

Learning to Fly: Computational Controller Design for Hybrid UAVs with Reinforcement Learning

JIE XU, Massachusetts Institute of Technology
TAO DU, Massachusetts Institute of Technology
MICHAEL FOSHEY, Massachusetts Institute of Technology
BEICHEN LI, Massachusetts Institute of Technology
BO ZHU, Dartmouth College
ADRIANA SCHULZ, University of Washington
WOJCIECH MATUSIK, Massachusetts Institute of Technology

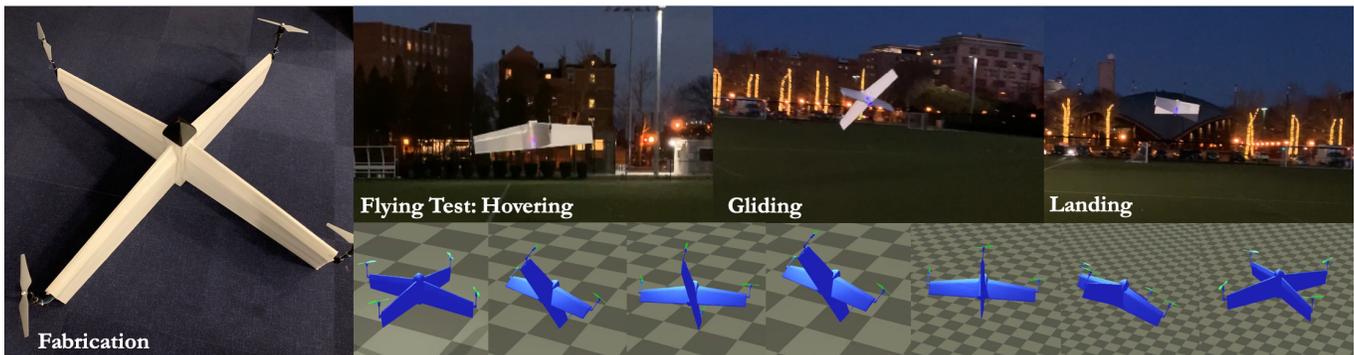


Fig. 1. We propose a computational approach to automatically design a mode-free, model-agnostic neural network controller for customized hybrid UAVs (left). Our novel reinforcement learning framework enables the training of a controller in a stochastically enhanced physics simulator (right bottom) and the crossing of the gap between simulation and reality (right top).

Hybrid unmanned aerial vehicles (UAV) combine advantages of multicopters and fixed-wing planes: vertical take-off, landing, and low energy use. However, hybrid UAVs are rarely used because controller design is challenging due to its complex, mixed dynamics. In this paper, we propose a method to automate this design process by training a mode-free, model-agnostic neural network controller for hybrid UAVs. We present a neural network controller design with a novel error convolution input trained by reinforcement learning. Our controller exhibits two key features: First, it does not distinguish among flying modes, and the same controller structure can be used for copters with various dynamics. Second, our controller works for real models without any additional parameter tuning process, closing the gap between virtual simulation and real fabrication. We demonstrate the efficacy of the proposed controller both in simulation and in our custom-built hybrid

UAVs (Figure 1, 8). The experiments show that the controller is robust to exploit the complex dynamics when both rotors and wings are active in flight tests.

CCS Concepts: • **Computing methodologies** → **Control methods**; **Reinforcement learning**; **Physical simulation**.

Additional Key Words and Phrases: hybrid UAVs, neural network controllers

ACM Reference Format:

Jie Xu, Tao Du, Michael Foshey, Beichen Li, Bo Zhu, Adriana Schulz, and Wojciech Matusik. 2019. Learning to Fly: Computational Controller Design for Hybrid UAVs with Reinforcement Learning. *ACM Trans. Graph.* 38, 4, Article 42 (July 2019), 12 pages. <https://doi.org/10.1145/3306346.3322940>

Authors' addresses: Jie Xu, Massachusetts Institute of Technology, jiex@csail.mit.edu; Tao Du, Massachusetts Institute of Technology; Michael Foshey, Massachusetts Institute of Technology; Beichen Li, Massachusetts Institute of Technology; Bo Zhu, Dartmouth College; Adriana Schulz, University of Washington; Wojciech Matusik, Massachusetts Institute of Technology.

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

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.
0730-0301/2019/7-ART42 \$15.00
<https://doi.org/10.1145/3306346.3322940>

1 INTRODUCTION

Multicopters are becoming increasingly popular due to their flight flexibility and stable hovering capability. Fixed-wing airplanes, on the other hand, are more energy efficient during level flight, making them better vehicles for long distance flights. In order to leverage the advantages of both designs, new hybrid UAVs which equip multicopters with a pair of fixed wings have piqued the interest of the aircraft community [Bapst et al. 2015; Hugh Stone et al. 2008; Oosedo et al. 2013; Ritz and D'Andrea 2017]. These vehicles can quickly switch between hover and flight modes, allowing for highly stable, energy efficient flights. There are many ways to combine

multicopters and fixed wings, resulting in hybrid UAVs with vastly different configurations.

Controlling such hybrid designs, however, is difficult because the aerodynamics changes dramatically during the flight. In typical designs, experts manually design controllers for both the *copter flight mode* and the *plane flight mode*, as well as controllers for transitioning between these modes. Designing controllers for the transition phase is especially challenging due to the complexity of the dynamics when rotors and wings are both active. Furthermore, due to the vast differences between various configurations of hybrid UAVs, controllers cannot be easily transferred between designs. In order to construct a new hybrid UAV, engineers have to redesign the controller architecture and its parameters. This makes the whole process extremely labor-intensive and therefore only a small fraction of the possible hybrid UAVs design space has been explored so far.

To address these challenges, we propose a method that automatically computes a neural network controller for a hybrid UAV. Compared with the traditional mode-based, model-specific controllers, our *unified* controller has the following immediate benefits: first, our mode-free, velocity-tracking controller works directly with the full dynamics during the flight. Our controller does not need to differentiate between the copter and flight modes or explicitly deal with the transition between modes. For example, the controller will automatically orient a tail-sitter hybrid UAV purely based on the input velocity – it will set it to a copter orientation for lower velocities and a plane orientation for higher velocities. Second, the architecture of our model-agnostic neural network controller can be used for almost any hybrid UAV configuration. In the future, this could allow knowledge acquired in training one hybrid UAV model to be directly transferred to another design’s controller.

A natural choice to train neural network controllers is using reinforcement learning, since these algorithms have been successfully applied on models with complex dynamics [Heess et al. 2017; Schulman et al. 2015, 2017; Won et al. 2017]. However, while this approach can be easily applied to train hybrid UAVs in a simulation environment, using these controllers for real flights is known to be challenging due to the gap between simulation and reality. To close this gap, we propose a novel error integral block in the state vector for the input of our network, enabling the elimination of steady-state drift errors that are challenging to remove manually. Additionally, we also devise a simulation infrastructure enhanced by noise and latency patterns and a reinforcement learning framework with novel reward functions for hybrid UAVs. We demonstrate the capabilities of our method by fabricating real hybrid UAVs and showing that our controller is robust for real flights.

In summary, our work has the following technical contributions:

- we present an automatic way for designing a *mode-free, model-agnostic* neural network controller for real hybrid UAVs.
- we propose a method to improve the robustness of the trained neural network controllers for hybrid UAVs and close the reality gap between simulation and real flight.

2 RELATED WORK

Computational Design for Fabrication. Computational design of complex functional mechanisms is a growing research field. This

body of work aims at (1) making it more efficient for experts to explore design spaces and find unprecedented solutions and (2) lowering the barrier of entry for design by casual users, enabling a new age of customization. Recent advances in the field propose new interactive techniques to design mechanical characters [Coros et al. 2013], robotic creatures [Megaro et al. 2015], creatures with wheels and legs [Geilinger et al. 2018], and multicopters [Du et al. 2016]. Our method builds upon this research and proposes a tool for designing hybrid UAVs by mixing and matching parts from an expert-designed collection [Desai et al. 2017; Schulz et al. 2017].

Computational design tools guide users in creating objects based on how they should function once manufactured. Therefore, the tools need to support *performance-driven* design exploration [Umetani et al. 2012] or design optimization [Bharaj et al. 2015; Umetani et al. 2014]. Performance evaluations are typically a bottleneck for these methods. This process is particularly challenging for cyberphysical systems whose performance depends not only on geometric considerations but also on motion planning and control. In the context of a hybrid UAV, measuring its flight performance requires designing an adequate controller in addition to designing its structure. Until now, the design process for hybrid UAV controllers had to be done manually due to the complexity of system dynamics. Our paper bridges this gap, proposing a new automatic method that can design controllers for a large variety of hybrid UAV designs.

Hybrid UAVs. Hybrid UAVs have demonstrated outstanding aerodynamic performance and energy saving capabilities. However, they also bring new challenges in controller design. A large variety of dynamic models and control algorithms have been developed to accommodate UAVs with different wing configurations (see e.g. [Cory and Tedrake 2008]). For example, a simple flat-plate wing model coupled with a TVLQR controller has demonstrated success in enabling precise control of a fixed-wing drone with fast locomotion and varying trajectories [Barry 2016]. This design supports a single flying mode and multiple predefined trajectories, relying on online search to select the best control outputs from the precomputed data.

Controlling a multi-mode UAV requires a complicated strategy due to the differences between flying modes, which rely on distinct flying mechanisms. A straightforward strategy is to aggregate multiple controllers with each one focused on a single flying mode. These modes can be organized as a state transition graph [Hugh Stone et al. 2008] or a cascading hierarchy [Bapst et al. 2015; Ritz and D’Andrea 2017]. A number of intermediate modes need to be incorporated to enable smooth transitions between different modes. A typical example of this approach is the design of a tail-sitter drone, which takes off vertically, hovers in the air statically, and glides horizontally [Oosedo et al. 2013]. Such methods produce reliable control mechanisms but lack applicability to different drone designs. Relatively little attention has been paid to a universal controller that can be applied to different UAV models.

Reinforcement Learning. Deep reinforcement learning has become a prevalent tool for training continuous controllers for a variety of robots to finish highly dynamic tasks in animation. Some recent successful examples include walkers [Heess et al. 2017], humanoid robots with acrobatic or athletic skills [Liu and Hodgins 2018; Peng et al. 2018a, 2017], robotic manipulators [Devin et al. 2017], and

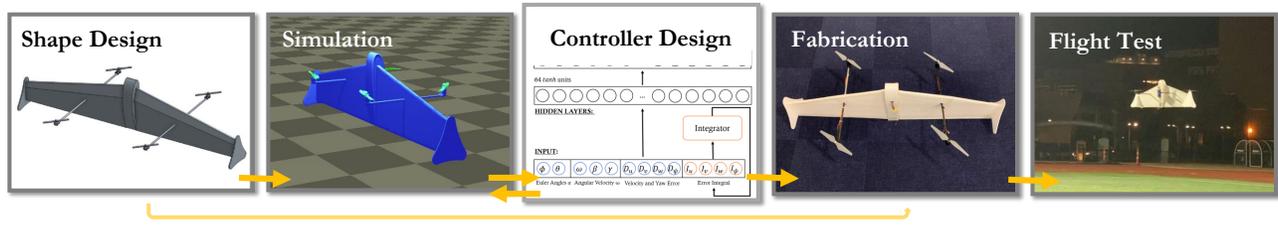


Fig. 2. In our system, users first design the geometry of a hybrid UAV (left most). The design is then used in our realistic simulator (middle left) to compute the corresponding neural network controller using reinforcement learning (middle). Once the training is done, we construct the corresponding hybrid UAV based on the geometry and controller specifications (middle right). Finally, we verify the UAV performance in real flight tests (right).

aerial robots [Won et al. 2017, 2018]. The work most related to ours is [Won et al. 2018], which trained a flying dragon actuated by flapping wings to follow a trajectory. Their flying robot modelled aerodynamic effects between wings and airflow, but the method was not tested on hardware platforms.

While reinforcement learning has produced convincing results in animation, successful demonstrations on real robots are rare due to discrepancies between simulated and real environments. A few notable papers have managed to close this gap for specific types of robots through domain randomization [Andrychowicz et al. 2018; Sadeghi and Levine 2017; Tan et al. 2018; Tobin et al. 2017], more accurate modeling [Bansal et al. 2016; Fu et al. 2016], or better sensing [Hwangbo et al. 2017]. Designing controllers for hybrid UAVs, however, poses new challenges not addressed in the previous work—these vehicles have a floating base, highly dynamic motion, and unsteady aerodynamic effects. Our experiments show that existing techniques are not sufficient for training a robust controller for a real hybrid UAV, which motivates us to propose neural network controllers with new inputs and a novel integral block. Moreover, controllers in most of the prior research worked for a specific robot design, while we present a method that works robustly for a variety of hybrid UAVs with significantly different topology and dynamics.

Our method is also related to [Peng et al. 2018b] in that both propose neural network controllers with time-dependent blocks to capture the mismatch between simulated and real environments. Comparing to the LSTM network controller in their method which uses a memory-based policy in a recurrent block, we encode the history information in a concise integral block, which helps reduce the workload of the training process significantly.

Our work draws inspiration from works on controlling real helicopters using reinforcement learning [Abbeel et al. 2007; Kim et al. 2004] and learning from demonstrations [Abbeel et al. 2010; Coates et al. 2008]. All the prior work and ours show controllers on real UAVs; however, the scopes are quite different. In the prior work, a functional controller on hardware is provided for a pilot expert to demonstrate flight patterns, then the method either learns an accurate dynamics model or learns an improved controller from the flight data. Our work is different in that we design controllers directly from the inaccurate dynamics of the hybrid UAVs without needing a real, working controller beforehand.

3 SYSTEM OVERVIEW

As shown in Figure 2, we propose a computational pipeline to design the controller for hybrid UAVs. Users design the geometry of a hybrid UAV by selecting and matching parts from a component data set. This data set contains parametric components designed by experts and composition rules that constrain how the components can be assembled [Du et al. 2016; Geilinger et al. 2018; Schulz et al. 2017]. We built a numerical simulator to reproduce the real experimental environment for UAVs that takes into account random sensor noise, variances in system identification, and delay in control signals (Section 4).

At the heart of our hybrid UAV control algorithm is a novel and unified neural network controller to support mode-free motion control (Section 5). We aim to tackle two long-standing challenges in the state-of-the-art hybrid UAV control algorithms. First, we want to automate the flying mode transitions for a single aircraft. To this end, we replace all the modes and their discrete transitions, which are usually hard-coded in a hybrid controller, with a mode-free continuous neural network. Second, we aim to automate the transfer from a network trained in a virtual simulator to a controller that can fly a real hybrid UAV. We formulate a new state vector as the input of our network by incorporating an error integral term, enabling a universal solution to eliminate the steady-state drift error that is challenging to fix even manually.

These two essential contributions, in conjunction with our aircraft simulator enhanced with noise perturbations (Section 4) and the reinforcement learning framework with a novel reward function (Section 6), establish a full pipeline to automate the processes of controller design and parameter tuning for a customized hybrid UAV. Finally, we evaluate the efficacy of the neural network controller, state representation, and the fully automated pipeline in numerical settings (Section 8.2). We finish the pipeline by manufacturing the vehicle based on the geometry specifications, set its controller code based on the manufactured vehicle, and conduct real flight experiments (Section 9).

4 DYNAMIC MODEL

The state variables of the hybrid UAV system are represented by a twelve-dimensional vector in phase space:

$$\mathbf{q} = (x, y, z, \phi, \theta, \psi, u, v, w, \omega, \beta, \gamma)^T \in \mathbf{R}^{12}, \quad (1)$$

which can be subdivided into four three-dimensional vectors in sequence. These four state vectors include position $\mathbf{x} = (x, y, z)^T$, the Euler angles (i.e. roll, pitch and yaw) $\boldsymbol{\phi} = (\phi, \theta, \psi)^T$, velocity $\mathbf{u} = (u, v, w)^T$, and angular velocity $\boldsymbol{\omega} = (\omega, \beta, \gamma)^T$.

Following the Newton's second law, the time derivative of the system's linear momentum can be written as

$$m\ddot{\mathbf{x}} = \mathbf{f}^G + \mathbf{f}^T + \mathbf{f}^L + \mathbf{f}^D, \quad (2)$$

with m as mass, \mathbf{f}^G as gravity, \mathbf{f}^T as thrust, \mathbf{f}^L as wing lift and \mathbf{f}^D as wing drag.

Following Euler's rotation equations, the time derivative of the system's angular momentum can be written as

$$\mathbf{J}\dot{\boldsymbol{\omega}} + \boldsymbol{\omega} \times \mathbf{J}\boldsymbol{\omega} = \sum_{i=1}^N (\lambda_i T_i \mathbf{d}_i + \mathbf{r}_i \times T_i \mathbf{d}_i) + \sum_{j=1}^M [s_j \times (\mathbf{f}_j^L + \mathbf{f}_j^D)], \quad (3)$$

with inertial tensor \mathbf{J} , propeller torque-thrust ratio λ , propeller thrust T , propeller direction \mathbf{d} , propeller position in body frame \mathbf{r} , and the mass center of wing in body frame \mathbf{s} . The number N and M denote the number of propellers and wings, respectively.

We follow the flat-plane model [Cory and Tedrake 2008] to approximate the lift and drag on each wing as:

$$\mathbf{f}^L = \rho |\mathbf{u}|^2 \cos(\alpha) \sin(\alpha) \mathbf{S} \quad (4)$$

$$\mathbf{f}^D = \rho |\mathbf{u}|^2 \sin^2(\alpha) \mathbf{S} \quad (5)$$

with ρ as the air density, α as the angle of attack, and \mathbf{S} as the surface area of each wing. The aerodynamic interactions among the propellers, wings, and fuselage are ignored in our dynamic model.

4.1 Numerical Simulation

We build a rigid-body simulator incorporating this simplified aerodynamic model to support the controller training process in a virtual environment. Our simulator mimics a closed feedback loop. In particular, the controller receives its input from both the pilot and the sensor and generates the control signals which are then used to update the rotor thrusts at each time step. The phase variables are updated with an explicit Euler time integration using Equations 2 to 5. We also model noise in a real feedback control loop by applying different numerical perturbations in the system such as adding noise to sensor readings, perturbing mass and inertia, adding noise to lift and drag, and perturbing the latency of control signals and the time step length of the control loop.

5 NEURAL NETWORK CONTROLLER

We design a neural network controller that is mode-free and model-agnostic to automate the control parameter tuning processes for various customized UAV designs. In particular, we incorporate an error integral term in our state vector to enable the transfer from a computationally trained model to a real copter.

5.1 Network Architecture

The architecture of our novel neural network controller is illustrated in Figure 3. We design a neural network with two fully-connected hidden layers. Each layer is composed of 64 \tanh units. The input state $\mathbf{s} \in \mathbf{R}^{13}$ is represented by a vector with fixed dimensions and the output of the network is an N -dimensional vector \mathbf{a} specifying

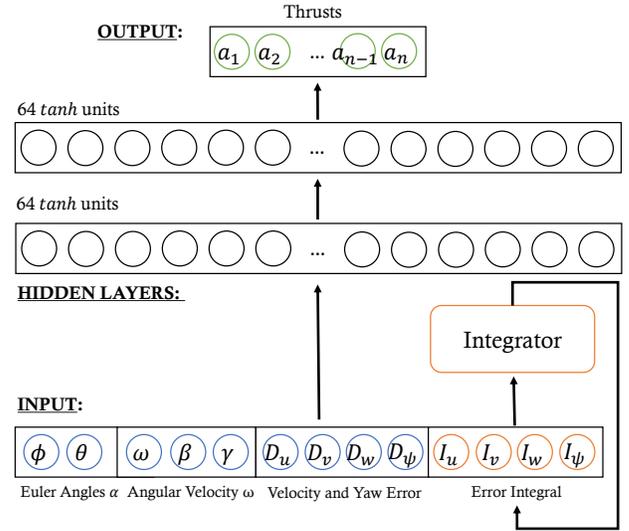


Fig. 3. The neural network controller takes a thirteen-dimensional vector consisting of angle, angular velocity, error, and its integral as input. The neural network consists of two layers with 64 \tanh units each and a linear layer to output the final thrust for each rotor.

the control signal for each rotor, with N as the total number of rotors equipped on the aircraft.

5.2 State Variables

As illustrated in Figure 3, we define the input state vector of a hybrid UAV for our neural network controller as:

$$\mathbf{s} = (\phi, \theta, \omega, \beta, \gamma, D_u, D_v, D_w, D_\psi, I_u, I_v, I_w, I_\psi)^T \in \mathbf{R}^{13}, \quad (6)$$

where D_u, D_v, D_w are the differences between the target and the current velocity components in a local coordinate system, D_ψ is the difference between the target and the current yaw angle, and I_u, I_v, I_w, I_ψ are the convolutional integrals of D_u, D_v, D_w , and D_ψ , respectively. Notice that each component of \mathbf{s} is a function of time t . We omit the symbol t in the following context for conciseness.

Compared to a standard state vector of a hybrid UAV (such as \mathbf{q} in Equation 1), we design our network input vector following these philosophies: first, we remove position from the state because we track velocity only for UAV motion control. We have chosen to use a velocity-tracking controller not only because it works as a unified control representation for various types of hybrid UAVs where an on-board positioning system is not always available, but also because it is a more intuitive framework for the pilot. Second, we record the velocity difference instead of recording both vectors to reduce the dimension of the space. Third, we append a four-dimensional integral term to the end to compensate for accumulated error terms regarding both velocity and yaw angle (see Section 5.3).

If we further define $\bar{\boldsymbol{\phi}} = (\phi, \theta)^T$, $\mathbf{D} = (D_u, D_v, D_w, D_\psi)^T$, and $\mathbf{I} = (I_u, I_v, I_w, I_\psi)^T$, we can rewrite \mathbf{s} in a blockwise manner as

$$\mathbf{s} = (\bar{\boldsymbol{\phi}}^T, \boldsymbol{\omega}^T, \mathbf{D}^T, \mathbf{I}^T)^T, \quad (7)$$

where now the state vector is a combination of the roll and pitch angles, the angular velocity, an error vector consisting of linear velocity and yaw, and the time integral of the error vector.

Finally, we give the definition of \mathbf{D} as the difference between the target state vector and the current state vector in a local frame specified by the current yaw angle:

$$\mathbf{D} = \begin{bmatrix} \mathbf{R}_z(-\psi) & 0 \\ 0 & 1 \end{bmatrix} \left(\begin{bmatrix} \hat{\mathbf{u}} \\ \hat{\psi} \end{bmatrix} - \begin{bmatrix} \mathbf{u} \\ \psi \end{bmatrix} \right). \quad (8)$$

The first three components of the target vector are the target velocity $\hat{\mathbf{u}}$ and the last component is the target yaw angle $\hat{\psi}$. Both the target velocity and the target yaw are in world space. Similarly, the current state vector is composed of the current velocity and the current yaw angle in world space. The difference between these two vectors are then transformed into a local space by a homogeneous matrix specified by a rotation with the angle of $-\psi$ along the world-down axis (we use North-East-Down coordinates for both world and body frames). We want to highlight that this rigid transformation always aligns the motion in the forward and right directions with the components of D_u and D_v , which is the key to enabling an efficient learning process for the network.

5.3 Error Integral

A distinguishing feature of our state vector compared with previous work is the introduction of the temporally accumulated errors of linear velocity and yaw angle. We will later show that this error term is the critical factor that enables effective transfer of control policies from one aircraft configuration to another (see Section 8.2).

We define the error integral term I as the convolution of \mathbf{D} and a smoothing kernel δ :

$$\mathbf{I}(t) = \int_0^t \mathbf{D}(\tau) \delta(t - \tau) d\tau, \quad (9)$$

The discrete form of Equation 9 can be obtained by discretizing the time axis with n time steps as

$$\mathbf{I}(n) = \sum_{i=1}^n \mathbf{D}(i) \eta^{n-i}, \quad (10)$$

in which a constant decay coefficient η is used to incrementally attenuate the impact of the errors in previous time steps, as well as to avoid numerical explosion after a long time interval.

We can further obtain a recursive expression of Equation 10 to update \mathbf{I}_n based on \mathbf{I}_{n-1} :

$$\mathbf{I}_n = \eta \mathbf{I}_{n-1} + \mathbf{D}_n. \quad (11)$$

We use Equation 10 to characterize the drift between different system configurations. In our network training experiments (see Section 8.2), we observed that it is challenging to transfer the control parameters from a virtual trained example to a real fabricated UAV. Adding noise into the simulation model can only help marginally. For example, a controller that works for a simulated UAV in tracking a certain target velocity trajectory will exhibit a constant drift when tracking the same trajectory for a real UAV with slightly different physical configurations. This constant drift issue is named the *steady-state error* in control theory [Aström and Murray 2010].

The key idea of this error integral is motivated by the classical PID controller. Recall that a traditional PID controller consists of three terms: P is proportional to the current tracking error, I is an integration of past error over time, and D is the derivative of error which is an estimate of the future trend of the system. The integral term I is designed to eliminate the system's steady-state error by consistently penalizing the accumulated drifts over time. Inspired by that, we append an error integral term to the state input of our neural network model to mimic the behavior of PID in minimizing the accumulated systematic error. Such drift is recorded in every time step and penalized as an energy term in the objective function (see Section 6) when training the neural network.

6 REINFORCEMENT LEARNING

6.1 Framework

We train our controller using a standard reinforcement learning framework. Our control problem is modeled by a Markov decision process problem $(\mathcal{S}, \mathcal{A}, \mathcal{P}, R, s_0, \gamma, H)$, where \mathcal{S} is the set of states, \mathcal{A} is the set of actions, $\mathcal{P}(s'|s, a)$ is the transition function which is a distribution of next states s' given the current s and action a , $R(s, a)$ is a scalar reward the agent can get after taking action a from state s , s_0 is the start state, $\gamma \in [0, 1)$ is a discount factor, and H is the maximum horizon. A policy $\pi_\theta(a|s)$ models the probability of choosing action a in state s . The goal is to find the parameters θ of a policy $\pi_\theta(a|s)$ that maximizes the expected finite-horizon reward:

$$J(\theta) = \mathbb{E}_{\tau \sim p_\theta(\tau)} \left[\sum_{t=0}^H \gamma^t R_t \right], \quad (12)$$

where $p_\theta(\tau) = p(s_0) \prod_{t=0}^{H-1} \mathcal{P}(s_{t+1}|s_t, a_t) \pi(a_t|s_t)$ is the probability distribution over all trajectories $\tau = (s_0, a_0, \dots, s_H, a_H)$, and $p(s_0)$ is the initial state distribution.

In our work, the state s is a 13 dimensional vector defined in Equation 6, and the action a consists of N values representing the thrust of each rotor. The policy π is our neural network controller. We use Proximal Policy Optimization [Schulman et al. 2017] (PPO), the state-of-the-art policy gradient method, to train the controller.

6.2 Reward

Our reinforcement learning reward R at each step t consists of five terms – the velocity tracking error, energy efficiency, flight stability, error integral, and orientation. The philosophy of this reward function is to enable the UAV to match a prescribed trajectory in the velocity space and meanwhile to optimize other objectives of motion. We give the definition of the reward function for the current time instant as:

$$R(\mathbf{s}, \mathbf{a}) = d - w_v c_v(\mathbf{s}) - w_a c_a(\mathbf{a}) - w_\omega c_\omega(\mathbf{s}) - w_I c_I(\mathbf{s}) - w_\psi c_\psi(\mathbf{s}). \quad (13)$$

In Equation 13, d is a constant alive bonus, c_v is the velocity cost, c_a is the action cost, c_ω is the stability cost, c_I is the integral cost, and c_ψ is the orientation cost, and $w_v, w_a, w_\omega, w_I, w_\psi$ are their corresponding weights. The role of the alive bonus d is to keep the controller alive as long as possible by providing a positive reward at each time step.

The need for the first three terms stems directly from the definition of a velocity-tracking controller for hybrid UAVs. The velocity

cost c_v penalizes the deviation of the current velocity from the target velocity as $c_v(\mathbf{s}) = \|\mathbf{v}_t - \hat{\mathbf{v}}\|_2^2$. The action cost $c_a(\mathbf{a})$ aims to maximize the leverage of the wing lift forces by penalizing the thrust output: $c_a(\mathbf{a}) = \|\mathbf{a}\|_2^2$ (see the usage of the wing lift in Figure 7). The stability cost $c_\omega(\mathbf{s})$ prevents the jittering motion by minimizing angular velocity: $c_\omega(\mathbf{s}) = \|\omega\|_2^2$. The fourth term, the integral cost c_I , attempts to zero out the steady state error by discouraging large integral error: $c_I(\mathbf{s}) = \|I_u^2 + I_v^2 + I_w^2 + I_\psi^2\|_2^2$. As previously discussed, this term is essential for bridging the reality gap.

Finally, the orientation cost $c_\psi = D_\psi^2$ encourages the controller to fly along the desired orientation as straight as possible by penalizing large D_ψ . As our velocity is defined in a re-oriented coordinate system, a UAV that travels along an arbitrary curve with its heading consistent with the tangent direction and its velocity being of the same magnitude as $\hat{\mathbf{v}}$ also has near-zero velocity cost. The term c_ψ takes the yaw angle ψ at each time step into account to penalize those undesired turning motions. As the control of a UAV is affected not by the actual yaw value but by the yaw error, we always set $\dot{\psi}$ to be zero during training. We show the ablation test for the orientation cost in Table 2.

7 HYBRID UAV MODELS

To verify that our method is general for different UAV configurations, we developed a CAD interface within a parametric CAD tool [OnShape 2019] to design a variety of different models. We designed a number of subcomponents such as wings, fuselages, and propeller mounts. We define a handful of parameters for each component that control shape and size variations. Users can create new designs by mixing, matching and manipulating parameters of the components. Once the design is finished, system identification (Table 1) is measured and computed based on the assembly of its specific components. The same numbers are used in training and simulation. We created the following virtual hybrid UAV models (Figure 4) to test the algorithm, which we briefly describe below.

Quad-plane. Our quad-plane is modeled based on the corresponding real plane we built. We use the system identification data for its physical parameters in simulation. The four vertical rotors are primarily used to cancel out gravity, and the front rotor is the major source of forward propulsion. For simplicity, we ignore the influence of the propellers on the aerodynamic models of the wings.

Tail-sitter. The virtual tail-sitter model is equipped with a flat wing and two pairs of rotors. To enter the plane mode, the tail-sitter is tilted forward to allow the wings to generate lift. This transition is challenging because the UAV needs to change its pose aggressively in a timely fashion: as the tail-sitter leans forward, the vertical net thrust quickly decays while the wings may yet not be ready to generate enough lift to compensate. This can easily cause the UAV to lose altitude if the controller is slow to respond.

X-wing. The X-wing model is a variant of the aforementioned tail-sitter, and therefore brings similar challenges when transitioning from hovering to gliding, even in simulation. This UAV consists of two orthogonal pairs of wings and four rotors at the end of each wing. Like a classic quadcopter, the design is symmetric around the axis of the fuselage, giving it good performance in copter mode.

Table 1. System identifications of the hybrid UAVs. Numbers are in SI units.

model	quad-plane	tail-sitter	x-wing
mass	1.779	1.108	1.532
diagonal inertia x	0.194	0.13	0.184
diagonal inertia y	0.064	0.025	0.191
diagonal inertia z	0.253	0.149	0.336
wingspan	1.38	1.43	1.51
max thrust	7	7	7

Double-wing. This example is another variant of a tail-sitter but with two pairs of parallel wings. We compute the lift and drag of each wing separately, assuming no interference in between. Rotors are evenly distributed on both sides and are connected by two rods orthogonal to the wings. Two fuselages are modeled between each pair of wings. Transitioning to the plane mode requires the fuselage to lean forward until the wings start to generate sufficient lift.

Asymmetric tail-sitter. We design this non-traditional hybrid UAV models to test the robustness of our controller algorithm. Essentially, this is the same tail-sitter as the second example with randomly translated (within 0.1m) and rotated (up to 15 degrees) rotors.

Note that the dynamics of these models is significantly different: for example, in order to make use of the wings when advancing, the quad-plane model needs to simply turn on the front rotor and mildly adjust the other four rotors, while the tail-sitter model has to lean forward drastically so that the wings become horizontal.

8 TRAINING AND EVALUATION

We train our controller for each hybrid UAV model in a physics simulator enhanced by a set of perturbation terms. We report the implementation details for the training process in Section 8.1. We evaluate our controller in Section 8.2 by conducting a series of ablation tests to demonstrate the efficacy of each reward term.

8.1 Training

We implemented our algorithm in Python and C++ using the PPO algorithm in OpenAI Baselines [Dhariwal et al. 2017]. At the beginning of each trajectory, a target velocity and the initial state are randomly generated, and the UAV attempts to track this velocity in a simulator at 100Hz. Rewards are collected at each time step. The simulation is terminated after a fixed duration (20s in our experiments) or the UAV cannot maintain a stable flying status, after which the simulation restarts with a new target velocity and initial state. The epoch terminates when the total number of simulated time steps reaches 4096. We ran PPO for 30 million time steps using the same hyperparameters in Table 3 for all models. Training each model took 70 minutes with 32 parallel MPI calls on a workstation of 112 CPU cores and 500G memory, and convergence was usually achieved well before the training ended (Figure 5).

During training, we added sensor noise, signal latency, and perturbed dynamic models to improve the robustness of the controller as suggested in previous work [Peng et al. 2018b; Tan et al. 2018]. We summarized their specifications in Table 4. In terms of perturbing the dynamic model, we further constrained the change of thrust at

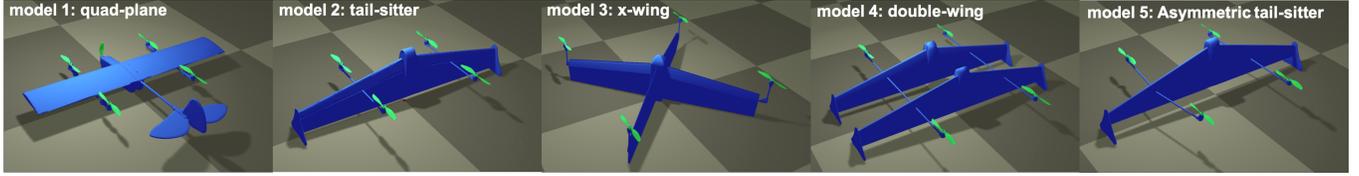


Fig. 4. Five hybrid UAV models used in our training and evaluation process. Note that their dynamics are quite different, making it a non-trivial task to design a control scheme that works for all of them.

Table 2. The average velocity tracking error (smaller is better) and survival time (larger is better) of different controllers in the task of tracking random velocities and two trajectories. Bold number is the best in each row. The meaning of each column is as follows: ours: enabling all features; w/o orientation cost: training without the orientation cost in the reward function; w/o noise or latency: training without sensor noise, signal latency, or perturbation in the dynamic model; w/o integral block: training without the integral block from the neural network controller and the integral term in the reward function; w/o velocity difference: using the target and sensed velocities instead of their difference as the input to the neural network in training.

metric	model	ours	w/o orientation cost	w/o noise or latency	w/o integral block	w/o velocity difference
RandVel velocity tracking error (m/s)	1	0.35	0.45	0.49	0.42	0.51
	2	0.44	0.40	0.39	0.40	0.37
	3	0.45	0.51	0.51	0.52	0.51
	4	0.41	0.57	0.50	0.57	0.43
	5	0.34	1.14	0.49	0.56	0.46
RandVel survival time (s) (maximum: 20s)	1	20.00	19.98	20.00	19.61	19.98
	2	20.00	20.00	11.67	19.62	19.95
	3	20.00	19.90	14.25	20.00	20.00
	4	17.71	19.24	13.13	19.26	20.00
	5	20.00	14.64	11.66	19.66	20.00
VelTraj 1 velocity tracking error (m/s)	1	0.47	0.59	0.64	0.55	0.61
	2	0.57	0.73	0.81	0.61	0.53
	3	0.73	0.69	0.96	0.78	0.66
	4	0.70	0.84	0.93	0.89	0.63
	5	0.62	0.60	0.57	0.58	0.63
VelTraj 1 survival time (s) (maximum: 50s)	1	50.00	50.00	49.07	50.00	50.00
	2	50.00	47.54	10.55	50.00	47.35
	3	50.00	14.76	10.44	50.00	50.00
	4	40.66	14.12	10.90	34.43	50.00
	5	44.83	10.73	10.60	43.81	50.00
VelTraj 2 velocity tracking error (m/s)	1	0.39	0.95	0.56	0.65	0.58
	2	0.43	0.59	0.45	0.41	0.45
	3	0.53	0.66	0.57	0.69	0.55
	4	0.53	0.53	0.49	0.60	0.67
	5	0.44	0.61	0.48	0.61	0.52
VelTraj 2 survival time (s) (maximum: 100s)	1	100.00	39.71	93.33	100.00	100.00
	2	100.00	71.20	69.63	100.00	97.69
	3	100.00	65.02	70.80	73.78	100.00
	4	89.05	47.87	65.30	76.87	97.70
	5	92.57	63.61	70.60	91.25	100.00

each time step to be below $0.4N$. As a result, the system had to learn to avoid generating unreasonably high frequency control signals.

8.2 Evaluation

We conducted ablation study on all our five virtual models (Figure 4) in simulation to validate the necessity of the new components we introduced in the training pipeline. The ablation test shows that our

unique reward definition and the integral block are the core reasons for our success in narrowing the reality gap.

Evaluation Metrics. To evaluate the performance of a given controller in simulation, we consider the following metrics: first, we randomly generated a target velocity and simulated the UAV for up to 20s at 400Hz, the control frequency on real UAV platforms. The

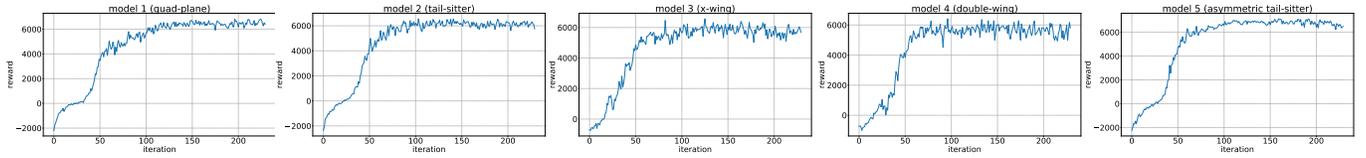


Fig. 5. The training progress on 5 hybrid UAV models using 30 million timesteps in 70 minutes.

Table 3. Hyperparameters used in training all models. Please refer to [Dhariwal et al. 2017], Section 5, and Section 6 for the meaning of each parameter.

name	value	name	value
timesteps_per_actorbatch	4096	schedule	linear
clip_param	0.2	η	0.999
entcoeff	0	d	4
optim_epochs	10	w_v	0.5
optim_stepsize	3e-4	w_a	2e-3
optim_batchsize	64	w_ω	5e-2
gamma	0.995	w_I	6e-2
lam	0.95	w_ψ	1.5
adam_epsilon	1e-5		

Table 4. A summary of noise we added in the training process. All magnitudes use SI units. Percentage numbers mean the noise is relative.

category	noise distribution	parameter range
sensor noise	Gaussian	$ \mu \leq 0.2, 0.005 \leq \sigma \leq 0.05$
mass perturbation	uniform	$\leq 5\%$
inertia perturbation	uniform	$\leq 40\%$
lift and drag	uniform	10% to 20%
signal delay	Gaussian	$\mu = 0.04, \sigma = 0.01$
time step delay	Gaussian	$\mu = 0.01, \sigma = 0.005$

simulation stopped early once the controller was incapable of maintaining a steady flight. We report the velocity tracking error and the survival time averaged in 50 repeated experiments. We refer to this as the RandVel task. Second, we designed two velocity trajectories by concatenating multiple one-dimensional constant speed tracking tasks: advancing, backing, ascending, descending, and side sliding at a constant rate. The UAV was instructed to finish the selected tasks in order. We refer to them as the VelTraj 1 and VelTraj 2 tasks respectively. In VelTraj 1, we designed a 50s velocity trajectory and commanded the UAV first to climb up at a constant rate, followed by advancing then landing. In VelTraj 2, the task is more complex, including advancing and back, side sliding and back, descending and ascending, and finally advancing at a high speed. We repeated the experiment 25 times for each task and report the average. We conducted all tasks in our simulation with noise modelling.

In the remaining section, we report the ablation study on the following components we introduced in the training pipeline: (1) the orientation cost in rewards, (2) the sensor noise, signal latency,

and the dynamics noise in simulation, (3) the integral block in the neural network controller, and (4) the usage of velocity difference as the input to the controller. We remove each component from the pipeline, retrain the controller using the same hyperparameters, and analyze the performance of the controller using the evaluation metrics above. This gives us 5 training environments to compare in total. We report the performance in these tasks in Table 2.

The Orientation Cost. As can be seen from Column 3 and 4 of Table 2, penalizing large orientation error lengthens the survival time in general. Among all the 15 experiments, including the orientation cost in rewards improved the flight time in 12 of them. Specifically, it brought significant performance boost for VelTraj 2, which is the most complex task among the three. Moreover, by comparing Column 4 and 5 with Column 6 and 7, we can see the orientation cost and the noise are the two most indispensable factors for increasing the flight time. Based on these observations, we conclude that by discouraging large heading angles, it is easier for the UAV to maintain a steady flight.

Compared with its influence on the flight time, the improvement on velocity tracking is less notable, part of the reason being that the orientation cost does not consist of any velocity quantities.

Sensor Noise, Latency, and Perturbation on Dynamics. Column 5 of Table 2 summarizes the performance of the controller after disabling all types of noise mentioned above. Together with Column 3, Column 5 shows that disabling noise in training has very strong negative effects on both metrics. In fact, after excluding noise, the performance became the worst in 10 and 5 out of the 15 experiments on flight time and velocity tracking error respectively.

Another interesting observation is that disabling noise is the least effective on model 1 among all models, especially for the flight time. Noting that the other three models are all variants of tail-sitters, this implies the quad-plane dynamic model is more resistant to noise. This is not surprising: a tail-sitter needs to change its pitch aggressively to switch between plane and copter modes, and therefore even a moderate perturbation may easily flip it over.

Integral Block. Quantitative results on the performance of the integral block are reported in Column 6 of Table 2. Comparing Column 6 with Column 3, we observe that adding the integral block suppressed the velocity residual moderately in 12 out of the 15 experiments. This is consistent with the behavior of an integral controller in the classic PID control theory, i.e., it can zero out the steady-state error given enough time. The velocity tracking errors in our experiments did not end up being zero because of the noisy environment as well as the fact that we were chasing a changeable target velocity in VelTraj tasks.



Fig. 6. The fabricated hybrid UAV models. From left to right: quad-plane (model 1), tail-sitter (model 2), x-wing (model 3).

We also notice the improvement of the survival time after introducing the integral block, but it is less significant. As a result, it is a good complement to the aforementioned orientation cost, which increases the flight time by a large margin but is less effective in reducing velocity tracking errors.

The combination of the orientation cost, the integral block, and the noise and latency model is the core reason for successfully bridging the simulation to reality gap when we transfer the trained controller to hardware, as shown in our video. Furthermore, the performance drop between Column 3 and Column 4 to 6 indicates that all three components are indispensable.

Velocity Inputs. In this experiment, we fed the desired and actual velocity as opposed to their difference into the neural network controller and report the results in Column 7 of Table 2. Among all the 30 experiments, the resulting controller outperformed ours in 9 experiments and was dominated by ours in fifteen, with the remaining six being the same. This modification to the velocity inputs is essentially nothing but a simple arithmetic operation. Given that our neural network has two hidden layers of 64 neurons, the network may easily absorb this operation into its control policy, which explains why this modification results in a controller with comparable performance to ours. Nevertheless, we chose to use our original design in all experiments because it led to a more compact input representation.

9 FABRICATION AND FLIGHT TESTS

We fabricate three UAVs using the design models (including the quad-plane, the tail-sitter, and the x-wing, see Figure 6) discussed in the previous section. First, we will discuss the fabrication process in Section 9.1, and then we will present a set of flying tests and log analyses in Section 9.2 to demonstrate the aerodynamic performance of our hybrid UAVs.

9.1 Fabrication

After the shape design is fixed and the controller is finalized in simulation, we generate a corresponding fabrication plan. We fabricated three UAV models to verify our trained controller in real flight. For the hybrid UAVs, fabrication processes include 3D printing, laser cutting and manual assembly. Structural stiffness and weight optimization are paramount to ensure the flight performance of the hybrid UAV. During the materials selection process, we optimize the design to limit weight and maximize structural stiffness. To

ensure structural stiffness, each hybrid UAV has an internal chassis constructed from carbon fiber tubing and fiberglass reinforced 3D printed nylon connectors. To construct the wings, we used two different architectures: one built from laser cut foam, and one built from a balsa wood frame and heat shrink plastic film. Laser cut foam wing are more durable, while the heat shrunk wings allow for more effective lift generation. The designer can choose from one of these techniques depending on preferred performance.

The flight controller runs our modified version of the open source software ArduCopter [ArduPilot 2016] on a Pixhawk flight computer hardware [Meier et al. 2012]. The hybrid UAVs' propellers are driven by brushless electric motors.

We briefly describe the three fabricated UAVs below:

Model 1: Quad-plane. Our custom-built quad-plane (Figure 6 left) is based on a classic fixed-wing plane model plus overlapping the rotor configuration of a standard quadcopter. We first built the fuselage, the wing, and the tail and assembled them into a fixed-wing model. We then added four rotors on both sides of the wing symmetrically. The wings and the tail did not have articulated control surfaces, so they generated lift and drag passively. The attitude of the UAV was primarily controlled by the net effects of all five rotors.

Model 2: Tail-sitter. The tail-sitter model (Figure 6 middle) is a mixture of a delta-wing plane and a quadcopter. We first assembled the fuselage and the pair of wings, after which we placed two rods with rotors on both ends in a position perpendicular to the wings and equally distant to the body. We added two wingtips so that the tail-sitter could take off from a vertical position on land.

Model 3: X-wing. The x-wing plane (Figure 6 right) consists of two pairs of wings perpendicular to each other. All wings were attached to the four sides of the central fuselage. We placed four rotors at the end of each wing and intentionally lifted them a little higher to reduce the interference between propellers and wings.

9.2 Flight Tests

During the flight test, the UAV first loitered for a few seconds to show that the controller can stabilize the body in the air, which is the key feature of a multicopter. We then instructed the UAV to advance directly from the loitering state, exhibiting the aerodynamic behavior of a fixed-wing plane. In contrast to other hybrid UAV controllers, which typically execute a manually designed mode transition from hovering to gliding, our neural network controller

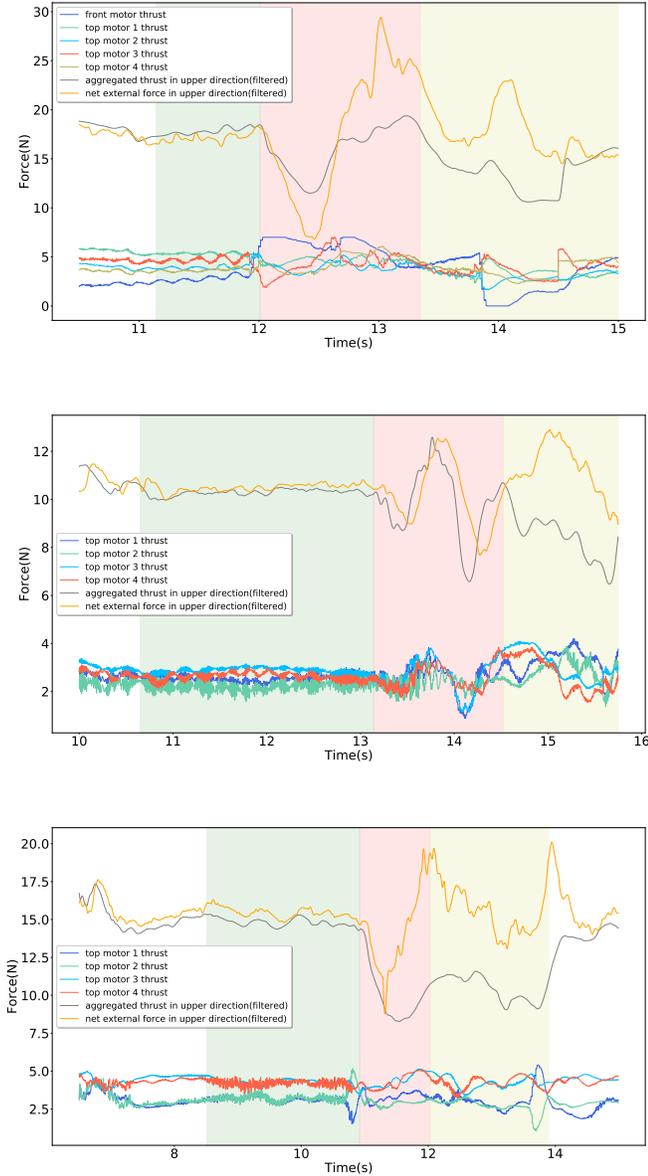


Fig. 7. The flight logs of the quad-plane (top), tail-sitter (middle) and X-wing (bottom). The shaded background represents the time interval when the UAV was loitering (green), transitioning (red), and starting to advance (yellow). The colored curves other than the orange ones correspond to the thrust from each rotor. The gray curve shows the vertical component of the net thrusts, and the orange curve presents the estimated net external force excluding gravity of the UAV. The difference between the orange and gray lines indicates the lift the wings contributed.

does not store any modes and can switch modes implicitly by simply updating target velocities. The UAV then pulled itself back to loiter and land. To demonstrate the generality of our approach, we perform the same flight tasks on three hybrid UAVs which have dramatically different dynamics. We highlight the transition between hovering

and advancing for all three models in Figure 8 and Figure 1. Note that for each model, the same controller was used both in simulation and in the flight video. We ask readers to refer to our video for the full flight tests.

Our flight logs (Figure 7) recorded the output thrust of each rotor over time during the flight, labeled by different colored curves. We computed the vertical component of the net thrust and plotted it as a gray curve in the log. To assist the analysis, we also plotted the net external force excluding gravity in the vertical direction in orange, which we estimated from the motion data. Assuming the UAV maintained its altitude all the time during the flight, we can use the margin between the orange and gray curves to estimate the lift from the wings. Finally, we highlighted the hovering, transition, and advancing using different background colors.

Model 1: Quad-plane. Traditionally, the front rotor of the quad-plane is solely used in the plane flight mode. Our flight log (Figure 7, top) shows that it also contributed to resisting gravity when the UAV was loitering. This novel usage of the front rotor came from the fact that the neural network controller was trained to intentionally fuse the copter and plane flight modes. We also notice that the wings started to provide lift as soon as the UAV entered its advancing mode (as expected), as the net thrust curve consistently lies below the reference weight line.

Model 2: Tail-sitter. During the loitering mode, the tail-sitter behaved quite stably as indicated by the almost constant output thrust from each rotor (Figure 7, middle). The transition from loitering to advancing is significantly harder than the quad-plane because it needs to lean forward by a large amount before the wings become active. Still, our neural network controller was capable of transitioning smoothly from loitering to advancing. From the flight log, we can see that both the rotors and wings were active when the vehicle was cruising afterwards, making the flight more efficient than a pure multicopter.

Model 3: X-wing. The X-wing shares quite a few similarities with the previous tail-sitter model, e.g., steady loitering mode and challenging transition. After entering the cruising mode, the margin between the vertical net force and the vertical net thrust is particularly large because it has the longest wingspan of the three models, making it possible to generate powerful lift.

10 DISCUSSION AND LIMITATIONS

The simulation experiments show that our *model-agnostic* method is able to control widely different configurations of hybrid UAVs and our *mode-free* controller enables flight in all six directions from an input velocity. The ablation tests show that there is a strong motivation for the components we introduce in the pipeline. Most importantly, borrowed from the classic PID control theory, the *integral block* decreases the velocity tracking error. Finally, our flight tests show how our proposed solution can bridge the reality gap.

Our method still leaves room for further developments. First, the analytic flat-plane model does not take into account complex aerodynamic effects between the airflow and the fuselage. We have chosen this approach because it allows real-time simulation in the training process. While such model is qualitatively good enough

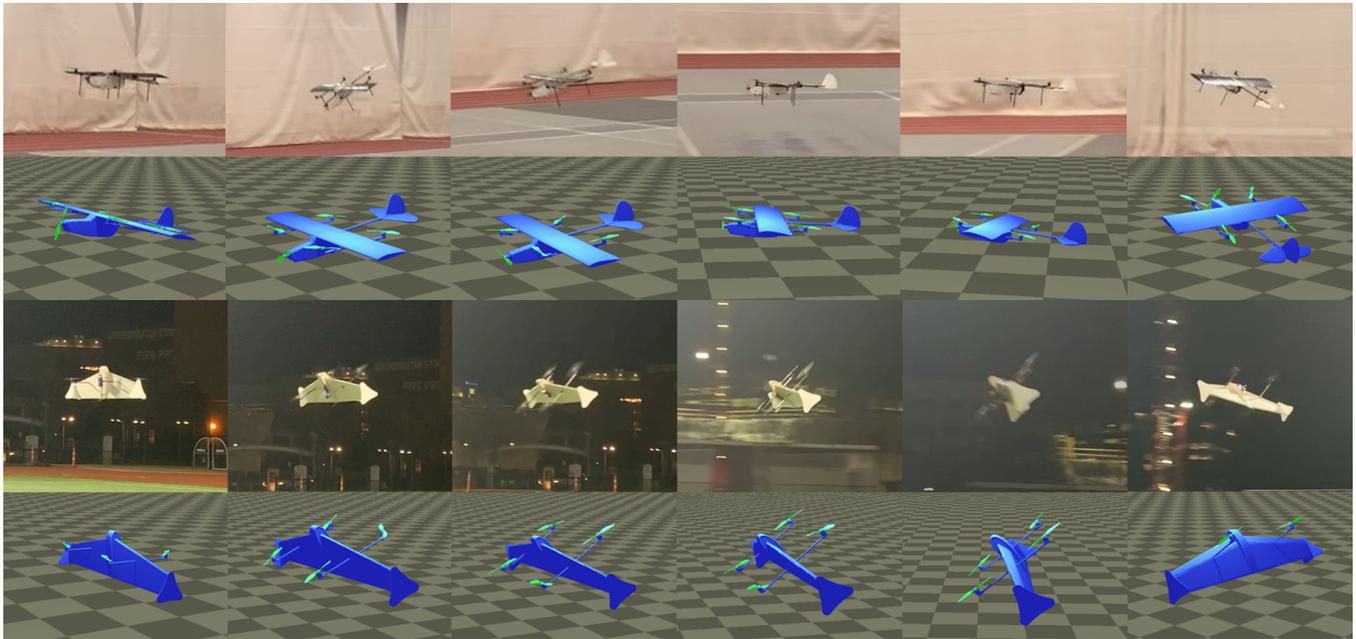


Fig. 8. The quad-plane (model 1) and the tail-sitter (model 2) are transitioning from loitering to gliding (row 1 and 3). Their corresponding flight simulation are shown below (row 2 and 4).

for building real fixed-wing plane models [Barry 2016], it would be interesting to investigate if a more accurate simulation could lead to better results in the case of hybrid UAVs.

Second, our proposed velocity tracking controller has no direct way to measure position, and therefore the deviation between a desired flight path and the actual trajectory of the hybrid UAV could be large. In the future, it would be valuable to add position sensors and design a trajectory tracking controller.

Third, it would be interesting to increase the maneuverability of the vehicles. While the UAVs can efficiently change pitch in order to maximize velocity, changes in the heading direction are slower because our method trains the copter to advance by keeping yaw at zero. Further training would be necessary to allow sharp turns.

Fourth, we restricted the scope of the problem to wings as passive actuators. Future work should investigate models with control surfaces (e.g., rudders) on the wings and tails.

Finally, it would be interesting to explore the parametric component database to develop algorithms for automatic design optimization based on flying capabilities, as in previous work [Du et al. 2016]. In theory, one could couple the proposed automatic method for computing controllers with simulation to measure flight performance for a given input geometry. While this would make performance evaluation quite slow, it could be coupled with pre-computation techniques for interactive design exploration [Schulz et al. 2017]. A more interesting approach would be to develop *concurrent* optimization methods directly in the reinforcement learning pipeline. While some approaches have been suggested [Schaff et al. 2018], further study is still required in this direction.

11 CONCLUSIONS

In this work, we propose a new method for automatically designing controllers for hybrid UAVs. As opposed to standard approaches where separate controllers are developed for each flight mode and their transitions, our method is *mode-free*—given only the input velocity, the controller will automatically orient the model to transition between loitering and advancing. Moreover, we show that our method is *model-agnostic* by demonstrating how the same neural network structure and learning algorithm can be used for vehicles with vastly different configurations. This result can enable further work in transferring knowledge between different hybrid UAV designs or other types of robots.

With a new input representation and a novel reward function, we were able to narrow the reality gap that is common in reinforcement learning approaches. We have shown how our learned controller can be used for flying different manufactured designs. A key technical contribution that enabled this result was the incorporation of a new *integral block*, inspired by PID controllers, which can compensate for system identification errors. We expect that this proposed solution will find application in many other domains, expanding the capabilities of reinforcement learning algorithms.

ACKNOWLEDGMENTS

We thank Ryan Gulland, Mohammed Al Jashmi, and Abdelgadir Balla for providing insightful ideas of hybrid UAVs designs. This work was supported by Air Force Research Laboratory’s sponsorship of Julia: A Fresh Approach to Technical Computing and Data Processing (Sponsor Award ID FA8750-15-2-0272, MIT Award ID 024831-00003).

REFERENCES

- Pieter Abbeel, Adam Coates, and Andrew Y. Ng. 2010. Autonomous Helicopter Aerobatics through Apprenticeship Learning. *The International Journal of Robotics Research* 29, 13 (Nov. 2010), 1608–1639. <https://doi.org/10.1177/0278364910371999>
- Pieter Abbeel, Adam Coates, Morgan Quigley, and Andrew Y. Ng. 2007. An Application of Reinforcement Learning to Aerobatic Helicopter Flight. In *Proceedings of the 19th International Conference on Neural Information Processing Systems (NIPS '07)*. 1–8. <http://dl.acm.org/citation.cfm?id=2976456.2976457>
- Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Jozefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, et al. 2018. Learning Dexterous In-Hand Manipulation. Retrieved April 22, 2019 from <https://arxiv.org/abs/1808.00177>
- ArduPilot. 2016. ArduPilot Open Source Autopilot. Retrieved April 22, 2019 from <http://ardupilot.org/>
- Karl Johan Aström and Richard M. Murray. 2010. *Feedback Systems: an Introduction for Scientists and Engineers*. Princeton University Press.
- Somil Bansal, Anayo K. Akametalu, Frank J. Jiang, Forrest Laine, and Claire J. Tomlin. 2016. Learning Quadrotor Dynamics Using Neural Network for Flight Control. In *Proceedings of 2016 IEEE 55th Conference on Decision and Control (CDC '16)*. 4653–4660. <https://doi.org/10.1109/CDC.2016.7798978>
- Roman Bapst, Robin Ritz, Lorenz Meier, and Marc Pollefeys. 2015. Design and Implementation of an Unmanned Tail-sitter. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '15)*. IEEE, 1885–1890. <https://doi.org/10.1109/IROS.2015.7353624>
- Andrew Barry. 2016. *High-Speed Autonomous Obstacle Avoidance with Pushbroom Stereo*. Ph.D. Dissertation. Massachusetts Institute of Technology, Cambridge, MA.
- Gaurav Bharaj, David I. W. Levin, James Tompkin, Yun Fei, Hanspeter Pfister, Wojciech Matusik, and Changxi Zheng. 2015. Computational Design of Metallophone Contact Sounds. *ACM Trans. Graph.* 34, 6, Article 223 (Oct. 2015), 13 pages. <https://doi.org/10.1145/2816795.2818108>
- Adam Coates, Pieter Abbeel, and Andrew Y. Ng. 2008. Learning for Control from Multiple Demonstrations. In *Proceedings of the 25th International Conference on Machine Learning (ICML '08)*. ACM, New York, NY, USA, 144–151. <https://doi.org/10.1145/1390156.1390175>
- Stelian Coros, Bernhard Thomaszewski, Gioacchino Noris, Shinjiro Sueda, Moira Forberg, Robert W. Sumner, Wojciech Matusik, and Bernd Bickel. 2013. Computational Design of Mechanical Characters. *ACM Trans. Graph.* 32, 4, Article 83 (July 2013), 12 pages. <https://doi.org/10.1145/2461912.2461953>
- Rick Cory and Russ Tedrake. 2008. Experiments in Fixed-Wing UAV Perching. In *AIAA Guidance, Navigation and Control Conference and Exhibit*. 7256.
- Ruta Desai, Ye Yuan, and Stelian Coros. 2017. Computational Abstractions for Interactive Design of Robotic Devices. In *2017 IEEE International Conference on Robotics and Automation (ICRA '17)*. 1196–1203. <https://doi.org/10.1109/ICRA.2017.7989143>
- Coline Devin, Abhishek Gupta, Trevor Darrell, Pieter Abbeel, and Sergey Levine. 2017. Learning Modular Neural Network Policies for Multi-Task and Multi-Robot Transfer. In *2017 IEEE International Conference on Robotics and Automation (ICRA '17)*. 2169–2176. <https://doi.org/10.1109/ICRA.2017.7989250>
- Prafulla Dhariwal, Filip Wolski, Oleg Klimov, Alex Nichol, Matthias Plappert, Alec Radford, John Schulman, Szymon Sidor, Yuhuai Wu, and Peter Zhokhov. 2017. OpenAI Baselines. Retrieved April 22, 2019 from <https://github.com/openai/baselines>
- Tao Du, Adriana Schulz, Bo Zhu, Bernd Bickel, and Wojciech Matusik. 2016. Computational Multicopter Design. *ACM Trans. Graph.* 35, 6, Article 227 (Nov. 2016), 10 pages. <https://doi.org/10.1145/2980179.2982427>
- Justin Fu, Sergey Levine, and Pieter Abbeel. 2016. One-shot Learning of manipulation skills with online dynamics adaptation and neural network priors. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '16)*. IEEE, 4019–4026.
- Moritz Geilinger, Roi Poranne, Ruta Desai, Bernhard Thomaszewski, and Stelian Coros. 2018. Skaterbots: Optimization-based Design and Motion Synthesis for Robotic Creatures with Legs and Wheels. *ACM Trans. Graph.* 37, 4, Article 160 (July 2018), 12 pages. <https://doi.org/10.1145/3197517.3201368>
- Nicolas Heess, Srinivasan Sriram, Jay Lemmon, Josh Merel, Greg Wayne, Yuval Tassa, Tom Erez, Ziyu Wang, Ali Eslami, Martin Riedmiller, et al. 2017. Emergence of Locomotion Behaviours in Rich Environments. *arXiv preprint arXiv:1707.02286* (2017).
- R. Hugh Stone, Peter Anderson, Colin Hutchison, Allen Tsai, Peter Gibbens, and K. C. Wong. 2008. Flight Testing of the T-Wing Tail-Sitter Unmanned Air Vehicle. *Journal of Aircraft - J AIRCRAFT* 45 (Mar. 2008), 673–685. <https://doi.org/10.2514/1.32750>
- Jemin Hwangbo, Inkyu Sa, Roland Siegwart, and Marco Hutter. 2017. Control of a Quadrotor With Reinforcement Learning. *IEEE Robotics and Automation Letters* 2, 4 (Oct. 2017), 2096–2103. <https://doi.org/10.1109/LRA.2017.2720851>
- H Jin Kim, Michael I Jordan, Shankar Sastry, and Andrew Y Ng. 2004. Autonomous helicopter flight via reinforcement learning. In *Advances in neural information processing systems (NIPS '04)*. 799–806.
- Libin Liu and Jessica Hodgins. 2018. Learning Basketball Dribbling Skills Using Trajectory Optimization and Deep Reinforcement Learning. *ACM Trans. Graph.* 37, 4, Article 142 (July 2018), 14 pages. <https://doi.org/10.1145/3197517.3201315>
- Vittorio Megaro, Bernhard Thomaszewski, Maurizio Nitti, Otmar Hilliges, Markus Gross, and Stelian Coros. 2015. Interactive Design of 3D-printable Robotic Creatures. *ACM Trans. Graph.* 34, 6, Article 216 (Oct. 2015), 9 pages. <https://doi.org/10.1145/2816795.2818137>
- Lorenz Meier, Petri Tanskanen, Lionel Heng, Gim Hee Lee, Friedrich Fraundorfer, and Marc Pollefeys. 2012. PIXHAWK: A Micro Aerial Vehicle Design for Autonomous Flight using Onboard Computer Vision. *Autonomous Robots* 33, 1-2 (2012), 21–39. <http://www.onshape.com/>
- Atsushi Oosedo, Satoko Abiko, Atsushi Konno, Takuya Koizumi, Tatuya Furui, and Masaru Uchiyama. 2013. Development of a quad rotor tail-sitter VTOL UAV without control surfaces and experimental verification. In *2013 IEEE International Conference on Robotics and Automation (ICRA '13)*. 317–322. <https://doi.org/10.1109/ICRA.2013.6630594>
- Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. 2018a. DeepMimic: Example-guided Deep Reinforcement Learning of Physics-based Character Skills. *ACM Trans. Graph.* 37, 4, Article 143 (July 2018), 14 pages. <https://doi.org/10.1145/3197517.3201311>
- Xue Bin Peng, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel. 2018b. Sim-to-Real Transfer of Robotic Control with Dynamics Randomization. In *2018 IEEE International Conference on Robotics and Automation (ICRA '18)*. 1–8. <https://doi.org/10.1109/ICRA.2018.8460528>
- Xue Bin Peng, Glen Berseth, Kangkang Yin, and Michiel Van De Panne. 2017. DeepLoco: Dynamic Locomotion Skills Using Hierarchical Deep Reinforcement Learning. *ACM Trans. Graph.* 36, 4, Article 41 (July 2017), 13 pages. <https://doi.org/10.1145/3072959.3073602>
- Robin Ritz and Raffaello D'Andrea. 2017. A Global Controller for Flying Wing Tailsitter Vehicles. In *2017 IEEE International Conference on Robotics and Automation (ICRA '17)*. 2731–2738. <https://doi.org/10.1109/ICRA.2017.7989318>
- Fereshteh Sadeghi and Sergey Levine. 2017. CAD2RL: Real Single-Image Flight Without a Single Real Image. In *Proceedings of Robotics: Science and Systems*. Cambridge, Massachusetts. <https://doi.org/10.15607/RSS.2017.XIII.034>
- Charles Schaff, David Yunis, Ayan Chakrabarti, and Matthew R Walter. 2018. Jointly Learning to Construct and Control Agents using Deep Reinforcement Learning. *arXiv preprint arXiv:1801.01432* (2018).
- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. 2015. Trust Region Policy Optimization. In *International Conference on Machine Learning (ICML '15)*. 1889–1897.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- Adriana Schulz, Jie Xu, Bo Zhu, Changxi Zheng, Eitan Grinspun, and Wojciech Matusik. 2017. Interactive Design Space Exploration and Optimization for CAD Models. *ACM Trans. Graph.* 36, 4, Article 157 (July 2017), 14 pages. <https://doi.org/10.1145/3072959.3073688>
- Jie Tan, Tingnan Zhang, Erwin Coumans, Atil Iscen, Yunfei Bai, Danijar Hafner, Steven Bohez, and Vincent Vanhoucke. 2018. Sim-to-Real: Learning Agile Locomotion For Quadraped Robots. In *Proceedings of Robotics: Science and Systems*. Pittsburgh, Pennsylvania. <https://doi.org/10.15607/RSS.2018.XIV.010>
- Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. 2017. Domain Randomization for Transferring Deep Neural Networks from Simulation to the Real World. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '17)*. IEEE, 23–30.
- Nobuyuki Umetani, Takeo Igarashi, and Niloy J Mitra. 2012. Guided Exploration of Physically Valid Shapes for Furniture Design. *ACM Trans. Graph.* 31, 4 (2012), 86–1.
- Nobuyuki Umetani, Yuki Koyama, Ryan Schmidt, and Takeo Igarashi. 2014. Pteromys: Interactive Design and Optimization of Free-formed Free-flight Model Airplanes. *ACM Trans. Graph.* 33, 4, Article 65 (July 2014), 10 pages. <https://doi.org/10.1145/2601097.2601129>
- Jungdam Won, Jongho Park, Kwanyu Kim, and Jehee Lee. 2017. How to Train Your Dragon: Example-guided Control of Flapping Flight. *ACM Trans. Graph.* 36, 6, Article 198 (Nov. 2017), 13 pages. <https://doi.org/10.1145/3130800.3130833>
- Jungdam Won, Jungnam Park, and Jehee Lee. 2018. Aerobatics Control of Flying Creatures via Self-regulated Learning. *ACM Trans. Graph.* 37, 6, Article 181 (Dec. 2018), 10 pages. <https://doi.org/10.1145/3272127.3275023>