for UAVs. It helps design effective policies, and also characterize them on an onboard computer using the HIL methodology and quality-of-flight metrics. With that in mind, it is possible to start optimizing algorithms for UAVs, treating the entire UAV and its operation as a system.

4 Air Learning

In this section, we describe the various Air Learning components. The different stages are shown in Fig. 2, which allows researchers to develop and benchmark learning algorithms for autonomous UAVs. Air Learning consists of six key components: an environment generator, an algorithm exploration framework, closed-loop real-time hardware in the loop setup, an energy and power model for UAVs, quality of flight metrics that are conscious of the UAV’s resource constraints, and a runtime system that orchestrates all of these components. By using all these components in unison, Air Learning allows us to fine-tune algorithms for the underlying hardware carefully.

4.1 Environment generator

Learning algorithms are data hungry, and the availability of high-quality data is vital for the learning process. Also, an environment that is good to learn from should include different scenarios that are challenging for the robot. By adding these challenging situations, they learn to solve those challenges. For instance, for teaching a robot to navigate obstacles, the data set should have a wide variety of obstacles (materials, textures, speeds, etc.) during the training process.

We designed an environment generator specifically targeted for autonomous UAVs. Air Learning’s environment generator creates high fidelity photo-realistic environments for the UAVs to fly in. The environment generator is built on top of UE4 and uses the AirSim UE4 (Shah et al., 2017) plugin for the UAV model and flight physics. The environment generator with the AirSim plugin is exposed as OpenAI gym interface.

The environment generator has different configuration knobs for generating challenging environments. The configuration knobs available in the current version can be classified

Fig. 2 Air Learning toolset for deep RL benchmarking in autonomous aerial machines. Our toolset consists of four main components. First, it has a configurable random environment generator built on top of UE4, a photo-realistic game engine that can be used to create a variety of different randomized environments. Second, the random environment generators are integrated with AirSim, OpenAI gym, and baselines for agile development and prototyping different state of the art reinforcement learning algorithms and policies for autonomous aerial vehicles. Third, its backend uses tools like Keras/Tensorflow that allow the design and exploration of different policies. Lastly, Air Learning uses the “hardware in the loop” methodology for characterizing the performance of the learned policies on real embedded hardware platforms. In short, it is an interdisciplinary tool that allows researchers to work from algorithm to hardware with the intent of enabling intra- and inter-layer understanding of execution. It also outputs a set of “Quality-of-Flight” metrics to understand execution

into two categories. The first category includes the parameters that can be controlled via a game configuration file. The second category consists of the parameters that can be controlled outside the game configuration file. The full list of parameters that can be controlled are tabulated in Table 2. Figure 3 shows some examples of a randomly generated arena using the environment generator. For more information on these parameters, please refer “Appendix” section.

4.2 Algorithm exploration

Deep reinforcement learning is still a nascent field that is rapidly evolving. Hence, there is significant infrastructure overhead to integrate random environment generator and evaluate new deep reinforcement learning algorithms for UAVs.

So, we expose our random environment generator and AirSim UE4 plugin as an OpenAI gym interface and integrate it popular reinforcement learning framework with stable baselines (Hill et al., 2018), which is based on OpenAI baselines.⁴ To expose our random environment generator into an OpenAI gym interface, we extend the work of Air-Gym (Kjell, 2018) to add support for environment randomization, a wide range of sensors (Depth image, Inertial Measurement Unit (IMU) data, RGB image, etc.) from AirSim and support exploring multimodal policies.

We seed the Air Learning algorithm suite with two popular and commonly used reinforcement learning algorithms. The first is Deep Q Network (DQN) (Mnih et al., 2013) and

⁴ We also support Keras-RL, another widely used RL framework.

Table 2 List of configurations available in current version of Air Learning environment generator

Parameter	Format	Description
<i>Arena Size</i>	[length, width, height]	Spawns a rectangular arena of "length" x "width" x "height"
<i>Wall Colors</i>	[R, G, B]	The colors of the wall of in [Red, Green, Blue] color format
<i>Asset</i>	<folder name>	Air Learning allows any UE4 asset to be imported into the project
<i># Static Obstacles</i>	Scalar Integer	The number of static obstacles in the arena
<i># Dynamic Obstacles</i>	Scalar Integer	The number of the dynamic obstacle in the arena
<i>Seed</i>	Scalar Integer	Seed value used in randomization
<i>Minimum Distance</i>	Scalar Integer	Minimum distance between two obstacle in the arena
<i>Goal Position</i>	[X, Y, Z]	Sets the goal position in X, Y and Z coordinates
<i>Velocity</i>	[V_{max} , V_{min}]	Velocity of the dynamic obstacle between V_{min} and V_{max}
<i>Materials</i>	<folder name>	Any UE4 material can be assigned to the UE4 asset
<i>Textures</i>	<folder name>	Any UE4 Texture can be assigned to the UE4 asset

Fig. 3 The environment generator generates different arena sizes with configurable wall texture colors, obstacles, obstacle materials etc. **a** arena with crimson colored walls with dimensions $50 \text{ m} \times 50 \text{ m} \times 5 \text{ m}$. The arena can be small or several miles long. The wall texture color is specified as an $[R, G, B]$ tuple, which allows the generator to create any color in the visible spectrum. **b** Some of the UE4 asset used in Air Learning. Any UE4 asset can be imported and Air Learning environment generator will randomly select and spawn it in the arena. **c** Arena with random obstacles. The positions of the obstacles can be changed every episode or a rate specified by the user (Color figure online)

Fig. 4 **a** Normalized reward during training for algorithm exploration between PPO-NC and DQN-NC. **b** Normalized reward during training for algorithm exploration between PPO-C and DQN-C. We find that the DQN agent performs better than the PPO agent irrespective of whether the agent was trained using curriculum learning or non-curriculum learning. The rewards are averaged over five runs with random seeds

the second is Proximal Policy Optimization (PPO) (Schulman et al., 2017). DQN falls into the discrete action algorithms where the action space is high-level commands (‘move forward,’ ‘move left’ e.t.c.,) and Proximal Policy Optimization falls into the continuous action algorithms (e.g., policy predicts the continuous value of velocity vector). For each of the algorithm variants, we also support an option to train the agent using curriculum learning (Bengio et al., 2009). For both these algorithms, we keep the observation space, policy architecture and reward structure same and compare agent performance. The environment configuration used in the training of PPO/DQN, the policy architecture, the reward function, is described in the appendix (“Appendix B” section).

Figure 4a shows the normalized reward of the DQN agent (DQN-NC) and PPO agent (PPO-NC) trained using non-curriculum learning. One of the observations is that the PPO agent trained using non-curriculum learning consistently accrues negative reward throughout the training duration. In contrast, the DQN agent trained using non-curriculum learning starts at the same as the PPO agent but the DQN agent accrues more reward beginning in the 2000th episode.

Figure 4b shows the normalized episodic reward for the DQN (DQN-C) and PPO (PPO-C) agents trained using curriculum learning. We observe a similar trend as we saw with the agents trained using non-curriculum learning where the DQN agent outperforms the PPO agent. However, in this case, the PPO agent has a positive total reward. But the DQN agent starts to accrue more reward starting from the 1000th episode. Also, the slight dip in the reward at 3800th is due to the curriculum's change (increased difficulty).

Reflecting on the results, we gathered in Fig. 4a, b, continuous action reinforcement learning algorithms such as PPO have generally been known to show promising results for low-level flight controller tasks that are used for stabilizing UAVs (Hwangbo et al., 2017). However, as our results indicate, applying these algorithms for a complex task, such as end-to-end navigation in a photo-realistic simulator, can be challenging for a couple of reasons.

First, we believe that the action space for the PPO agent limits the exploration compared to the DQN agent. For the PPO agent, the action space is the components of velocity vector v_x and v_y whose value can vary from $[-5 \text{ m/s}, 5 \text{ m/s}]$. Having such an action space can be a constraining factor for PPO. For instance, if the agent observes an obstacle at the front, it needs to take action such that it moves right or left. Now for PPO agent, since the action space is continuous values of $[V_x, V_y]$, for it to move forward in the x -direction, the V_x can be any positive number while the V_y component has to be '0'. It can be quite challenging for the PPO agent (or continuous action algorithm) to learn this behavior, and it might require a much more sophisticated reward function that identifies these scenarios and rewards or penalizes these behaviors accordingly. In contrast, for the DQN agent, the action space is much simpler since it has to only yaw (i.e., move left or right) and then move forward or vice versa.

Second, in our evaluation, we keep the reward function, input observation and the policy architecture same for DQN and PPO agent. We choose to fix these because we want to focus on showcasing the capability of the Air Learning infrastructure. Since RL algorithms are sensitive to hyperparameters and the choice of the reward function, it could be possible that our reward function, policy architecture could have inadvertently favored the DQN agent compared to the PPO agent. The sensitivity of the RL algorithms to the policy and reward is still an open research problem (Judah et al., 2014; Su et al., 2015).

The takeaway is that we can do algorithm exploratory studies with Air Learning. For high-level task like point-to-point navigation, discrete action reinforcement learning algorithms like DQN allows more flexibility compared to continuous action reinforcement learning algorithms like PPO. We also demonstrate that incorporating techniques such as curriculum learning can be beneficial to the overall learning.

4.3 Policy exploration

Another essential aspect of deep reinforcement learning is the policy, which determines the best action to take. Given a particular state the policy needs to maximize the reward. A neural network approximates the policies. To assist the researchers in exploring effective policies, we use Keras/TensorFlow (Chollet, 2015) as the machine learning back-end tool. Later on, we demonstrate how one can do algorithm and policy explorations for tasks like autonomous navigation though Air Learning is by no means limited to this task alone.

4.4 Hardware exploration

Often aerial roboticists port the algorithm onto UAVs to validate the functionality of the algorithms. These UAVs can be custom built (NVIDAA-AI-IOT, 2015) or commercially available off-the-shelf (COTS) UAVs (Hummingbird, 2018; Intel, 2018) but mostly have fixed hardware that can be used as onboard compute. A critical shortcoming of this approach is that the roboticist cannot experiment with hardware changes. More powerful hardware may (or may not) unlock additional capabilities during flight, but there is no way to know until the hardware is available on a real UAV so that the roboticist can physically experiment with the platform.

Reasons for wanting to do such exploration includes understanding the computational requirements of the system, quantifying the energy consumption implications as a result of interactions between the algorithm and the hardware, and so forth. Such evaluation is crucial to determine whether an algorithm is, in fact, feasible when ported to a real UAV with a specific hardware configuration and battery constraints.

For instance, a Parrot Bepop (Parrot, 2019) comes with a P7 dual-core CPU Cortex A9 and a Quad core GPU. It is not possible to fly the UAV assuming a different piece of hardware, such as the NVIDIA Xavier (NVIDIA, 2019) processor that is significantly more powerful; at the time of this writing there is no COTS UAV that contains the Xavier platform. So, one would have to wait until a commercially viable platform is available. However, using Air Learning, one can experiment how the UAV would behave with a Xavier since the UAV is flying virtually.

Hardware exploration in Air Learning allows for evaluation of the best reinforcement learning algorithm and its policy on different hardware. It is not limited by the onboard compute available on the real robot. Once the best algorithm and policy are determined, Air Learning allows for characterizing the performance of these algorithms and policies on different types of hardware platforms. It also enables to carefully fine-tune and co-design algorithms and policy while being mindful of the resource constraints and other limitation of the hardware.

A HIL simulation combines the benefits of the real design and the simulation by allowing them to interact with one another as shown in Fig. 5. There are three core components in Air Learning's HIL methodology: (1) a high-end desktop that simulates a virtual environment flying the UAV (*top*); (2) an embedded system that runs the operating system, the deep reinforcement learning algorithms, policies and associated software stack (*left*); and (3) a flight controller that controls the flight of the UAV in the simulated environment (*right*).

The simulated environment models the various sensors (RGB/Depth Cameras), actuators (rotors), and the physical world surrounding the agent (Obstacles). This data is fed into the reinforcement learning algorithms that are running on the embedded companion computer, which processes the input and outputs flight commands to the flight controller. The controller then communicates those commands into the virtual UAV flying inside the simulated game environment.

The interaction between the three components is what allows us to evaluate the algorithms and policy on various embedded computing platforms. The HIL setup we present allows for the swap-ability of the embedded platform under test. The methodology enables us to effectively measure both the performance and energy of the agent holistically and more accurately, since one can evaluate how well an algorithm performs on a variety of different platforms.

Fig. 5 Hardware-in-the-loop (HIL) simulation in Air Learning

In our evaluation, which we discuss later, we use a Raspberry Pi (Ras-Pi 4) as the embedded hardware platform to evaluate the best performing deep reinforcement learning algorithm and its associated policy. The HIL setup includes running the environment generator on a high-end desktop with a GPU. The reinforcement learning algorithm and its associated policy run on the Ras-Pi 4. The state information (Depth image, RGB image, IMU) are requested by Ras-Pi 3 using AirSim Plugins APIs which involves an RPC (remote procedural calls) over TCP/IP network (both high-end desktop and Ras-Pi 4 are connected by ethernet). The policy evaluates the actions based on the state information it received from the high-end desktop. The actions are relayed back to the high-end desktop through AirSim flight controller API's.

4.5 Energy model in AirSim plugin

In Air Learning, we use the energy simulator we developed in our prior work (Boroujerdian et al., 2018). The AirSim plugin is extended with a battery and energy model. The energy model is a function of UAVs velocity, acceleration. The values of velocity and acceleration are continuously sampled and using these we estimate the power as proposed in this work (Tseng et al., 2017). The power is calculated using the following formula:

$$\begin{aligned}
 P = & \begin{bmatrix} \beta_1 \\ \beta_2 \\ \beta_3 \end{bmatrix}^T \begin{bmatrix} \|\vec{v}_{xy}\| \\ \|\vec{a}_{xy}\| \\ \|\vec{v}_{xy}\| \|\vec{a}_{xy}\| \end{bmatrix} + \begin{bmatrix} \beta_4 \\ \beta_5 \\ \beta_6 \end{bmatrix}^T \begin{bmatrix} \|\vec{v}_z\| \\ \|\vec{a}_z\| \\ \|\vec{v}_z\| \|\vec{a}_z\| \end{bmatrix} \\
 & + \begin{bmatrix} \beta_7 \\ \beta_8 \\ \beta_9 \end{bmatrix}^T \begin{bmatrix} m \\ \vec{v}_{xy} \cdot \vec{w}_{xy} \\ 1 \end{bmatrix}
 \end{aligned} \quad (1)$$

In Eq. 1, v_{xy} and a_{xy} are the velocity and acceleration in the horizontal direction. v_z and a_z denotes the velocity and acceleration in the z direction. m denotes the mass of the payload. β_1 to β_9 are the coefficients based on the model of the UAV used in the simulation. For the energy calculation model, we use the columb counter technique as described in prior work (Kumar et al., 2016). The simulator computes the total number of columb that has passed over the battery over every cycle.

Using the energy model Air Learning allows us to monitor the energy continuously during training or during the evaluation of the reinforcement learning algorithm.

4.6 Quality of flight metrics

Reinforcement learning algorithms are often evaluated based on success rate where the success rate is based on whether the algorithm completed the mission. This metric only captures the functionality of the algorithm and grossly ignores how well the algorithm performs in the real world. In the real world, there are additional constraints for a UAV, such as the limited onboard compute capability and battery capacity.

Hence, we need additional metrics that can quantify the performance of learning algorithms more holistically. To this end, Air Learning introduces Quality-of-Flight (QoF) metrics that not only captures the functionality of the algorithm but also how well they perform when ported to onboard compute in real UAVs. For instance, the algorithm and policies are only useful if they accomplish the goals within finite energy available in the UAVs. Hence, algorithms and policies need to be evaluated on the metrics that describe the quality of flight such as mission time, distance flown, etc. In the first version of Air Learning, we consider the following metrics.

Success rate: The percentage of time the UAV reaches the goal state without collisions and running out of battery. Ideally, this number will be close to 100% as it reflects the algorithms' functionality, taking into account resource constraints.

Time to completion: The total time UAV spends finishing a mission within the simulated world.

Energy consumed: The total energy spent while carrying out the mission. Limited battery available onboard constrains the mission time. Hence, monitoring energy usage is of utmost importance for autonomous aerial vehicles, and therefore should be a measure of policy's efficiency.

Distance traveled: Total distance flown while carrying out the mission. This metric is the average length of the trajectory that can be used to measure how well the policy did.

4.7 Runtime system

The final part is the runtime system that orchestrates the overall execution. The runtime system starts the game engine with the correct configuration of the environment before the agent starts. It also monitors the episodic progress of the reinforcement learning algorithm and ensures that before starting a new episode that it randomizes the different parameters, so the agent statistically gets a new environment. It also has resiliency built into it to resume the training in case any one of the components (for example UE4 engine) crashes.

In summary, using Air Learning environment generator, researchers can develop various challenging scenarios to design better learning algorithms. Using Air Learning interfaces to OpenAI gym, stable-baselines and TensorFlow backend, they can rapidly evaluate different reinforcement learning algorithms and their associated policies. Using Air Learning HIL methodology and QoF metrics, they can benchmark the performance of learning algorithms and policies on resource-constrained onboard compute platforms.

5 Experimental evaluation prelude

The next few sections focus heavily on how Air Learning can be used to demonstrate its value. As a prelude, this section presents the highlights to focus on the big picture.

Policy evaluation (Sect. 6): We show how Air Learning can be used to explore different reinforcement learning based policies. We use the best algorithm determined during the algorithm exploration step and use that algorithm to explore the best policy. In this work, we use Air Learning environment generator to generate three environments, namely, *No Obstacles*, *Static Obstacles*, and *Dynamic Obstacles*. These three environments create a varying level of difficulties by changing the number of static and dynamic obstacles in the environments for the autonomous navigation task.

We also show how Air Learning allows end users to perform benchmarking of the policies by showing two examples. In the first example, we show how well the policies trained in one environment generalize to the other environments. In the second example, we show to which of the sensor inputs the policy is most sensitive towards. This insight can be used while designing the network architecture of the policy. For instance, we show that image input has the highest sensitivity amongst other inputs. Hence a future iteration of the policy can have more feature extractors (increasing the depth of filters) dedicated to the image input.

System evaluation (Sect. 7): We show the importance of benchmarking algorithm performance on resource-constrained hardware such as what is typical of a UAV compute platform. In this work, we use a Raspberry Pi 4 (Ras-Pi 4) as an example of resource-constrained hardware. We use the best policies determined in the policy exploration step (Sect. 6) and use that to compare the performance between Intel Core-i9 and Ras-Pi 4 using HIL and the QoF metrics available in Air Learning. We also show how to artificially degrade the performance of the Intel Core-i9 to show how compute performance can potentially affect the behavior of a policy when it is ported over to a real aerial robot.

In summary, using these focused studies, we demonstrate how Air Learning can be used by researchers to design and benchmark algorithm-hardware interactions in autonomous aerial vehicles, as shown previously in Fig. 2.

6 Policy exploration

In this section, we perform policy exploration for the DQN agent with curriculum learning (Bengio et al., 2009). The policy exploration phase aims to determine the best neural network policy architecture for each of the tasks (i.e., autonomous navigation) in different environments with and without obstacles.

We start with a basic template architecture, as shown in Fig. 6. The architecture is multi-modal and takes depth image, velocity, and position data as its input. Using this template, we sweep two parameters, namely *# Layers* and *# Filters* (making the policy wider and deeper). To simplify the search, for convolution layers, we restrict filter sizes to 3×3 with stride 1. This choice ensures that there is no loss of pixel information. Likewise, for fully-connected layers, *# Filter* parameter denotes the number of hidden neurons in that layer. The choice of using *# Layers* and *# Filters* parameters to control both the convolution and fully-connected layers is to manage the complexity of searching over large NN hyperparameters design space.

The *# Layers* and *# Filters* and the template policy architecture can be used to construct a variety of different policies. For example, a tuple of (*# Filters* = 32, *# Layers* = 5) will result in a policy architecture where there five convolution layers with 32 filters (with 3×3 filters) followed by five fully-connected layers with 32 hidden neurons each. For each of the navigation tasks (in different environments), we sweep the template parameters (*# Layers* and *# Filters*) to explore multiple policy architectures for the DQN agent.

6.1 Training and testing methodology

The training and testing methodology for the DQN agent running in the different environments is described below.

Fig. 6 The network architecture template for the policies used in DQN agents. We sweep the *# Layers* and *# Filters* parameters in the network architecture template. Both the agents take a depth image, velocity vector, and position vector as inputs. The depth image is passed through *# Layers* of convolutions layers with *# Filters* each. *# Layers* and *# Filters* are variables that we sweep. We also use a uniform filter size of (3×3) with stride of 1. The combined vector space is passed to the *# Layers* of fully connected network, each with *# Filters* hidden units. The choice of using *# Layers* and *# Filters* parameters to control both the convolution and fully-connected layers is to manage the complexity of searching over large NN hyperparameters design space. The action space determines the number of hidden units in the last fully connected layer. For the DQN agent, we have twenty-five actions

Environments: For the point-to-point autonomous navigation task for UAVs, we create three randomly generated environments, namely, *No Obstacles*, *Static Obstacles*, and *Dynamic Obstacles*, with varying levels of static obstacles and dynamic obstacles. The environment size for all three levels is 50 m $\times$ 50 m. For the *No Obstacles* environment, there are no obstacles in the main arena, but the goal position is changed every episode. For *Static Obstacles*, the number of obstacles varies from five to ten, and it is changed every four episodes. The end goal and position of the obstacles are changed every episode. For *Dynamic Obstacles*, along with five static obstacles, we introduce up to five dynamic obstacles of whose velocities range from 1 to 2.5 m/s. The obstacles and goals are placed in random locations every episode to ensure that the policy does not over-fit.

Training methodology: We train the DQN agent using curriculum learning in the environments described above. We use the same methodology described in “Appendix B” section, where we checkpoint policy in each zone for the three environments. The hardware used in training is an Intel Core-i9 CPU with an Nvidia GTX 2080-TI GPU.

Testing methodology: For testing the policies, we evaluate the checkpoints saved in the final zone. Each policy is evaluated on 100 randomly generated goal/obstacle configuration (controlled by the ‘Seed’ parameter in Table 2). The same 100 randomly generated environment configurations are used across different policy evaluations. The hardware we use for testing the policies is the same as the hardware used for training them (Intel Core-i9 with Nvidia GTX 2080-TI).

6.2 Policy selection

The policy architecture search for *No Obstacles*, *Static Obstacles*, and *Dynamic Obstacles* are shown in Fig. 8. Figure 7a–c show the success rate for different policy architecture searched for the DQN agent trained using curriculum learning on *No Obstacles*, *Static Obstacles*, and *Dynamic Obstacles* environments, respectively. In the figures, the x-axis corresponds # Filter sizes (32, 48, or 64) and the y-axis corresponds to the # Layers (2, 3, 4, 5, and 6) for *No Obstacles/Static Obstacles* environments and # Layers (5, 6, 7, 8, 9) for *Dynamic Obstacles* environment. The reason for sweeping different (larger) policies is because “Dynamic Obstacles” will be a harder task, and a deeper policy might help improve the success rate compared to a shallow policy. Each cell corresponds to a unique policy architecture based on the template defined in Fig. 7. The value in each cell corresponds to the success rate for the best policy architecture. The $\pm$ denotes the standard deviation (error bounds) across five seeds. For instance, in Fig. 7a, the best performing policy architecture with # Filters of 32 and # Layers of 2 results in a 72% success rate. The success rate across five seeds results in a standard deviation of $\pm$ of 8% error. For evaluation, we always choose the best performing policy (i.e., the policy that achieves best success rate).

Based on the policy architecture search, we notice that as the task complexity increases (obstacle density increases), a larger policy improves the task success rate. For instance, in the *No Obstacles* case (Fig. 7a), the policy with # Filters of 32 and # Layers of 5 achieves the highest success rate of 91%. Even though we name the environment No Obstacles, the UAV agent can still collide with the arena walls, which lowers the success rate. For *Static Obstacles* case (Fig. 7b), the policy with # Filters of 48 and # Layers of 4 achieves the best success rate of 84%. Likewise, for *Dynamic Obstacles* case (Fig. 7c), the policy architecture with # Filters of 32 and # Layers of 7 achieves the best success rate of 61%.

Fig. 7 **a**, **b**, and **c** show the policy architecture search for the *No Obstacles*, *Static Obstacles*, and *Dynamic Obstacles* environments. Each cell shows the success rate for the policies for # Layers and # Filters’ corresponding values. The success rate is evaluated in *Zone 3*, which is the region that is not used during training. Each policy is evaluated on the same 100 randomly generated environment configuration (controlled by the ‘Seed’ parameter described in Table 2). The policy architecture with the highest success rate is chosen as the best policy for DQN agents in the environment with no obstacles, static obstacles, and dynamic obstacles. The standard deviation error across multiple seeds are denoted by ($\pm$) sign. For the *No Obstacles* environment, the policy with # Layers of five and # Filters of 32 is chosen as the best performing policy. Likewise, for the *Dynamic Obstacles* environment, the policy architecture with # Layers of 7 and # Filter of 32 is chosen as the best policy

The success rate loss in *Static Obstacles* and *Dynamic Obstacles* cases can be attributed to an increase in the possibility of collisions with static and dynamic obstacles.

6.3 Success rate across the different environments

To study how a policy trained in one environment performs in other environments, we take the best policy trained in the *No Obstacles* environment and evaluate it on the *Static Obstacles* and *Dynamic Obstacles* environments. We do the same for the best policy trained on *Dynamic Obstacles* and assess it on the *No Obstacles* and *Static Obstacles* environments.

The results for the generalization study are tabulated in Table 3. We see that the policy trained in the *No Obstacles* environment has a steep drop in success rate from 91 to 53% in *Static Obstacles* and 32% in *Dynamic Obstacles* environment, respectively. In contrast, we observe that the policy trained in the *Dynamic Obstacles* environment has an increased success rate from 61 to 89% in the *No Obstacles* and 74% in the *Static Obstacles* environment, respectively.

The drop in the success rate for the policy trained in the *No Obstacles* environment is expected because, during its training, the agent might not have encountered a variety of obstacles (static and dynamic obstacles) to learn from as it might have encountered in the

Table 3 Evaluation of the best-performing policies trained in one environment tested in another environment

Policy (# Layers, # Filters)	No obstacles	Static obstacles	Dynamic obstacles
<i>No Obstacles</i> (5,32)	0.91	0.53	0.32
<i>Dynamic Obstacles</i> (7,32)	0.89	0.74	0.61

We evaluate the best performing policy (7 Layers, 32 Filters) trained on *Dynamic Obstacles* in *No Obstacles* and *Static Obstacles* environment. Likewise, we also evaluate the best performing policy (5 Layers, 32 Filters) trained in the *No Obstacles* environment in the *Static Obstacles* and *Dynamic Obstacles* environments

other two environments. The same reasoning can also apply to the improvement in the success rate observed in the policy trained in the *Dynamic Obstacles* environment when it is evaluated on the *No Obstacles* and *Static Obstacles* environments.

In general, the agent performs best in the environment where it is trained, which is expected. But we also observe that training an agent in a more challenging environment can yield good results when evaluating in a much less challenging environment. Hence, having a random environment generator, such as what we have enabled in Air Learning, can help the policy generalize well by creating a wide variety of different experiences for the agent to experience during training.

6.4 Success rate sensitivity to sensor input ablation

In doing policy exploration, one is also interested in studying the policy's sensitivity towards a particular sensor input. So we ablate the sensor inputs to the policy to understand the effects. We ablate the policy's inputs one by one and see the impact of various ablation and its success rate. It is important to note that we do not re-train the policy with ablated inputs. This is to perform reliability study and simulate the real-world scenario if a particular sensing modality is corrupted.

The policy architecture we used for the DQN agent in this work is multi-modal in nature which receives depth image, velocity measurement V_t and position vector X_t as inputs. The V_t is a 1-dimensional vector of the form $[v_x, v_y, v_z]$ where v_x, v_y, v_z are the components of velocity vector in x, y and z directions at time 't'. The X_t is a 1-dimensional vector of the form $[X_{goal}, Y_{goal}, D_{goal}]$, where X_{goal} and Y_{goal} are the relative 'x' and 'y' distance with respect to the goal position and D_{goal} is the euclidean distance to the goal from the agent's current position.

The baseline success rate we use in this study is when all the three inputs are fed to the policy. The velocity ablation study refers to removing the velocity input measurements from policy inputs. Likewise, the position ablation study and depth image ablation study refer to removing the position vector and depth image from the policy's input stream. The results of various input ablation studies are plotted in Fig. 8.

For the *No Obstacles* environment, the policy success rate drops from 91% to 53% when velocity measurements are ablated. When the depth image is ablated, we find that the success rate drops to 7%, and when the position vector is ablated, the success rate drops to 42%. Similarly, for *Static Obstacles*, we find that if the depth image input is ablated, it fails to reach the destination. Likewise, when the velocity and position inputs are ablated, we observe the success rate drops from 84% to 33%. Similarly, we see a similar observation in a *Dynamic Obstacles* environment where the success rate drops to 0% when the depth image is ablated.

The depth image is the highest contributor to the policy's success, whereas the velocity input is significant but least among the other two inputs. The drop in the policy success rate due to depth image ablation is evident from policy architecture since maximum features in the flatten layer are contributed by the depth image than velocity and position (both 1×3 vectors). Another interesting observation is that when the position input is ablated, the agent also loses the information about its goal. The lack of goal position results in an exploration policy capable of avoiding obstacles (due to depth image input). In *No Obstacles* environment (where there are no obstacles except walls), the agent is free to explore unless it collides with the walls or exhaust maximum allowed steps. Due to the exploration,

Fig. 8 The effect of ablating the sensor inputs on the success rate. We observe that the depth image contributes the most to the policy's success, whereas velocity input affects the least in the success. All the policy evaluations are in Zone3 on the Intel Core-i9 platform

the agent reaches the goal position 42 out of 100 times. Our results are in line with prior work (Duisterhof et al., 2019; Palacin et al., 2005) where such random action-based exploration yields some amount of success. However, in a cluttered environment, random exploration may result in sub-optimal performance due to a higher probability of collision or exhausting maximum allowed steps (a proxy for limited battery energy).

Using Air Learning, researchers can gain better insights into how reliable a particular set of inputs in the case of sensor failures. The reliability studies and its impact on learning algorithms are essential given the kind of application the autonomous aerial vehicles are targeted. Also, understanding the sensitivity of a particular input towards success can lead to better policies where more feature extraction can be assigned to those inputs.

7 System evaluation

This section demonstrates how Air Learning can benchmark the algorithm and policy's performance on a resource-constrained onboard compute platform, post-training. We use the HIL methodology (Sect. 4.4) and QoF metrics (Sect. 4.6) for benchmarking the DQN agent and its policy. We evaluate them on the three different randomly generated environments described in Sect. 6.

7.1 Experimental setup

The experimental setup has two components, namely, the server and System Under Test (SUT), as shown in Fig. 9. The server component is responsible for rendering the environment (for example, *No Obstacles*). The server consists of an 18 core Intel Core-i9 processor with an Nvidia RTX-2080. The SUT component is the system on which we want to evaluate the policy. The SUT is the proxy for the onboard compute system used in UAVs. In this work, we compare the policies' performance on two systems, namely Intel Core-i9 and Ras-Pi 4. The key differences between the Intel Core-i9 and Ras-Pi 4 platform are tabulated in Table 4. The systems are vastly different in their performance capabilities and represent ends of the performance spectrum.

Three latencies affect the overall processing time. The first is t_1 , which is the latency to extract the state information (Depth Image, RGB Image, etc.) from the server. The state information is fetched from the server to the SUT. The communication protocol used

Fig. 9 The Experimental setup for policy evaluation on two different platforms. The platform under test is called the System Under Test (SUT). The environments are rendered on a server with Intel Core-i9 with Nvidia RTX 2080. Clock speed is a function in the AirSim plugin, which speeds up the environment time relative to the real world clock. In our evaluation, we set the clock speed to 2X. Time t_1 is the time it takes to get the state information from the environment to the SUT. We use an Intel Core-i9 and a Ras-Pi 4 as the two SUTs. Time t_2 is the time it takes to evaluate the forward pass of the neural network policy. This latency depends on the SUT. It is different for the Intel Core-i9 and the Ras-Pi 4. Time t_3 is the actuation time for which the control is applied

Table 4 The most pertinent System Under Test (SUT) specifications for the Intel Core-i9 and Ras-Pi 4 systems

Platform	Intel Core-i9	Ras-Pi 4
CPU cores	18	4 Arm-A72
CPU frequency	4.2 GHz	1.5 GHz
GPU	Nvidia GTX 2080 TI	None
Power	350 W	<1.7 W
Cost	\$1500	\$75

between the server and the SUT is TCP/IP. Initially, we found that ethernet adapter on Intel Core-i9 faster compared to the ethernet adapter on Ras-Pi 4. We make the t_1 latencies between Intel Core-i9 and Ras-Pi 4 same by adding artificial sleep for Intel Core-i9 platform.⁵

The second latency is t_2 , which is the policy evaluation time for the SUT (i.e., the Intel Core-i9 or the Ras-Pi 4). The policies are evaluated on the SUT, which predicts the output actions based on the input state information received from the server. The policy architecture used in this work has 40.3 Million (*No Obstacles* and *Static Obstacles*) and 161.77 Million parameters (*Dynamic Obstacles*). The t_2 latency for *No Obstacles* policy on Ras-Pi 4 is 396 ms, while on the desktop, equipped with GTX 2080 Ti GPU and Intel Core i9 CPU, it is 11 ms. The desktop is 36x times faster.

The third latency is t_3 . Once the policies are evaluated, it predicts actions. These actions are converted to the low-level actuation using the AirSim flight controller APIs.⁶ These APIs have a duration parameter that controls the duration of a particular action must be applied. This duration parameter is denoted by t_3 , and it is kept the same for both SUTs.

To evaluate the impact of the SUT performance on the overall learning behavior, we keep the t_1 and t_3 latencies constant for both Intel Core-i9 and Ras-Pi 4 systems. We focus only on the difference in the policy evaluation time (i.e., t_2) and study how it affects the overall performance time. Using this setup, we evaluate the best policy determined in Sect. 6 for environments with no obstacles, static obstacles, and dynamic obstacles.

⁵ The sleep latency value that was added to Intel Core-i9 was determined by doing a ping test with the packet size equal to the size of the data (Depth Image) we fetch from the server and averaged it over 50 iterations.

⁶ <https://github.com/Microsoft/AirSim/blob/master/docs/apis.md>.

7.2 Desktop vs. embedded SUT performance

In Table 5, we compare the performance of the policy on a Intel Core-i9 (high-end desktop) and the Ras-Pi 4. We evaluate the best policy on the *No Obstacles*, *Static Obstacles* and *Dynamic Obstacles* environments described previously in Sect. 6.

In the *No Obstacles* case, the policy running on the high-end desktop is 11% more successful compared to the policy running on the Ras-Pi 4. The flight time to reach the goal on the desktop is, on average, 25.29 s, whereas on the Ras-Pi 4, it is 37.37 s, which yields a performance gap of around 47.76%. The distance flown for the same policy on the desktop is 27.59 m, whereas on the Ras-Pi 4, it is 33.06 m, which contributes to a difference of 19.82%. Finally, the desktop consumes an average of 20 kJ of energy, while the Ras-Pi 4 consumes 25.4 kJ, which is 29.48% more energy.

In the *Static Obstacles* case, the policy running on the desktop is 13% more successful than the policy running on Ras-Pi 4. The flight time to reach the goal on the high-end desktop is, on average, 30.25 s, whereas on the Ras-Pi 4, it is 34.44 s. That yields a performance gap of around 13.85%. For the distance flown, the policy running on the desktop has a trajectory length of 28.7 m, whereas the same policy on the Ras-Pi 4 has a trajectory length of 32.57 m. This contributes to a difference in 13.4%. For energy, the policy running on the desktop on an average consumes 19.2 kJ of energy, while policy running on Ras-Pi 4 on an average consumes 23.90 KJ of energy, which is about 32% more energy.

In the *Dynamic Obstacles* case, the success rate between the desktop and the Ras-Pi 4 is 6%. The flight time to reach the goal on the desktop is, on average, 21.48 s, whereas on the Ras-Pi 4, it is 35.36 s, yielding a performance gap of around 64.61%. For the distance flown, the policy running on the desktop has a trajectory length of 23.51 m, whereas the same policy running on Ras-Pi 4 has a trajectory length of 32.86 m. This contributes to a difference in 40%. For energy, the policy running on the desktop, on average, consumes 18.76 kJ of energy while policy running on Ras-Pi 4 consumes 24.31 KJ of energy, which is about 30% more energy.

Overall, across the three different environments, the policy evaluated on the Ras-Pi 4 achieves a success rate that is within 13% compared to the policy assessed on the desktop. While some degradation in performance is expected, the magnitude of the deterioration is more severe for the other QoF metrics, such as flight time, energy, and distance flown. This difference is significant to note because when the policies are ported to resource-constrained compute like the Ras-Pi 4 (a proxy for onboard compute in real UAVs), they could perform worse, such as being unable to finish the mission due to low battery.

In summary, the takeaway is that evaluations of policies solely on a high-end machine do not accurately reflect the real-time performance on an embedded compute system such as those available on UAVs. Hence, relying on success rate as the sole metric is insufficient, though this is by and a large state of the art means to report success. Using Air Learning and its HIL methodology and QoF metrics, we can understand to what extent the choice of onboard compute affects the performance of the algorithm.

7.3 Root-cause analysis of SUT performance differences

It is important to understand why the policy performs differently on the Intel Core i9 versus the Ras-Pi 4. So, we perform two experiments. First, we plot the policy trajectories on the Ras-Pi 4 and compare it to the Intel Core-i9 to understand if there is a flight path difference. Visualizing the trajectories helps us build intuition about the variations between the

Table 5 Inference time, success rate, and Quality of Flight (QoF) metrics between Intel Core i9 desktop and Ras-Pi 4 for *No Obstacles*, *Static Obstacles*, and *Dynamic Obstacles*

Metric	No obstacles				Static obstacles				Dynamic obstacles			
	Intel Core i9		Ras-Pi 4		Intel Core i9		Ras-Pi 4		Intel Core i9		Ras-Pi 4	
	Intel Core i9	Perf. Gap (%)	Ras-Pi 4	Perf. Gap (%)	Intel Core i9	Perf. Gap (%)	Ras-Pi 4	Perf. Gap (%)	Intel Core i9	Perf. Gap (%)	Ras-Pi 4	Perf. Gap (%)
Inference latency (ms) (↓)	11.00		396	3500	10		542.28	5322.8	9.3		948.92	10103.4
Success rate (%) (↑)	91		80.00	11.00	84.00		71.00	13.00	61.00		55.00	6.00
<i>QOF metrics</i>												
Flight time (s) (↓)	25.29		37.37	47.76	30.258		34.44	13.85	21.48		35.36	64.61
Distance flown (m) (↓)	27.59		33.06	19.82	28.70		32.57	13.4	23.51		32.86	39.77
Energy (kJ) (↓)	19.68		25.483	29.48	19.2		23.90	24.47	18.76		24.31	29.58

The policy under evaluation is the best policy obtained from policy evaluation (Sect. 6). Perf. Gap (%) is used to quantify the hardware gap

Fig. 10 Figures **a**, **b**, **c** compare the trajectories of Ras-Pi 4 and Intel Core-i9. The red columns in **b** and **c** denotes the position of static obstacles (Color figure online)

two platforms. Second, we take an Intel Core-i9 platform and degrade its performance by adding artificial sleep such that the policy evaluation times are similar to that of Ras-Pi 4. This helps us validate if the processing time is giving rise to the QoF metric discrepancy.

To plot the trajectories, we fix the end goal's position, obstacles, and evaluate 100 trajectories with the same configuration in the *No Obstacles*, *Static Obstacles*, and *Dynamic Obstacles* environments. The trajectories are shown in Fig. 10a–c. They are representative of repeated trajectories between the start and end goals. The trajectories between the desktop and Ras-Pi 4 are very different—the desktop trajectory orients towards the goal and proceeds directly. The Ras-Pi 4 trajectory starts toward the goal, but then drifts, resulting in a longer trajectory. This is likely a result of the actions taken because of stale sensory information, due to the longer inference time; recall there is a 20× difference in the inference time between the desktop and Ras-Pi 4 (Sect. 7.1 and Table 5).

To further root-cause and test whether the (slower) processing time (t_2) is giving rise to the long trajectories, we take the best performing policy trained on the high-end desktop in the *Static Obstacles* environment and gradually degrade the policy's evaluation time by introducing artificial sleep times into the program.⁷ Sleep time injection allows us to model the big differences in the behavior of the same policy and its sensitivity to the onboard compute performance.

Table 6 shows the effect of degrading the compute performance on policy evaluation. The baseline is the performance on the high-end Intel Core-i9 desktop. Intel Core-i9 (150 ms) and Intel Core-i9 (300 ms) are the scenarios where the performance of Intel Core-i9 is degraded by 150 ms and 300 ms, respectively. As performance deteriorates from 3 ms to 300 ms, the flight time degrades by 97%, the trajectory distance degrades by 21%, and energy degrades by 43%.

We visualize degradation impact by plotting the same policy's trajectories on the baseline Intel Core-i9 system and the degraded versions of Intel Core-i9 systems (150 ms and 300 ms). The trajectory results are shown in Fig. 11. As we artificially degrade, the drift in trajectories gets wider, which increases the trajectory length to reach the goal position, thus degrading the QoF metrics. We also see that the trajectory of the degraded Intel Core i9 closely resembles the Ras-Pi 4 trajectory.

⁷ Adding artificial sleep into the high-end desktop is a simple first-order approximation of the Ras-Pi 4 system. In reality, we cannot fully equate the high-end desktop to the Ras-Pi 4 since there are other differences (e.g., system architecture, memory sub-system, and power).

Table 6 Degradation in policy evaluation using artificially injected program sleep (proxy for performance degradation)

Metric	Core i9	Core i9 (150 ms)	Core i9 (300 ms)
Inference latency (ms) (↓)	11.00	150.00	300.00
<i>QoF metrics</i>			
Flight time (s) (↓)	24.08	32.38	47.59
Distance flown (m) (↓)	25.64	28.27	31.11
Energy (kJ) (↓)	19.09	23.18	27.28

The policy is the best performing policy trained on ‘Static Obstacles.’ The baseline is Intel Core i9 without any artificial sleep. Intel Core i9 (150 ms latency) and Intel Core i9 (300 ms latency) represent scenarios where 150 ms and 300 ms of artificially injected delay added to the policy evaluation

Fig. 11 Trajectory visualization of the best-performing policy on Intel Core i9 and artificially degraded versions of Intel Core i9 (150 ms) and Intel Core i9 (300 ms)

In summary, the onboard compute choice and algorithm profoundly affect the resulting UAV behavior and shape of the trajectory. Additional quality of flight metrics (energy, distance, etc.) captures the differences better than just the success rate. Moreover, evaluations done purely on a high-end desktop might show lower energy consumption in a mission, but when the solution is ported to real robots, the solution might consume more energy due to the sub-par performance of the onboard compute. Using the hardware-in-the-loop (HIL) methodology allows us to identify these differences and other performance bottlenecks that arise due to the onboard compute without having to port things to the real robots. Hence, a tool like Air Learning with its HIL methodology helps identify such differences at the early stage.

In the next section, we show how Air Learning HIL can mitigate the hardware gap and characterize the end-to-end learning algorithms and model these characteristics to create robust and performance-aware policies.

8 Mitigating the hardware gap

In this section, we demonstrate how Air Learning HIL technique can be used to minimize the hardware gap due to differences in the training hardware and deployment hardware (onboard compute). To that end, we propose a general methodology where we train a policy on the high-end machine with added latencies to mimic the onboard compute's performance. Using this method, we show that it minimizes the hardware gap from 38% to less than 0.5% on flight time metric, 16.03% to 1.37% on the trajectory length metric, and 15.49% to 0.1% on the energy of flight metric.

One way to mitigate the hardware gap is to directly train the policy on the onboard computer available in the robot (Ahn et al., 2020; Ha et al., 2018; Kalashnikov et al., 2018). Though on-device RL training is practical for ground-based or fixed robots to overcome the 'sim2real gap' (Boeing & Bräunl, 2012; Koos et al., 2010), in the context of UAVs, training the RL policy on-device during flight has logistical limitations and not scalable (As explained in Sect. 2). Moreover, some of the onboard computers on these UAVs don't have the necessary hardware resources required for on-device RL training. For instance, most hobbyist drones and research UAV platforms (e.g., CrazyFlie) are typically powered by microcontrollers and have a total of 1 MB memory. For most vision based navigation, its storage space is insufficient for the policy weights. Hence these resource constraints make RL training on-device extremely challenging.

To overcome these resource constraint limitations and enable on-device RL training for UAVs, we introduce a methodology that uses HIL for training the RL policy. This methodology allows us to train the RL policy on a high-end machine (e.g., Intel Core-i9 with GPUs) while capturing the latencies incurred in processing the policy in the onboard computer. We describe the details of the methodology below.

8.1 Methodology

The methodology is divided into three phases, namely 'Phase 1', 'Phase 2', and 'Phase 3' as shown in Fig. 12. In Phase 1 (Fig. 12a), we use the HIL to determine three specific latencies namely t_1 , t_2 , and t_3 defined in Sect. 7.1. We capture the latency distribution when the policies are run on-device (e.g., Ras-Pi). The distribution captures the variation in the decision-making times when the policy is deployed in the onboard computer.

Once the latency distribution is captured, we calculate the maximum achievable velocities for safe navigation based on the decision-making time (Liu et al., 2016). This is to ensure that the drone can navigate safely without colliding with an obstacle. We evaluate the maximum safe velocity the aerial robot can travel based on the visual

Fig. 12 A three-phase methodology for mitigating the hardware gap using hardware-in-the-loop training. **a** In phase 1, we use the hardware-in-the-loop methodology on a candidate policy to get the policy's latency distribution on target hardware (Ras-Pi 4). We use prior work (Krishnan et al., 2020) as the cyber-physical model to determine the upper bound for maximum velocity. **b** In phase 2, we use the latency distribution to randomly sample the delay that needs to be added to the policy's training. **c** In phase 3, the HIL trained policy is deployed on the target hardware for evaluation

performance model proposed in this work (Krishnan et al., 2020). The model considers the time to action latency, drones' physics (e.g., thrust-to-weight ratio, sensing distance, etc.) to determine the drone's maximum safe velocity.

In phase 2 (Fig. 12b), we train the policy by adding extra delays sampled from the latency distribution determined in phase 1. The decision-making loop's added delays mimic the typical processing delay when the policy is deployed on the resource-constrained onboard computer. The policy's action space is also scaled based on the maximum velocity achievable based on the decision-making time (Krishnan et al., 2020; Liu et al., 2016).

Once the policy is trained, in phase 3 (Fig. 12c), we deployed the trained policy on the onboard compute (Ras-Pi) and evaluate its performance and quality of flight metrics.

8.2 Experimental setup and evaluation

To validate the methodology, we train a policy on the Static Obstacle environment with at most two to three obstacles. The candidate architecture policy is 5 Layers with 32 Filters based on the template defined in Fig. 7.

We use the HIL setup described in Fig. 9 to evaluate the decision making latency on Ras-Pi 4, which is our target resource-constrained hardware platform. The simulation environment is rendered on the Intel Core-i9 server. We deploy a randomly initialized policy on the Ras-Pi 4 at this stage to benchmark the latencies. We do a rollout of 1000 steps using HIL to capture the variations in decision-making times.

On the high-end server (Intel core-i9 with GTX 2080 TI), we train the candidate policy for the same task (i.e., Static Obstacles) with added delay element in the decision-making loop. The delay element's actual value is randomly sampled from the latency distribution obtained for the candidate policy (5 Layers with 32 Filters) running on the Ras-Pi 4. Also, based on the maximum value of the latency from the distribution, we estimate the upper limit for the safe velocity for drone (Krishnan et al., 2020; Liu et al., 2016). This upper limit in safe velocity is then used to scale the action space such that at any point, the drone's velocity at each step does not exceed the maximum safe velocity.

Once the candidate policy's training with added latency is complete, we deploy the policy on the Ras-Pi 4 platform (target resource-constrained onboard compute). We use the HIL methodology to evaluate the quality of flight metrics on Ras-Pi 4. The comparison in trajectories between Core-i9 and Ras-Pi 4 is shown in Fig. 13. The two trajectories are very similar to each other and do not suffer from larger drifts that were seen before. Table 7 compares the quality of flight metric. The performance gap (denoted by "Perf Gap") is reduced from 38% to less than 0.5% on the flight time metric, 16.03% to 1.37% on the trajectory length metric, and 15.49% to 0.1% on the energy of flight metric.

In summary, we show that training policy with added delay to mimic the target platform can be used to minimize the hardware gap and performance difference between the training machine and the resource-constrained onboard compute.

Fig. 13 Comparison of trajectory for a policy that uses mitigation technique (denoted by the label “With mitigation”) with the policy that does not use mitigation technique (represented by the label “Without mitigation”). The policy trained on the training machine (denoted by the label “HIL”) is also plotted for comparison. Using the mitigation technique, we reduced the trajectory length degradation from 34.15 to 29.03 m (to within 1.37%)

9 Future work

The Air Learning toolset and benchmark that we built can be used for solving several open problems related to UAVs which spans multiple disciplines. The goal of this work was to demonstrate the breadth of Air Learning as an interdisciplinary tool. In the future, Air Learning can be used to address numerous other questions, including but not limited to the following.

Environments: In this work, we focus primarily on UAV navigation for indoor applications (Khosiawan & Nielsen, 2016). Future work can extend Air Learning’s environment generator to explore new robust reinforcement learning policies for UAV control under harsh environmental conditions. For instance, AirSim weather APIs can be coupled with Air Learning environment generator to explore new reinforcement learning algorithms for UAV control with different weather conditions.⁸

Algorithm design: Reinforcement algorithms are susceptible to hyperparameter tuning, policy architecture, and reward function. Future work could involve using techniques such as AutoML (Zoph et al., 2017) and AutoRL (Chiang et al., 2019) to determine the best hyperparameters, and explore new policy architectures for different UAV tasks.

Policy exploration: We designed a simple multi-modal policy and kept the policy architecture same across DQN and PPO agent. In future work, one could explore other types of policy architectures, such as LSTM (Bakker, 2002) and recurrent reinforcement learning (Li et al., 2015). Another future work could expand our work by exploring energy efficient policies by using the capability available in Air Learning to monitor energy consumption continuously. Energy-aware policies can be associated with open problems in mobile robots, such as charging station problem (Kundu & Saha, 2018).

System optimization studies: Future work on the system optimization can be classified into two categories. First, one can perform a thorough workload characterization for reducing the reinforcement learning training time. System optimizations will speed up the training process, thus allowing us to build more complex policies and strategies (OpenAI, 2018) for solving open problems in UAVs. Second, path to building custom hardware accelerators

⁸ AirSim plugin weather APIs can be found here: https://github.com/microsoft/AirSim/blob/master/PythonClient/computer_vision/weather.py.

Table 7 Evaluation of quality of flight between Ras-Pi 4 and Intel Core-i9 with and without mitigation

Metric	Intel Core i9 (Without mitigation)	Intel Core i9 (Training Machine)	Ras-Pi 4 (With mitigation)	Ras-Pi (Without mitigation)	Perf Gap (With mitigation)	Perf Gap (Without mitigation)
Inference latency (ms)	11	400	396	396	1	1
Success rate	85	85	83	73	2.4	16.43
<i>Quality of flight</i>						
Flight time (s)	24.08	32.62	32.44	44.93	0.5	37.73
Distance flown (m)	25.64	29.43	29.031	34.15	1.37	16.03
Energy (kJ)	19.09	24.59	24.57	28.40	0.1	15.49

After using the methodology to minimize the hardware gap, we were able to reduce the gap from 38% to less than 0.5% on the flight time, 16.03% to 1.37% on the trajectory length, and 15.49% to 0.1% on the energy of flight

to improve the onboard compute performance can be explored. Having specialized hardware onboard would allow better real-time performance for UAVs.

10 Conclusion

We present Air Learning, a deep RL gym and cross-disciplinary toolset, which enables deep RL research for resource constraint systems, and an end-to-end holistic applied RL research for autonomous aerial vehicles. We use Air Learning to compare the performance of two reinforcement learning algorithms namely DQN and PPO on a configurable environment with varying static and dynamic obstacles. We show that for an end to end autonomous navigation task, DQN performs better than PPO for a fixed observation inputs, policy architecture and reward function. We show that the curriculum learning based DQN agent has a better success rate compared to non-curriculum learning based DQN agent with the same number of experience (steps). We then use the best policy trained using curriculum learning and expose the difference in the behavior of aerial robot by quantifying the performance of the policy using HIL methodology on a resource-constrained Ras-Pi 4. We evaluate the performance of the best policy using quality of flight metrics such as flight time, energy consumed and total distance traveled. We show that there is a non-trivial behavior change and up to 40% difference in the performance of policy evaluated in high-end desktop and resource-constrained Ras-Pi 4. We also artificially degrade the performance of the high-end desktop where we trained the policy. We observe a similar variation in the trajectory as well as other QoF metrics as observed in Ras-Pi 4 thereby showing how the onboard compute performance can affect the behavior of policies when ported to real UAVs. We also show the impact of energy QoF on the success rate of the mission. Finally, we propose a mitigation technique using the HIL technique that minimizes the hardware gap from 38% to less than 0.5% on the flight time metric, 16.03% to 1.37% on the trajectory length metric, and 15.49% to 0.1% on the energy of flight metric.

Appendix

A Air Learning environment generator knobs

Here we list the parameters available in Air Learning environment generator in detail. These parameters are exposed as a game configuration file which can be modified by the end user during runtime. These parameters can also be part of the deep RL training setup where it can be changed before the onset of new episode.⁹

Arena size: The `Arena Size` is the total volume available in the environment. It is represented by `[length, width, height]` tuple. A large arena size means the UAV has to cover more distance in reaching the goal which directly impacts its energy and mission success. Figure 3a an arena size of 50 m x 50 m x 5 m. The arena can be customized by adding materials, which we describe in the “Materials” section.

Wall color: The `Wall Color` parameter can be used to set the wall colors of the arena. The parameter takes `[R, G, B]` tuple as input. By setting different values of `[R, G, B]`, any color in the visible spectrum can be applied to the walls. The neural network

⁹ <https://bit.ly/38WL2CA>.

policies show sensitivity towards different colors (Zeiler & Fergus, 2013) and varying these color during training can help the policy to generalize well.

Asset: An *Asset* in Air Learning is a mesh in UE4 (Epic, 2018). Any asset that is available in the project can be used as a static obstacle, dynamic obstacle, or both. At simulation startup, Air Learning uses these assets as either a static or dynamic obstacle. The number of assets that will be spawned in the arena will be equal to the `#Static Obstacle` and `#Dynamic Obstacle` parameter. By having the ability to spawn any asset as an obstacle, the UAV agent can generalize to avoid collision with different types of obstacle. Figure 3b shows some of the sample assets used in this work.

Number of obstacles: The `# Static Obstacles` is a parameter that describes the total number of static objects that is spawned in the environment. Figure 3c shows some of the assets used as random obstacles by the environment generator. Using this parameter, we can generate environments ranging from very dense to very sparse obstacles. Depending upon the value of this parameter, the navigation complexity can be easy or difficult. A large number of obstacles increases the collision probability and can be used for stressing the efficacy of reinforcement learning algorithms.

Minimum distance: The `Minimum distance` is a parameter that controls the minimum distance between two static objects in the arena. This parameter in conjunction with `# Static Obstacles` is what determines congestion.

Goal position: The `Goal Position` is a parameter that specifies the destination coordinate that the UAV must reach. The `Goal Position` coordinates should always be inside the arena, and there is error checking for input errors. Similar to `# Static Obstacles`, it increases task complexity.

Number of dynamic obstacles: The `# Dynamic Obstacles` is a parameter that describes the total number of obstacles that can move in the environment.

Velocity: The `Velocity` parameter is a tuple of the form $[V_{\min}, V_{\max}]$ that works with `# Dynamic Obstacles`. The environment generator randomly chooses a value from this range for the velocity of a dynamic obstacle. This coupled with the `# Dynamic Obstacles` helps control how dynamic and challenging the environment is for the aerial robot.

Seed: The `Seed` parameter is used for randomizing the different parameters in the environment. By setting the same ‘Seed’ value, we can reproduce (and randomize) the environment (obstacle position, goal position, etc.).

As mentioned previously, there is a second category of parameters that can be configured. These are not included in the configuration file. Instead, they are controlled by putting files into folders. Details about them are as follows.

Textures: A *Texture* is an image that is used on an UE4 asset (Games, 2018). They are mapped to the surfaces of any given asset. At startup, the environment generator applies textures to matching assets. Textures and materials (below point) help the training algorithm capture different object features, which is important to help the algorithm generalize.

Materials: A *Material* is a UE4 asset (Epic, 2018) that can be applied to meshes to control the visual look of the scene. Material is usually made of multiple textures to create a particular visual effect for the asset. At simulation startup, Air Learning environment generator applies materials to matching assets.

Materials can help training algorithms on two fronts. First, neural network policy has a sensitivity to capture various material features in the objects (Tobin et al., 2017; Zeiler & Fergus, 2013). For instance, the type of material affects how light interacts with the surface, and as a result, an RL based robot that is relying on images as input can learn different things (and act differently) under different materials and the textures that it observes. Second, they can make it challenging for the algorithms using image-based inputs. For

instance, shiny and transparent objects are harder to detect (Berger et al., 2017; Lai & Fuh, 2015).

In summary, Air Learning's environment generator allows any UE4 asset to be loaded into the project, and provides flexibility in the choice of obstacles, materials, and texture. These features are essential to provide a safe sandbox environment where to train and evaluate various deep reinforcement learning algorithms and policies that can generalize well.

B Algorithm exploration

We explore two RL algorithm types for end-to-end navigation task in autonomous UAVs. The choice of the seed algorithm we used in this work can be classified into discrete action algorithms and continuous action algorithm. For discrete action reinforcement learning algorithm, we use Deep Q Networks (DQN), and for the continuous action algorithm, we use Proximal Policy Optimization (PPO). For both these algorithms, we keep the observation space, policy architecture and reward structure same and compare agent performance.

Training methodology

The training methodology, policy architecture, reward function, and action space for PPO and DQN agent with and without curriculum learning is described below.

Non-curriculum learning: We train the DQN agent and PPO agent on the environment with static obstacles. To determine the baseline performance for both the algorithms, we train each agent to 1 Million steps using non-curriculum learning. For non-curriculum learning, we randomize the position of the goal and obstacles every episode to be anywhere in the arena. Simply put, the entire arena acts like one zone as shown in Fig. 14a. The checkpoints are saved every 50,000 steps and use the last saved checkpoint after 1 Million steps.

Fig. 14 Zoning used in the training methodology for curriculum learning and non-curriculum learning. Here we show the top view of our environment in wireframe mode (Games, 2018) available in UE4. **a** In non-curriculum learning, the end goal is randomly placed anywhere in the arena. Unlike curriculum learning, the entire arena is one zone. **b** In curriculum learning, we split the arena into virtual partitions, and the end goal is placed within a specific zone and gradually moved higher zone once it succeeds in more than 50% over 1000 latest episode

Curriculum learning: To improve the baseline performance for DQNs and PPO, we employ the curriculum learning (Bengio et al., 2009) approach where the goal position is progressively moved farther away from the starting point of the agent. To implement this, we divide the entire arena into multiple zones namely Zone 0, Zone 1 and Zone 2 as shown in Fig. 14b. Zone 0 corresponds to the region that is within 16 m from the UAV starting position and Zone 1 and Zone 2 are within 32 m and 48 m respectively. Initially, the position of goal for the UAV is determined randomly such that the goal position lies within Zone 0. Once the UAV agent achieves 50% success over a rolling window of past 1000 episodes, the position of the goal expands to Zone 1 and so forth. To make sure that the agent does not forget learning in the previous zone, the goal position in the next zone is inclusive of previous zones. We train the agent to progress until Zone 2. Both the agents (PPO and DQN) are trained for 1 Million steps. We checkpoint the policy at every zone so that it can be evaluated on how well it has learned to navigate across all three zones.

Policy architecture: The policy architecture for both PPO and DQN agent used is multi-modal in nature. It receives depth image, velocity vector (V_t) and position vector (X_t) as inputs as shown in Fig. 15. The V_t is a 1-dimensional vector of the form $[v_x, v_y, v_z]$ where v_x, v_y, v_z are the components of velocity vector in x, y and z directions at time ' t '. The X_t is a 1-dimensional vector of the form $[X_{goal}, Y_{goal}, D_{goal}]$, where X_{goal} is the difference in the x -coordinate of the goal and x -coordinate of the agent's current position, Y_{goal} is the difference in the y -coordinate of the goal and y -coordinate of the agent's current position and D_{goal} is the euclidean distance to the goal from the agent's current position.

The depth image is processed by four convolutions layers whose filter depth and size are 32 (4×4), 64 (4×4), 128 (2×2), and 64 (1×1) respectively. As an example, in a 32 (4×4) filter, 32 is the depth of the filter and (4×4) is the size of the filter. The fourth layer's output is flattened and concatenated with the velocity vector (V_t) and position vector (X_t). The combined inputs are then fed to three layers of fully connected layers with 256 hidden units each. The action space for the agent determines the number of hidden units in the final fully connected layer. For the DQN agent, we have twenty-five discrete actions whereas, for PPO agent, we have two actions. Hence, the final layer for the DQN agent has twenty-five hidden units, and PPO agent has two hidden units.

Fig. 15 The network architecture for the policy in the PPO and DQN agents. Both the agents take a depth image, velocity vector, and position vector as inputs. The depth image has four layers of convolutions after which the results are concatenated with the velocity and position vectors. In a 32 (4×4) convolution filter, 32 is the depth of the filter and (4×4) is the filter size. The combined vector space is applied to the three layers of a fully connected network, each with 256 hidden units. The action space determines the number of hidden units in the last fully connected layer. For DQN agent, we have twenty-five actions, and for PPO agent we have two actions which control the velocity of the UAV agent in X and Y direction

For DQN agent, the activation used for all convolution and the fully connected layer is *ReLU*, and for PPO agent, we use *ReLU* except for the last layer where we use *Tanh* for producing continuous values.

Action space: The action space for DQN consists of twenty-five discrete actions. Out of these twenty-five action spaces, ten actions are for moving forward with different fixed velocities ranging from 1 to 5 m/s, five actions are for moving backward, five actions for yawing right with fixed yaw rates of 108 °, 54 °, 27 °, 13.5 ° and 6.75 ° and another five actions for yawing left with fixed yaw rates of − 216 °, − 108 °, − 54 °, − 27 ° and − 13.5 °. At each time step, the policy takes observation space as inputs and outputs one of the twenty-five actions based on the observation. The high-level actions are mapped to low-level flight commands using the flight controller shown in Fig. 5 and as it is implemented.¹⁰

The action space for PPO on the other hand consists of velocity components v_x (velocity in x -direction) and v_y (velocity in y -direction). At each time step, the policy takes observation as the input and generates continuous values for v_x and v_y . The values of v_x and v_y are scaled such that values of the magnitude of velocity lie anywhere between 1 and 5 m/s. We use the `MaxDegreeOfFreedom` option in the AirSim API that calculates the yaw rates automatically to make sure the drone is pointed in the direction it moves.

Reward: The reward function for both PPO agent and DQN agent are kept the same and is defined as follows.

$$r = 1000 * \alpha - 1000 * \beta - D_g + D_c * \gamma \quad (2)$$

α is a binary variable where ‘1’ denotes if the goal is reached else it is ‘0’. β is also a binary variable where ‘1’ denotes if there is a collision with walls, obstacles or ground else it is ‘0’. D_g is the distance to the goal at any time steps from the agents’ current position. If the agent is going away from the goal, the distance to the goal increases thus penalizing the agent. γ is also a binary variable which is set to ‘1’ if the agent is closer to the goal. D_c is the distance correction which is applied to penalize the agent if it chooses actions which speed up the agent away from the goal. The distance correction term is defined as follows:

$$D_c = (V_{max} - V_{now}) * t_{max} \quad (3)$$

V_{max} is the maximum velocity possible for the agent which for DQN is fixed at 5 m/s and for PPO the outputs are scaled to lie between 1 and 5 m/s. V_{now} is the current velocity of the agent and t_{max} is the duration of the actuation.

Policy architecture vs runtime latency tradeoffs

Air Learning HIL can also be used to understand the tradeoff between the policy selection and the onboard hardware. In this section, we study the latency tradeoffs for various policies trained for point-to-point navigation policies in No Obstacle, Static Obstacle, and Dynamic Obstacle environments.

Figure 16 shows the latency tradeoff between the size of the policy and the latency to run on Ras-Pi 4. As the policy becomes wider/deeper, we can see that it increases the policy execution time, translating to increased decision making time. Hence while selecting a policy architecture, one must also account for the hardware latency.

¹⁰ https://microsoft.github.io/AirSim/docs/simple_flight/.

Fig. 16 **a** Understanding the trade-offs in latencies between various policy trained for No obstacles and Static obstacles environment. **b** Latencies for various policies trained for Dynamic Obstacles environment. The latencies are averaged over 1000 runs

Acknowledgements The effort at Harvard University and The University of Texas at Austin was sponsored by support from Intel.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., & Zheng, X. (2015). TensorFlow: Large-scale machine learning on heterogeneous systems. Software available from www.tensorflow.org
- Adiprawita, W., Ahmad, A. S., & Semibiring, J. (2008). Hardware in the loop simulator in UAV rapid development life cycle. *CoRR*, vol. abs/0804.3874.
- Ahn, M., Zhu, H., Hartikainen, K., Ponte, H., Gupta, A., Levine, S., & Kumar, V. (2020). Robel: Robotics benchmarks for learning with low-cost robots. In *Conference on robot learning* (pp. 1300–1313). PMLR.
- Bakker, B. (2002). Reinforcement learning with long short-term memory. *Advances in Neural Information Processing Systems*, 5, 1475–1482.
- Bellemare, M. G., Naddaf, Y., Veness, J., & Bowling, M. (2015). The arcade learning environment: An evaluation platform for general agents. In *Proceedings of the 24th international conference on artificial intelligence, IJCAI'15* (pp. 4148–4152). AAAI Press.
- Bengio, Y., Louradour, J., Collobert, R., & Weston, J. (2009). Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning* (pp. 41–48). ACM.
- Berger, K., Voorhies, R., & Matthies, L. H. (2017). Depth from stereo polarization in specular scenes for urban robotics. In *2017 IEEE international conference on robotics and automation (ICRA)* (pp. 1966–1973). IEEE.
- Boeing, A., & Bräunl, T. (2012). Leveraging multiple simulators for crossing the reality gap. In *2012 12th international conference on control automation robotics & vision (ICARCV)* (pp. 1113–1119). IEEE.

- Bojarski, M., Testa, D. D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., Jackel, L. D., Monfort, M., Muller, U., Zhang, J., Zhang, X., Zhao, J., & Zieba, K. (2016). End to end learning for self-driving cars. *CoRR*, vol. abs/1604.07316.
- Boroujerdian, B., Genc, H., Krishnan, S., Cui, W., Faust, A., & Reddi, V. (2018). Mavbench: Micro aerial vehicle benchmarking. In *2018 51st annual IEEE/ACM international symposium on micro-architecture (MICRO)* (pp. 894–907). IEEE.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., & Zaremba, W. (2016). Openai gym. *CoRR*, vol. abs/1606.01540.
- Chiang, H.-T.L., Faust, A., Fiser, M., & Francis, A. (2019). Learning navigation behaviors end-to-end with autorl. *IEEE Robotics and Automation Letters*, 4(2), 2007–2014.
- Chiang, H. L., Faust, A., Fiser, M., & Francis, A. (2019). Learning navigation behaviors end-to-end with autorl. *IEEE Robotics and Automation Letters*, 4, 2007–2014.
- Chollet, F. (2015). Keras. <https://github.com/fchollet/keras>
- Crazyflie. (2018). Crazyflie 2.0. <https://www.bitcraze.io/crazyflie-2/>
- DJI. (2018). DJI-mavic pro. <https://www.dji.com/mavic>
- Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., & Koltun, V. (2017). CARLA: An open urban driving simulator. In *Proceedings of the 1st annual conference on robot learning* (pp. 1–16).
- Duisterhof, B. P., Krishnan, S., Cruz, J. J., Banbury, C. R., Fu, W., Faust, A., de Croon, G. C. H. E., & Reddi, V. J. (2019). Learning to seek: Autonomous source seeking with deep reinforcement learning onboard a nano drone microcontroller. *CoRR*, vol. abs/1909.11236.
- Epic, G. (2018). Ue4 materials. <https://docs.unrealengine.com/en-US/Engine/Basics/AssetsAndPackages>
- Epic, G. (2018). Wire frame. <https://docs.unrealengine.com/en-us/Engine/Rendering/Materials>
- Faust, A., Palunko, I., Cruz, P., Fierro, R., & Tapia, L. (2017). Automated aerial suspended cargo delivery through reinforcement learning. *Artificial Intelligence*, 247, 381–398.
- Games, E. (2018). Ue4 textures. <https://docs.unrealengine.com/en-us/Engine/Content/Types/Textures>
- Games, E. (2018). Wire frame. <https://docs.unrealengine.com/en-us/Engine/UI/LevelEditor/Viewports/ViewModes>
- Gandhi, D., Pinto, L., & Gupta, A. (2017). Learning to fly by crashing. *CoRR*, vol. abs/1704.05588.
- Giusti, A., Guzzi, J., Ciresan, D. C., He, F.-L., Rodríguez, J. P., Fontana, F., et al. (2016). A machine learning approach to visual perception of forest trails for mobile robots. *IEEE Robotics and Automation Letters*, 1(2), 661–667.
- Goodchild, A., & Toy, J. (2018). Delivery by drone: An evaluation of unmanned aerial vehicle technology in reducing CO2 emissions in the delivery service industry. *Transportation Research Part D: Transport and Environment*, 61, 58–67.
- Gu, S., Holly, E., Lillicrap, T. P., & Levine, S. (2016). Deep reinforcement learning for robotic manipulation. *CoRR*, vol. abs/1610.00633.
- Ha, S., Kim, J., & Yamane, K. (2018). Automated deep reinforcement learning environment for hardware of a modular legged robot. In *2018 15th international conference on ubiquitous robots (UR)* (pp. 348–354). IEEE.
- Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P., & Levine, S. (2018). Soft actor-critic algorithms and applications. arXiv preprint [arXiv:1812.05905](https://arxiv.org/abs/1812.05905)
- Hill, A., Raffin, A., Ernestus, M., Gleave, A., Traore, R., Dhariwal, P., Hesse, C., Klimov, O., Nichol, A., Plappert, M., Radford, A., Schulman, J., Sidor, S., & Wu, Y. (2018). Stable baselines. <https://github.com/hill-a/stable-baselines>
- Hummingbird, A. (2018). Asctec hummingbird. <http://www.asctec.de/en/uav-uas-drones-rpas-roav/asctec-hummingbird/>
- Hwangbo, J., Sa, I., Siegwart, R., & Hutter, M. (2017). Control of a quadrotor with reinforcement learning. *IEEE Robotics and Automation Letters*, 2(4), 2096–2103.
- Intel. (2018). Intel aero ready to fly drone. <https://www.intel.com/content/www/us/en/products/drones/aero-ready-to-fly.html>
- Judah, K., Fern, A. P., Tadepalli, P., Goetschalckx, R. (2014). Imitation learning with demonstrations and shaping rewards. In *Twenty-eighth AAAI conference on artificial intelligence*.
- Kalashnikov, D., Irpan, A., Pastor, P., Ibarz, J., Herzog, A., Jang, E., Quillen, D., Holly, E., Kalakrishnan, M., Vanhoucke, V., & Levine, S. (2018). Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation. arXiv preprint [arXiv:1806.10293](https://arxiv.org/abs/1806.10293)
- Khosiawan, Y., & Nielsen, I. (2016). A system of uav application in indoor environment. *Production & Manufacturing Research*, 4(1), 2–22.
- Kjell, K. (2018). Airgym. <http://github.com/Kjell-K/AirGym>
- Koch, W., Mancuso, R., West, R., & Bestavros, A. (2018). Reinforcement learning for uav attitude control.

- Koenig, N., & Howard, A. (2004). Design and use paradigms for gazebo, an open-source multi-robot simulator. *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 3, 2149–2154.
- Koos, S., Mouret, J.-B., & Doncieux, S. (2010). Crossing the reality gap in evolutionary robotics by promoting transferable controllers. In *Proceedings of the 12th annual conference on genetic and evolutionary computation* (pp. 119–126). ACM.
- Kretchmar, R. M. (2000). *A synthesis of reinforcement learning and robust control theory*. Colorado State University Fort Collins.
- Krishnan, S., Wan, Z., Bhardwaj, K., Whatmough, P., Faust, A., Wei, G.-Y., et al. (2020). The sky is not the limit: A visual performance model for cyber-physical co-design in autonomous machines. *IEEE Computer Architecture Letters*, 19(1), 38–42.
- Kumar, K. R., Sastry, V., Sekhar, O. C., Mohanta, D., Rajesh, D., & Varma, M. P. C. (2016). Design and fabrication of coulomb counter for estimation of soc of battery. In *2016 IEEE international conference on power electronics, drives and energy systems (PEDES)* (pp. 1–6). IEEE.
- Kundu, T., & Saha, I. (2018). Charging station placement for indoor robotic applications. In *2018 IEEE international conference on robotics and automation (ICRA)* (pp. 3029–3036). IEEE.
- Lai, P.-J., & Fuh, C.-S. (2015). Transparent object detection using regions with convolutional neural network. In *IPPR conference on computer vision, graphics, and image processing* (pp. 1–8).
- Li, X., Li, L., Gao, J., He, X., Chen, J., Deng, L., & He, J. (2015). Recurrent reinforcement learning: A hybrid approach. *CoRR*, vol. abs/1509.03044.
- Liu, S., Watterson, M., Tang, S., & Kumar, V. (2016). High speed navigation for quadrotors with limited onboard sensing. In *2016 IEEE international conference on robotics and automation (ICRA)* (pp. 1484–1491). IEEE.
- Locobot. (2018). An open source low cost robot. <http://www.locobot.org/>.
- Mahmood, A. R., Korenkevych, D., Komer, B. J., & Bergstra, J. (2018). Setting up a reinforcement learning task with a real-world robot. *CoRR*, vol. abs/1803.07067.
- Mahmood, A. R., Korenkevych, D., Vasani, G., Ma, W., & Bergstra, J. (2018). Benchmarking reinforcement learning algorithms on real-world robots. *CoRR*, vol. abs/1809.07731.
- Menard, M., & Wagstaff, B. (2015). *Game development with unity*. Nelson Education.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., & Riedmiller, M. (2013). Playing atari with deep reinforcement learning. arXiv preprint [arXiv:1312.5602](https://arxiv.org/abs/1312.5602)
- Murali, A., Chen, T., Alwala, K. V., Gandhi, D., Pinto, L., Gupta, S., & Gupta, A. (2019). Pyrobot: An open-source robotics framework for research and benchmarking. arXiv preprint [arXiv:1906.08236](https://arxiv.org/abs/1906.08236)
- NVIDAA-AI-IOT. (2015). NVIDIA-AI-IOT/redtail. <https://github.com/NVIDIA-AI-IOT/redtail/wiki/Sky-pad-TBS-Discovery-Setup>
- NVIDIA. (2019). NVIDIA Xavier. <https://developer.nvidia.com/embedded/buy/jetson-agx-xavier-devkit>
- OpenAI. (2018). Openai five. <https://blog.openai.com/openai-five/>
- Palacin, J., Palleja, T., Valganón, I., Pernia, R., & Roca, J. (2005). Measuring coverage performances of a floor cleaning mobile robot using a vision system. In *Proceedings of the 2005 IEEE international conference on robotics and automation* (pp. 4236–4241). IEEE.
- Parrot. (2019). Parrot bebop-2. <https://www.parrot.com/us/drones/parrot-bebop-2-fpv?ref=#parrot-bebop-2-fpv-details>
- Peng, K., Feng, L., Hsieh, Y., Yang, T., Hsiung, S., Tsai, Y., & Kuo, C. (2017). Unmanned aerial vehicle for infrastructure inspection with image processing for quantification of measurement and formation of facade map. In *2017 international conference on applied system innovation (ICASI)* (pp. 1969–1972). IEEE.
- Plappert, M. (2016). Keras-rl. <https://github.com/keras-rl/keras-rl>
- Quillen, D., Jang, E., Nachum, O., Finn, C., Ibarz, J., & Levine, S. (2018). Deep reinforcement learning for vision-based robotic grasping: A simulated comparative evaluation of off-policy methods. *CoRR*, vol. abs/1802.10264.
- Riedmiller, M. (2012). 10 steps and some tricks to set up neural reinforcement controllers. In *Neural networks: Tricks of the trade* (pp. 735–757). Springer.
- Sadeghi, F., & Levine, S. (2016). (cad)2rl: Real single-image flight without a single real image. *CoRR*, vol. abs/1611.04201.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. *CoRR*, vol. abs/1707.06347.
- Shah, S., Dey, D., Lovett, C., & Kapoor, A. (2017). Air Sim: High-fidelity visual and physical simulation for autonomous vehicles. *CoRR*, vol. abs/1705.05065.
- Su, P.-H., Vandyke, D., Gasic, M., Mrksic, N., Wen, T.-H., & Young, S. (2015). Reward shaping with recurrent neural networks for speeding up on-line policy learning in spoken dialogue systems. arXiv preprint [arXiv:1508.03391](https://arxiv.org/abs/1508.03391)

- Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., & Abbeel, P. (2017). Domain randomization for transferring deep neural networks from simulation to the real world. *CoRR*, vol. abs/1703.06907.
- Todorov, E., Erez, T., & Tassa, Y. (2012). Mujoco: A physics engine for model-based control. In *2012 IEEE/RISJ international conference on intelligent robots and systems* (pp. 5026–5033).
- Travnik, J. B., Mathewson, K. W., Sutton, R. S., & Pilarski, P. M. (2018). Reactive reinforcement learning in asynchronous environments. *Frontiers in Robotics and AI*, 5, 79.
- Tseng, C., Chau, C., Elbassioni, K. M., & Khonji, M. (2017). Flight tour planning with recharging optimization for battery-operated autonomous drones. *CoRR*, vol. abs/1703.10049.
- Valcasara, N. (2015). *Unreal engine game development blueprints*. Packt Publishing Ltd.
- Waharte, S., & Trigoni, N. (2010). Supporting search and rescue operations with uavs. In *2010 international conference on emerging security technologies* (pp. 142–147). IEEE.
- Wu, B., Chen, W., Fan, Y., Zhang, Y., Hou, J., Liu, J., Huang, J., Liu, W., & Zhang, T. (2019). Tencent ml-images: A large-scale multi-label image database for visual representation learning. *CoRR*, vol. abs/1901.01703.
- Yahya, A., Li, A., Kalakrishnan, M., Chebotar, Y., & Levine, S. (2016). Collective robot reinforcement learning with distributed asynchronous guided policy search. *CoRR*, vol. abs/1610.00673.
- Zeiler, M. D., & Fergus, R. (2013). Visualizing and understanding convolutional networks. *CoRR*, vol. abs/1311.2901.
- Zoph, B., Vasudevan, V., Shlens, J., & Le, Q. V. (2017). Learning transferable architectures for scalable image recognition. *CoRR*, vol. abs/1707.07012.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Srivatsan Krishnan¹ · Behzad Boroujerdian² · William Fu¹ · Aleksandra Faust³ · Vijay Janapa Reddi^{1,2}

¹ Harvard University, Cambridge, USA

² The University of Texas at Austin, Austin, USA

³ Google Research, Mountain View, USA