



HAL
open science

Reinforcement learning with formal performance metrics for quadcopter attitude control under non-nominal contexts

Nicola Bernini, Mikhail Bessa, Rémi Delmas, Arthur Gold, Eric Goubault,
Romain Penneec, Sylvie Putot, François Sillion

► To cite this version:

Nicola Bernini, Mikhail Bessa, Rémi Delmas, Arthur Gold, Eric Goubault, et al.. Reinforcement learning with formal performance metrics for quadcopter attitude control under non-nominal contexts. *Engineering Applications of Artificial Intelligence*, 2023, 127, 10.1016/j.engappai.2023.107090 . hal-04336289

HAL Id: hal-04336289

<https://polytechnique.hal.science/hal-04336289>

Submitted on 13 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Reinforcement Learning with Formal Performance Metrics for Quadcopter Attitude Control under Non-nominal Contexts

Nicola Bernini^a, Mikhail Bessa^a, Rémi Delmas^a, Arthur Gold^a, Eric Goubault^{b,*}, Romain Pennek^a, Sylvie Putot^b, François Sillion^a

^aUber ATCP, Paris, France

^bLIX, Ecole polytechnique, CNRS, IP-Paris, Palaiseau, France

Abstract

We explore the reinforcement learning approach to designing controllers by extensively discussing the case of a quadcopter attitude controller. We provide all details allowing to reproduce our approach, starting with a model of the dynamics of a *crazyflie 2.0* under various nominal and non-nominal conditions, including partial motor failures and wind gusts. We develop a robust form of a signal temporal logic to quantitatively evaluate the vehicle's behavior and measure the performance of controllers. The paper thoroughly describes the choices in training algorithms, neural net architecture, hyperparameters, observation space in view of the different performance metrics we have introduced. We discuss the robustness of the obtained controllers, both to partial loss of power for one rotor and to wind gusts and finish by drawing conclusions on practical controller design by reinforcement learning.

Keywords: Reinforcement learning, control, quadcopter dynamics, performance metrics, temporal logics

1. Introduction

Neural net based control is now widely used in control. For instance, reinforcement learning is known to be linked to optimal control [1]. Very impressive real-life experiments have shown how practical reinforcement learning and privileged learning can be [2], but have somehow masked the enormous amount of experiments and heuristics that had to be learned in the process. Indeed, we are still in need for a full understanding of what advantages and performances we can gain from learning-based control, and what level of formal guarantees we can reach, either at design or at verification time.

This paper extends our HSCC 2021 article [3] with a more complete description of several aspects including the modeling, lessons that have been learned, and most importantly the description of the logic that has been used for evaluating performances of our neural net controllers, as well as new results concerning some spurious correlations that appeared in all attitude controllers that we trained.

We concentrate here on low-level controls, and more specifically attitude control for quadcopters. These con-

trollers have the advantage of being understandable - performances being easily measurable -, well studied in the literature, and essential to all higher-level controls and path tracking algorithms. We focus on reinforcement learning (RL) methods, which are close to control and more particularly optimal control. Furthermore, RL has experienced tremendous progress over the past few years, with modern continuous state and action spaces training algorithms such as Soft Actor Critic (SAC) [4] and Twin-Delayed Deep Deterministic Policy Gradients (TD3) [5].

A common belief is that learning-based control would be more robust to perturbations than e.g. PIDs, or at least could be trained to be more robust. Indeed, even a rather small neural net can encode a much more complex feedback control function than a simple PID, but this is commonly believed to be at the expense of formal guarantees. Also, the current zoology of training methods and architecture choices makes it difficult to fully understand the range of possible results.

This paper studies some of these aspects on the fundamental case of an attitude controller for the *crazyflie 2.0* [6] quadcopter. We first present in Section 3 a non-linear ODE model for simulating the dynamics of a quadcopter, and extend it to account for partial motor failures, aerodynamic effects and wind gusts. We then present a flexible training platform with various neural net architectures and algorithms in Section 4, discuss performance evaluation using a robust signal temporal logic in Section 6, and describe our experimental setup in Section 7. Finally we discuss experimental results in Section 8.

This paper develops in detail the following research

*Corresponding author

Email addresses: nicola.bernini@gmail.com (Nicola Bernini), mikhail.bsa@gmail.com (Mikhail Bessa), remi.delmas.3000@gmail.com (Rémi Delmas), arthur.gold.ag@gmail.com (Arthur Gold), goubault@lix.polytechnique.fr (Eric Goubault), romain.pennek@gmail.com (Romain Pennek), putot@lix.polytechnique.fr (Sylvie Putot), francois.sillion@gmail.com (François Sillion)

items:

1. we develop a neural-net based control study case, after modeling a quadcopter’s dynamics, including aerodynamic effects and partial power loss on motors
2. we discuss the effect of the chosen training algorithm, neural net architecture, reduced observable state spaces and hyperparameters on the performance of the controller, and on the RL training process
3. we present our experimental platform, which allowed us to compare more than 16,000 parameter choices
4. we develop Signal Temporal Logic observers to assess controller performance in a precise manner
5. we demonstrate high-quality attitude control using RL, for a relevant set of queries
6. we show that these controllers have a certain built-in robustness in non-nominal cases, with respect to partial failures of actuators and perturbations such as wind gusts.
7. we discuss in details the lessons learned in reinforcement learning, while applying it to the problem of synthesizing quadcopter attitude controllers

2. Related work

This paper is based on, and compared with, the following work:

RL in control. Reinforcement learning in control has been advertised, since [7], for the possibility to be more adaptive than classical methods in control such as PIDs. RL’s close relationship with optimal control (the reward function is dual to the objective function) also makes it particularly appealing for applications to control, see e.g. [1].

Recently model-based reinforcement learning has been successfully used to train controllers without any initial knowledge of the dynamics and in a data-efficient way. For instance, in [8], a learning-based model predictive control algorithm has been used to synthesize a low level controller. In [9], a hybrid approach is proposed, combining the model based algorithm PILCO [10] and a classic controller like a PD or a LQR controller.

In this paper, we focus on model-free algorithms because of their generality and because we have high fidelity models available for quadrotors, such as the crazyflie 2.0 [6]. More specifically we concentrate on actor-critic learning which has undergone massive improvements over the last few years with DDPG [11], SAC [4], TD3 [5], and compare it with the popular PPO method [12].

The high dimensionality of the full Markovian observation space is a challenge for training, prompting for a study of different choices for the sets of states observed by RL: we consider sub-spaces of the full Markovian observation space, where we leave out the states which have the least effect on the dynamics of the quadcopter. This is linked to partially observed Markov Decision Processes and Non Markovian learning, see e.g. [13].

We also study the robustness of our neural nets, as well as the specific training of the neural net controller to be able to handle disturbances (wind gusts, partial motor failures). These issues may be linked to robust MDPs [14], but we have stuck to the classical (PO)MDP approach here, for which we have a wealth of tools and techniques available.

RL for quadcopters, and attitude control. Most papers have been focusing on higher-level control loops, with the notable exception of [15], which serves as the basis of our work. We improve the results of [15] by considering more recent training algorithms (SAC and TD3), finer performance measures, and refined physical models (in particular perturbations due to partial motor failures and wind gusts). The closest other works related to attitude control for quadcopter are [16], [17], [18] and [19].

In [16], the goal is to stabilize a quadcopter in hover mode, from various initial conditions (including initial angular rates). The authors also consider perturbations to the dynamics, which are more predictable than ours: motor lag and noise on sensors.

In [17], the objective is to control a quadcopter under cyber-attacks targeting its localization sensors (gyroscope and GPS) and motors. The authors consider (partial) motor failure (a limit on its maximal power, just like we do), but not wind gusts. Contrarily to most approaches including ours, their controller combines a classical controller and a neural net.

In [18] the authors discuss the training of a neural net controller for both attitude and position. They observe that it is difficult to train both aspects at the same time, whereas separating control in hover mode (acting mostly on the attitude) and control in position seems to work better. The learning process is based on a full state observation plus the difference with the target state. We extend this work first in discussing the simplification of the observed states, then in more rigorously defining observation metrics for offsets and overshoots.

In [19], the author considers neural nets for controlling roll, pitch, yaw rate and thrust, which is similar to the problem we are studying here, and attempts to train the controller such that it can accommodate motor and mass uncertainties within given bounds. In contrast, we deal with uncertainties such as wind gusts and motor failures, following known parametric models.

Signal Temporal Logics. The study of reinforcement learning under temporal logic specifications has gained a lot of interest in recent years. In a discrete and finite state setting, in [20] a linear-time temporal logics (LTL) property observer automaton is composed with the system MDP to allow blocking unsafe actions during training. In [21, 22] rewards are modulated depending on the observer state, and a model-free approach is proposed in [23] using Limit Deterministic Büchi Automata. *Shielding* [24] simultaneously trains an optimal controller and a *shield* that corrects

the LTL-formula violating actions. The method requires a fully explicit model of the environment MDP and builds the product of the original MDP with the property monitor. Later works extend shielding to the continuous [25, 26] and online [27] cases, assuming an embeddable predictive environment model is available, but only handle simple state invariants.

Temporal logics with quantitative semantics such as Metric Interval Temporal Logic (MITL) [28], Signal Temporal Logic (STL) [29] . . . , have been studied in relation with reinforcement learning. Robust interpretation yields a real number indicative of the distance to the falsification boundary. STL has seen numerous extensions improving expressiveness and signal classes [30, 31, 32, 33] as well as smooth differentiable semantics [34, 35, 36]. Solutions to well known dimension and magnitude mismatch in robust STL interpretation were proposed recently in [37] but have not yet been used in a RL setting. STL usages are varied: In [38], Q-learning is used to train a policy maximizing both the probability of satisfaction and the expected robustness of a given STL specification; The approach requires storing previously visited states in the MDP in addition to the original MDP state, yielding a high dimensional system and limiting learning efficiency. In [39] the authors derive barrier functions from robust temporal logic specifications, either to modulate rewards during training or to control the switch from an optimal and potentially unsafe controller to a safe backup controller [40].

In summary, existing methods focused on the training phase either suffer from dimensionality and combinatorial explosion, require expected robustness approximations, or are strongly tied to the Q-learning algorithms.

Considering our goal is to study a large hyper-parameter space for training controllers and we need to quantify controller performance rigorously, we used an expressive yet tractable variant of STL [32] to specify properties and assess trained controllers offline, separately after training. The next steps will be to start using STL-derived reward signals during training on the most promising architectures.

3. Modelling and control of a crazyflie 2 quadrotor

In this section, we present the dynamical model of the crazyflie quadrotor [41, 42] and we augment it with partial motor failures and wind gusts modelling.

3.1. Nominal model

We study the dynamics on the vertical axis and the pitch rate, roll rate and yaw rate control (4 degrees of freedom), with the following state variables: the vertical position in the world frame z , the linear velocity of the center of gravity in the body-fixed frame with respect to the inertial frame (u, v, w) , the angular orientation represented by the Euler angles (ϕ, θ, ψ) where ϕ is the roll angle θ is the pitch angle and ψ is the yaw angle, the attitude

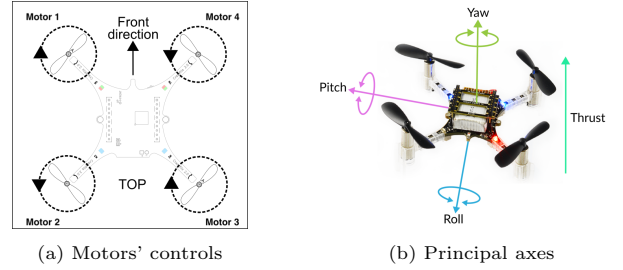


Figure 1: Crazyflie 2.0 – source: <http://www.bitcraze.io> [43] CC BY-SA 3.0

or angular velocity with respect to the body frame (p, q, r) with p the roll rate, q the pitch rate and r the yaw rate.

The Crazyflie 2.0 linear velocities are controlled through the angular velocities and the angular velocities are controlled through rotor thrust differential. For instance, to increase the pitch rate q , *Motor*₂ and *Motor*₃ rotor speeds should be higher than *Motor*₁ and *Motor*₄ (see Figure 1a). As there is symmetry, it works similarly for the roll rate p (with *Motor*₄ and *Motor*₃ vs. *Motor*₁ and *Motor*₂ instead). However, the yaw rate r is controlled through the gyroscopic effect. To make the quadcopter rotate clockwise in the x-y plane, the rotor speeds of the clockwise rotating motors (*Motor*₂ and *Motor*₄) should be higher than those of the counterclockwise rotating ones (*Motor*₁ and *Motor*₃).

Using Newton's equations given a *thrust* force and moments M_x , M_y and M_z exerted along the three axes of the quadcopter, and using the rotation matrix R from the body frame to the inertial frame,

$$R = \begin{pmatrix} c_\psi c_\theta & c_\psi s_\theta s_\phi - c_\phi s_\psi & s_\psi s_\phi + c_\psi c_\phi s_\theta \\ c_\theta s_\psi & c_\psi c_\phi + s_\psi s_\theta s_\phi & c_\phi s_\psi s_\theta - c_\psi s_\phi \\ -s_\theta & c_\theta s_\phi & c_\theta c_\phi \end{pmatrix}$$

(and R^{-1} is the transpose of R) the Translation-Rotation kinematics and dynamics [42] lead to a 10-dimensional non-linear dynamical system:

$$\begin{cases} \dot{z} = -s_\theta u + c_\theta s_\phi v + c_\theta c_\phi w & \dot{\theta} = c_\phi q - s_\phi r \\ \dot{u} = rv - qw + s_\theta g & \dot{\psi} = \frac{c_\phi}{c_\theta} r + \frac{s_\phi}{c_\theta} q \\ \dot{v} = -ru + pw - c_\theta s_\phi g & \dot{p} = \frac{I_y - I_z}{I_x} qr + \frac{1}{I_x} M_x \\ \dot{w} = qu - pv - c_\theta c_\phi g + \frac{F}{m} & \dot{q} = \frac{I_z - I_x}{I_y} pr + \frac{1}{I_y} M_y \\ \dot{\phi} = p + c_\phi t_\theta r + t_\theta s_\phi q & \dot{r} = \frac{I_x - I_y}{I_z} pq + \frac{1}{I_z} M_z \end{cases} \quad (1)$$

writing c_x as a short for $\cos(x)$, s_x for $\sin(x)$ and t_x for $\tan(x)$. F is the sum of the individual motor thrusts, and I_x , I_y , I_z are the quadcopter's moments of inertial around the x , y and z axes, respectively.

Instead of controlling directly each rotor speed, the four commands *thrust*, cmd_ϕ , cmd_ψ and cmd_θ are used to deduce the PWM (Pulse Width Modulation) values to apply to each motor, Equation (2):

$$PWM = \begin{bmatrix} PWM_1 \\ PWM_2 \\ PWM_3 \\ PWM_4 \end{bmatrix} = \begin{bmatrix} 1 & -1/2 & -1/2 & -1 \\ 1 & -1/2 & 1/2 & 1 \\ 1 & 1/2 & 1/2 & -1 \\ 1 & 1/2 & -1/2 & 1 \end{bmatrix} \begin{bmatrix} thrust \\ cmd_\phi \\ cmd_\theta \\ cmd_\psi \end{bmatrix} \quad (2)$$

PWMs are linked to rotation rates Ω : $\Omega = [\omega_1 \ \omega_2 \ \omega_3 \ \omega_4]^\top = C_1 PWM + C_2$. Finally, we deduce the input force and moments from the squared rotation rates, Equation (1), with force and momentum equations $[F \ M_x \ M_y \ M_z]^\top$ equal to:

$$\begin{bmatrix} C_T(C_1^2(cmd_\theta^2 + cmd_\phi^2 + 4cmd_\psi^2 + 4thrust^2) + 8C_1C_2thrust + 4C_2^2) \\ 4C_Td(C_1^2(cmd_\phi thrust - cmd_\theta cmd_\psi) + C_1C_2cmd_\phi) \\ 4C_Td(C_1^2(cmd_\theta thrust - cmd_\phi cmd_\psi) + C_1C_2cmd_\theta) \\ 2C_D(C_1^2(4cmd_\psi thrust - cmd_\phi cmd_\theta) + 4C_1C_2cmd_\psi) \end{bmatrix} \quad (3)$$

The physical and constant parameters we are using for the crazyflie are obtained by merging data from [42] and [6] and listed in Table 1:

3.2. Motor failure

We suppose that the quadcopter may experience a power loss on motor 1. This partial failure is modeled as a saturation of the maximum PWM, with a factor between 0.8 and 1.

Since quadcopter controls rely on differential thrust between motors, motor failures are very difficult to cope with. In order to keep a constant yaw when one motor is failing, the gyroscopic effect must be made equal to zero, for instance by having the two motors rotating in the opposite direction match the saturation of the faulty motor. The same idea applies to pitch and roll axes.

Therefore, if the failure is not too harsh, and the target states are not too demanding, it is *a priori* feasible to recover some control of the faulty quadrotor by saturating all four motors in the same way.

In this paper, we will look at two potential solutions to control in the presence of partial motor failure. The first one is to look at how robust a controller that has been designed for nominal cases (i.e. without partial motor failures) is. The other one is to train, using reinforcement learning, a controller optimized for a variety of non-nominal situations.

3.3. Wind gusts

3.3.1. Aerodynamic effects

In Equation (1), we neglected all aerodynamic effects. When we take into account aerodynamic forces, an extra force F^a is exerted on the quadcopter that depends on the wind speed and direction relative to the quadcopter, the angular velocities of the rotors and extra moments M_x^a , M_y^a and M_z^a . We follow the full aerodynamic model of [6]

with the coefficients measured for a crazyflie 2.0, where the effect of the wind on the structure is neglected with respect to the effect on the rotors, and the blade flipping effect (due to elasticity of the rotor) is also neglected.

The extra force F^a can be decomposed as the sum of the four extra aerodynamic forces on rotor i ($i = 1, \dots, 4$), that can be modelled as depending linearly on the rotors angular velocities, and linearly on the wind relative speed with respect to rotors. Other models [44] include blade flipping and other drag effects, but the induced drag we are modelling is the most important one for small quadrotors with rigid blades. We use $f^i = \Omega_i K W_i^r$ for the aerodynamic force exerted on rotor i in the inertial frame, where K is the drag coefficients matrix, W_i^r is the relative wind speed as seen from rotor i , in the body frame, i.e. $W_i^r = (u_i, v_i, w_i) - R^T W_a$ with W_a the absolute wind speed in the inertial frame, (u_i, v_i, w_i) being the linear velocities of the rotors in the body frame, R is the rotation matrix from the body frame to the inertial frame (R^T is its inverse, i.e. its transpose here), and Ω_i is the absolute value of the angular velocity of the i -th rotor.

The drag coefficients we are using for the crazyflie are one of the models of [6]:

$$K = \begin{pmatrix} -9.1785 & 0 & 0 \\ 0 & -9.1785 & 0 \\ 0 & 0 & -10.311 \end{pmatrix} 10^{-7} kg \cdot rad^{-1}$$

For the crazyflie, $\Omega_i = C_1 PWM_i + C_2$, where the expression PWM_i depends on $thrust$, cmd_ϕ , cmd_θ and cmd_ψ as given by Equation (2).

The linear velocities of rotors can be computed as follows:

$$\begin{pmatrix} u_j \\ v_j \\ w_j \end{pmatrix} = \begin{pmatrix} p \\ q \\ r \end{pmatrix} \times \begin{pmatrix} dc_j \\ ds_j \\ h \end{pmatrix} + \begin{pmatrix} u \\ v \\ w \end{pmatrix} = \begin{pmatrix} qh - rds_j + u \\ -ph + rdc_j + v \\ pds_j - qdc_j + w \end{pmatrix}$$

(u, v, w) are the linear velocities of the center of mass of the quadrotor in the body frame, (p, q, r) are the angular velocities of the quadrotor (see Section 3.1). d is the length of the arm linking the center of the drone to any of the four motors, and for $j \in \{1, 2, 3, 4\}$, $c_j = \sin(\frac{\pi}{2}(j-1) + \frac{3\pi}{4})$ and $s_j = \cos(\frac{\pi}{2}(j-1) + \frac{3\pi}{4})$ are such that (c_j, s_j, h) is the coordinate of rotor j in the body frame, with the center of mass being the origin.

Now, we add to the second term of Equation (1) for \dot{u} , \dot{v} , \dot{w} the aerodynamic force $F^a = (F_x^a, F_y^a, F_z^a)$ divided by m , and to moments of Equation (3), the aerodynamic moments $M^a = (M_x^a, M_y^a, M_z^a)$ with $F^a = f_1 + f_2 + f_3 + f_4$ and $M^a = (dc_1, ds_1, h) \wedge f_1 + (dc_2, ds_2, h) \wedge f_2 + (dc_3, ds_3, h) \wedge f_3 + (dc_4, ds_4, h) \wedge f_4$.

We derive the full dynamics of the quadcopter considering aerodynamic effects, and only write below the modified

Param	Description	Value	Unit
I_x	Inertia about x-axis	$1.657\,171 \times 10^{-5}$	$\text{kg} \cdot \text{m}^2$
I_y	Inertia about y-axis	$1.665\,560\,2 \times 10^{-5}$	$\text{kg} \cdot \text{m}^2$
I_z	Inertia about z-axis	$2.926\,165\,2 \times 10^{-5}$	$\text{kg} \cdot \text{m}^2$
m	Mass	0.028	kg
g	Gravity	9.81	$\text{m} \cdot \text{s}^{-2}$
C_T	Thrust Coefficient	1.285×10^{-8}	$\text{N} \cdot \text{rad}^{-2} \cdot \text{s}^2$
C_D	Torque Coefficient	7.645×10^{-11}	$\text{N} \cdot \text{rad}^{-2} \cdot \text{s}^2$
C_1	PWM to Ω factor	0.040 765 21	-
C_2	PWM to Ω bias	380.8359	-
h	z rotor wrt CoG	0.005	m
d	Arm length	$0.046/\sqrt{2}$	m
p_{max}	Maximum motor PWM	65 535	-

Table 1: Parameters for the crazyflie 2.0 model

equations:

$$\begin{cases} \dot{u} = rv - qw + s_\theta g + \frac{F_x^a}{m} \\ \dot{v} = -ru + pw - c_\theta s_\phi g + \frac{F_y^a}{m} \\ \dot{w} = qu - pv - c_\theta c_\phi g + \frac{F_z + F_z^a}{m} \\ \dot{p} = \frac{I_y - I_z}{I_x} qr + \frac{1}{I_x} (M_x + M_x^a) \\ \dot{q} = \frac{I_z - I_x}{I_y} pr + \frac{1}{I_y} (M_y + M_y^a) \\ \dot{r} = \frac{I_x - I_y}{I_z} pq + \frac{1}{I_z} (M_z + M_z^a) \end{cases} \quad (4)$$

3.3.2. Wind models

There are two main types of models in the literature, represented by e.g. Discrete Wind Gust and von Kármán Gust or Dryden Gust models. Von Kármán gusts and Dryden gusts are stochastic gust models (homogeneous and stationary gaussian processes) characterized by their power spectral densities for the wind's three components, Dryden gusts being an approximation of Von Kármán gusts.

The Discrete Wind Gusts model consists in a explicit and deterministic representation of wind gusts as half period cosine perturbations ([45], eq. (45)). We focus on this model because it is widely used for aircraft certification (using dozens of discrete wind gusts with different magnitudes and scales).

A discrete wind gust is characterized by its fixed direction, magnitude and scale, and lasts for a half period during which wind speed increases until it reaches its maximum intensity. The absolute wind velocity is given as a function of time as, using the same notations as in Section 3.3.1: $W_a(t) = \frac{A_g}{2} \left(1 - \cos\left(\frac{\pi(t-t_0)}{\delta}\right)\right) V_d$ if $t_0 \leq t \leq t_0 + 2\delta$, 0 otherwise, where A_g is the maximal magnitude of the wind gust, δ is the half life of the gust, and V_d is a normalized vector in R^3 , which is the wind (absolute) direction.

3.4. PID Control

As in [15], the objective is to train only the attitude controller, and not the altitude one. We therefore use a

PID for controlling z . We will also need some idea of what a standard PID controller may achieve in terms of performance, and robustness to wind gusts and failures. For this, we will primarily use one of the altitude and attitude PID controller implemented in the crazyflie 2.0. Given setpoints z_{sp} , p_{sp} , q_{sp} and r_{sp} , the quadrotor is controlled using a PID controller (called PID1 in the sequel) which is the one of [41]:

$$\begin{cases} thrust = 1000(25(2(z_{sp} - z) - w) \\ \quad + 15 \int (2(z_{sp} - z) - w) dt) + 36000 \\ cmd_\phi = 250(p_{sp} - p) + 500 \int (p_{sp} - p) dt \\ cmd_\theta = 250(q_{sp} - q) + 500 \int (q_{sp} - q) dt \\ cmd_\psi = 120(r_{sp} - r) + 16.7 \int (r_{sp} - r) dt \end{cases} \quad (5)$$

But as we will see, the attitude controller implemented in the crazyflie 2.0 is not very reactive, most probably for ensuring that the altitude is very securely controllable (since too much reactivity in pitch and roll means sudden loss of vertical speed). In order to give an idea of what we could observe as best performance, we also designed a specific PID for attitude, that we call PID2, which is much more reactive:

$$\begin{cases} thrust = 3000(z_{sp} - z) \\ \quad + 300 \int (z_{sp} - z) dt - 500\dot{z} + 48500 \\ cmd_\phi = 1000(p_{sp} - p) + 400 \int (p_{sp} - p) dt - 40\dot{p} \\ cmd_\theta = 1000(q_{sp} - q) + 400 \int (q_{sp} - q) dt - 40\dot{q} \\ cmd_\psi = 2000(r_{sp} - r) + 1000 \int (r_{sp} - r) dt - 100\dot{r} \end{cases} \quad (6)$$

4. Training

4.1. Underlying Markov decision process

Reinforcement learning is designed to solve Markov decision problems. At each discrete time step $k = 1, 2, \dots$, the controller observes the state x_k of the Markov process, selects action a_k , receives a reward r_k , and observes the

next state x_{k+1} . As we are dealing with Markov processes, the probability distributions for r_k and x_{k+1} depend only on x_k and a_k . Reinforcement learning tries to find a control policy, i.e. a mapping from states to actions, in the form of a neural net, that maximizes at each time step the expected discounted sum of future reward.

For the attitude control problem at hand, the set of Markovian states is *thrust*, p , q , r , $err_p = p_{sp} - p$, $err_q = q_{sp} - q$, $err_r = r_{sp} - r$ (where (p_{sp}, q_{sp}, r_{sp}) is the target state, or "plateau" we want to reach), in the nominal case (similarly to what is done in e.g. [18]). We will also consider partially observed Markov processes, with only subsets of states for improving sampling over smaller dimensional states, by leaving out those states which should have less influence on the dynamics: our first candidate is to leave out thrust, which appears only as second order terms in the moments calculation, Equation (3), and also, p , q , r that are second order in the formulation of the angular rates, again in Equation (3). We do not consider here adding past information, classical in non Markovian environments [13], that has been used for attitude control in e.g. [15], but increases the dimension by a large amount.

In the case of partial motor failure, we add the knowledge of the maximum thrust for faulty motor 1, as a continuous variable between 80% and 100%. In the case of aerodynamic effect and wind gusts, we add the knowledge of the maximal magnitude and direction (in the inertial frame) of the incoming wind. In both cases, it can effectively be argued that it is possible to detect failures in almost real time, and to measure (or be given from ground stations) maximum winds and corresponding directions, in almost real time as well. In the case of wind-gusts, Markovian states include also the linear velocities u , v and w , since wind gusts are only defined in the inertial frame, and the induced aerodynamic effects depend on relative wind speed.

With a view to solving optimal control problems (or Model-Predictive like control), we choose to use a reward function which is a measure of the distance between the current attitude (p, q, r) with (p_{sp}, q_{sp}, r_{sp}) , the target attitude (similar to the one used in [15]):

$$r(s) = -\min\left(1, \frac{1}{3\Omega_{max}} \|\Omega^* - \Omega\|\right)$$

Ω_{max} being the maximal angular rate that we want to reach for the quadcopter, and Ω is the angular rate vector (p, q, r) which is part of the full state s of the quadcopter.

4.2. Neural net architecture

Neural nets, such as multiple layer perceptrons (MLP) with RELU activation, can efficiently encode all piecewise-affine functions [46]. It is also known [47] that the solution to a quadratic optimal control (MPC) problem for linear-time invariant system is piecewise-affine. Furthermore, there are good indications that this applies more generally, in particular for non-linear systems [48]. This

naturally leads to thinking that MLPs with RELU networks are the prime candidates for controlling the attitude with distance to the objective as cost (or reward). In some ways, the resulting piecewise-affine function encodes various proportional gains that should be best adapted to different subdomains of states, so as to reach an optimal cumulated (and discounted, here with discount rate $\gamma = 0.99^1$) distance to the target angular rates, until the end of training.

In theory [49], one could find a good indication of the architecture of the neural net in such situations, but the bounds that are derived in [49] are not convenient for such a highly complex system. It is by no means obvious what architecture will behave best, both for training and for actual controller performance, although a few authors argue that deeper networks should be better, see e.g. [50].

Architectures that have been reported in the literature for similar problems are generally alike. In [16], the neural net is a Multi-Layer Perceptron (MLP) with two layers of 64 neurons each, and with tanh activation function. In [17], the part of the controller which is a neural net is a MLP with two layers of 96 neurons each and tanh activation function, whose input states (observation space) are all states plus the control. In [18], the hover mode neural net controller, which is the most comparable to our work, is a MLP with two layers of 400 and 300 neurons respectively, with RELU activation for hidden layers and tanh for the last layer. In [19], the resulting architecture is a two layers MLP with 128 neurons on each layer, and RELU activation function.

We will report experiments with one to four layers, and with 4, 8, 16, 32 or 64 neurons per layer, with RELU activation function (except for the rescaling of the output, using tanh). We limit the reporting of our experiments to these values since we observed that these were enough to find best (and worst) behaviours.

4.3. Training algorithms

The first three algorithms we are discussing in Section 7, DDPG [11], SAC [4] and TD3 [5] are all off-policy, actor-critic methods, which are generally considered to be better suited for control applications in robotics [7] (DDPG is used for instance in [17]). Because of its effectiveness in practice, observed by many authors, e.g. [15] for attitude control, we also compare with the on-policy Proximal Policy Optimisation [12], also used for similar applications in [16] and in [19].

DDPG is the historical method for continuous observation and action space applications to control, SAC and TD3 being improvements of DDPG. For instance, SAC regularizes the reward with an entropy term that is supposed to reduce the need to fine hyper-parameter tuning.

¹All other parameters, learning rates in particular are the standard ones of Stable Baselines 2.7.0

Let us now describe the training mechanism: we call *query signal* the function describing the prescribed angular rates at any given time. We model this signal by a constant plateau, of magnitude chosen randomly between -0.6 and 0.6 radians per second, and duration chosen randomly between 0.1 and 1 second. We are training over a time window of 1 second (a training episode) during which the query signal is a constant plateau followed by a value of 0 until the end of the episode. We chose to report on training where these query signals are used independently on pitch, roll and yaw. We tested joint queries as well but do not report specifically the corresponding results since we observed no significant difference.

Controls are updated every 0.03 seconds, and we simulate the full state of the quadrotor, using a Runge Kutta of order 4 on Equation (1) with a time step of 0.01 seconds.

The evaluation of the controller is made on similar query signals, but on time windows that last 20 seconds, with a query signal generated according to a more general class of queries (see below). Query signals on such longer time windows could also be considered for training : [15] refers to this approach as "continuous mode" and reports much poorer performance compared to the "episodic mode" with 1 second queries. We therefore decided to report only on episodic mode training.

Variable	Unit	Lower Bound	Higher Bound
z	m	-1000	+inf
u	$m \cdot s^{-1}$	-30	30
v	$m \cdot s^{-1}$	-30	30
w	$m \cdot s^{-1}$	-30	30
ϕ	rad	$-\pi$	π
θ	rad	$-\pi$	π
ψ	rad	$-\pi$	π
p	$rad \cdot s^{-1}$	-5π	5π
q	$rad \cdot s^{-1}$	-5π	5π
r	$rad \cdot s^{-1}$	-5π	5π
cmd_ϕ	PWM	-400	400
cmd_θ	PWM	-400	400
cmd_ψ	PWM	-1000	1000
F	N	0	52428

Table 2: State and action space bounds

Such query classes are characterised by three distributions A , D and S for respectively the amplitude and duration of stable plateaus, and the step amplitude between successive stable plateaus. These distributions are the same for each axis. We define three different classes of queries (where $U(a,b)$ denotes the Uniform distribution of support $[a,b]$): easy ($A = U(-0.2, 0.2)$, $D = U(0.5, 0.8)$, $S = U(0, 0.3)$), medium ($A = U(-0.4, 0.4)$, $D = U(0.2, 0.5)$, $S = U(0, 0.6)$) and hard ($A = U(-0.6, 0.6)$, $D = U(0.1, 0.2)$, $S = U(0, 0.9)$). Our query generator actually changes the joint distribution of amplitude and duration of stable plateaus by filtering out those queries which would

make the roll, pitch and yaw go through singular values in the Euler angles description of the dynamics.

The initial states are sampled in rather large intervals of values. These values as well as the maximal magnitudes of states are given in Table 2:

5. Robust Signal Temporal Logic with Aggregates

To formalize the behavioral properties of the closed-loop system we defined our own flavor of Signal Temporal Logic [51]. Our logic is mainly inspired by two preexisting works [52] and [53]. From [52] we reuse the notion of aggregate operators over sliding windows and extend it with a robust quantitative semantics, where the original paper only defined a classic Boolean semantics for the language. From [53] we reuse the idea of time-averaged robustness, and propose a new Until operator which combines both spatial robustness (instantaneous falsification margin at time t) and temporal robustness (robustness of the property to time delays over signals).

In this logic formulas are interpreted over piecewise-constant signals, whereas they were interpreted over piecewise-linear signal traces in [52]. Our logic's semantics can however be extended to piecewise-linear signals without significant issue.

5.1. Abstract Syntax

Terms, Formulas and Aggregates:

$$\tau ::= c \mid x \mid f(\tau_1, \dots, \tau_n) \quad (7)$$

$$\mid ite(\phi, \tau_1, \tau_2) \quad (8)$$

$$\mid \text{On}_{[a,b]} \psi_{\mathbb{R}} \quad (9)$$

$$\mid \psi_{\mathbb{R}} U_{[a,b]}^{d_{\mathbb{R}}} \phi \quad (10)$$

$$\mid \tau \downarrow U_{[a,b]}^{d_{\mathbb{R}}} \phi \quad (11)$$

$$\phi ::= \top \mid \perp \quad (12)$$

$$\mid \tau > 0 \quad (13)$$

$$\mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \quad (14)$$

$$\mid \text{On}_{[a,b]} \psi_{\mathbb{B}} \quad (15)$$

$$\mid \psi_{\mathbb{B}} U_{[a,b]}^{d_{\mathbb{B}}} \phi \quad (16)$$

$$\mid \phi_1 \downarrow U_{[a,b]}^{d_{\mathbb{B}}} \phi_2 \quad (17)$$

$$\mid \phi_1 \bar{U}_{[a,b]} \phi_2 \quad (18)$$

$$\psi_{\mathbb{R}} ::= \text{Min } \tau \mid \text{Max } \tau \quad (19)$$

$$\psi_{\mathbb{B}} ::= \text{Forall } \phi \mid \text{Exists } \phi \quad (20)$$

with $(a, b) \in \mathbb{R}^2$ and $a \leq b$, $d_{\mathbb{R}} \in \mathbb{R}$, $d_{\mathbb{B}} \in \mathbb{B}$.

A term τ is either: a constant c , a signal x or a combinatorial function f applied to a number of terms (7); an if-then-else selection of a term based on a Boolean condition 8; a value computed from a time interval $[a, b]$ using some numeric aggregate $\psi_{\mathbb{R}}$ (9); an "aggregate until" term $\psi_{\mathbb{R}} U_{[a,b]}^{d_{\mathbb{R}}} \phi$ which computes a real value over a time interval $[a, b]$ using a numeric aggregate $\psi_{\mathbb{R}}$ (10); or a "time-point

until" $\tau \downarrow U_{[a,b]}^{d_{\mathbb{R}}} \phi$, which samples the value of a term when a formula becomes true (11).

A formula ϕ is either: a logical constant true \top or false \perp (12); the comparison of a term to zero (13); the negation of a formula, or the conjunction or disjunction of a formula (14); an *aggregate* computed from a time interval $[a, b]$ using some logic aggregate $\psi_{\mathbb{B}}$ (15); an *aggregate until* formula $\psi_{\mathbb{B}} U_{[a,b]}^{d_{\mathbb{B}}} \phi$ which computes a truth value over a time interval $[a, b]$ using a logical aggregate $\psi_{\mathbb{B}}$ (16); a *sample until* $\phi \downarrow U_{[a,b]}^d \phi$, which samples the value a formula when some formula becomes true (17); or an *average until* of a formula ϕ_1 computed over time interval $[a, b]$ until ϕ_2 becomes satisfied (18). A *numeric aggregate* $\psi_{\mathbb{R}}$ is either the min or max of a term τ (19). A *logic aggregate* $\psi_{\mathbb{B}}$ is either the Forall or Exists of a formula ϕ (20).

In addition to these core operators, the logic provides a number of derived operators defined in terms of the core operators.

The *term lookup* operator is defined as follows:

$$D_a^d \tau = \tau \downarrow U_{[a,a]}^d \top \quad (21)$$

The *formula lookup* operator is defined as follows:

$$D_a^d \phi = \phi \downarrow U_{[a,a]}^d \top \quad (22)$$

The original STL's Globally, Finally and Until operators are defined as follows:

$$F_{[a,b]} \phi = \text{On}_{[a,b]} \text{Exists } \phi \quad (23)$$

$$G_{[a,b]} \phi = \text{On}_{[a,b]} \text{Forall } \phi \quad (24)$$

$$\phi_1 U_{[a,b]}^{STL} \phi_2 = (\text{Forall } \phi_1) U_{[a,b]}^{\perp} \phi_2 \quad (25)$$

5.2. Interpretation Structures

Terms and formulas are interpreted over total piecewise-constant functions $\sigma : \mathbb{R} \rightarrow \mathbb{R}^n$, which assign a value to a tuple of signals $X = (x_1, \dots, x_n)$ of size n at any time $t \in \mathbb{R}$. However, for practical reasons we only consider total piecewise-constant functions defined by a finite sequence of *breakpoints*:

$$Bkpts = \llbracket (t_j, X_j) \mid j \in [0, M-1], (t_j, X_j) \in (\mathbb{R}, \mathbb{R}^n), \rrbracket$$

where $t_j < t_{j+1}$ for all j , and by a default value $X_d \in \mathbb{R}^n$, as follows:

$$\sigma(t) = \begin{cases} X_d & \text{if } t \in (-\infty, 0) \\ X_j & \text{if } t \in [t_j, t_{j+1}) \\ X_{M-1} & \text{if } t \in [t_{M-1}, +\infty) \end{cases} \quad (26)$$

We use the following notations:

- $T_{\sigma} = \llbracket t_j \mid j \in [0, M-1] \rrbracket$ is its sequence of timesteps,
- $\sigma(x_i, t)$, by abuse of notation, is the i^{th} coordinate of $\sigma(t)$, i.e. the value of signal x_i at time t .

5.3. Standard Semantics

5.3.1. Term Semantics

Assuming some fixed trace σ , the interpretation function for terms

$$sem() : \tau \rightarrow \mathbb{R} \rightarrow \mathbb{R} \quad (27)$$

is defined inductively as follows:

$$sem(c)(t) = c \quad (28)$$

$$sem(x_i)(t) = \sigma(x_i, t) \quad (29)$$

$$sem(f(\tau_1, \dots, \tau_n))(t) = f(sem(\tau_1)(t), \dots, sem(\tau_n)(t)) \quad (30)$$

$$sem(ite(\phi, \tau_1, \tau_2))(t) = \begin{cases} sem(\tau_1)(t) & \text{if } sem(\phi)(t) \\ sem(\tau_2)(t) & \text{otherwise} \end{cases} \quad (31)$$

$$sem(\text{On}_{[a,b]} \psi_{\mathbb{R}})(t) = sem(\psi_{\mathbb{R}})([t+a, t+b]) \quad (32)$$

$$sem(\psi_{\mathbb{R}} U_{[a,b]}^d \phi)(t) = \begin{cases} sem(\psi_{\mathbb{R}})([t, t']), \text{ where} \\ t' \in [t+a, t+b] \text{ smallest} \\ \text{instant st. } sem(\phi)(t') = \top \\ d \text{ if no such } t' \text{ exists.} \end{cases} \quad (33)$$

$$sem(\tau \downarrow U_{[a,b]}^d \phi)(t) = \begin{cases} sem(\tau)(t'), \text{ where} \\ t' \in [t+a, t+b] \text{ smallest} \\ \text{instant st. } sem(\phi)(t') = \top \\ d \text{ if no such } t' \text{ exists.} \end{cases} \quad (34)$$

The semantics of numeric aggregates is defined over intervals as follows:

$$sem(\text{Min } \tau)([a, b]) = \min_{t \in [a,b] \cap T_{\sigma}} (sem(\phi)(t)) \quad (35)$$

$$sem(\text{Max } \tau)([a, b]) = \max_{t \in [a,b] \cap T_{\sigma}} (sem(\phi)(t)) \quad (36)$$

Since traces are total piecewise-constant functions, defined by a finite number of samples, and all operators have default values and are hence total, interpretation functions are also total function, and evaluating a numeric or logic aggregates requires inspecting only a finite number of timesteps and yields an exact result.

5.3.2. Formula Semantics

Assuming a fixed trace σ , formula semantics is given by the function:

$$sem() : \phi \rightarrow \mathbb{R} \rightarrow \mathbb{B} \quad (37)$$

defined inductively as follows:

$$\text{sem}(\top)(t) = \top \quad (38)$$

$$\text{sem}(\perp)(t) = \perp \quad (39)$$

$$\text{sem}(\tau > 0)(t) = \text{sem}(\tau)(t) > 0 \quad (40)$$

$$\text{sem}(\neg\phi)(t) = \top \text{ iff } \text{sem}(\phi)(t) = \perp \quad (41)$$

$$\text{sem}(\phi_1 \wedge \phi_2)(t) = \top \begin{cases} \text{iff } \text{sem}(\phi_1)(t) = \top \\ \text{and } \text{sem}(\phi_2)(t) = \top \end{cases} \quad (42)$$

$$\text{sem}(\phi_1 \vee \phi_2)(t) = \top \begin{cases} \text{iff } \text{sem}(\phi_1)(t) = \top \\ \text{or } \text{sem}(\phi_2)(t) = \top \end{cases} \quad (43)$$

$$\text{sem}(\text{On}_{[a,b]} \psi_{\mathbb{B}})(t) = \text{sem}(\psi_{\mathbb{B}})([t+a, t+b]) \quad (44)$$

$$\text{sem}(\psi_{\mathbb{B}} \text{U}_{[a,b]}^b \phi)(t) = \begin{cases} \text{sem}(\psi_{\mathbb{B}})([t, t']), \text{ where } t' \\ \text{in } [t+a, t+b] \text{ smallest} \\ \text{timestep st. } \text{sem}(\phi)(t') = \top \\ b \text{ if no such } t' \text{ exists.} \end{cases} \quad (45)$$

$$\text{sem}(\phi_1 \downarrow \text{U}_{[a,b]}^b \phi_2)(t) = \begin{cases} \text{sem}(\phi_1)([t, t']), \text{ where } t' \\ \text{in } [t+a, t+b] \text{ smallest} \\ \text{timestep st. } \text{sem}(\phi_2)(t') = \top \\ b \text{ if no such } t' \text{ exists.} \end{cases} \quad (46)$$

The classic semantics for the *average until* operator $\phi_1 \bar{\text{U}}_{[a,b]} \phi_2$ is defined exactly as the original *STL Until* semantics.

The semantics of logic aggregates is defined over intervals $[a, b]$ as follows:

$$\text{sem}(\text{Forall } \phi)([a, b]) = \bigwedge_{t \in [a, b]} \text{sem}(\phi)(t) \quad (47)$$

$$\text{sem}(\text{Exists } \phi)([a, b]) = \bigvee_{t \in [a, b]} \text{sem}(\phi)(t) \quad (48)$$

5.4. Robust semantics

The robust semantics

$$\rho() : \phi \rightarrow \mathbb{R} \rightarrow \mathbb{R} \quad (49)$$

only concerns Boolean formulas, and is defined inductively as follows:

$$\rho(\top)(t) = +\infty \quad (50)$$

$$\rho(\perp)(t) = -\infty \quad (51)$$

$$\rho(\tau > 0)(t) = \text{sem}(\tau)(t) \quad (52)$$

$$\rho(\neg\phi)(t) = -\rho(\phi)(t) \quad (53)$$

$$\rho(\phi_1 \wedge \phi_2)(t) = \min(\rho(\phi_1)(t), \rho(\phi_2)(t)) \quad (54)$$

$$\rho(\phi_1 \vee \phi_2)(t) = \max(\rho(\phi_1)(t), \rho(\phi_2)(t)) \quad (55)$$

$$\rho(\text{On}_{[a,b]} \psi_{\mathbb{B}})(t) = \rho(\psi_{\mathbb{B}})([t+a, t+b]) \quad (56)$$

$$\rho(\psi_{\mathbb{B}} \text{U}_{[a,b]}^b \phi)(t) = \begin{cases} \rho(\psi_{\mathbb{B}})([t, t']) \text{ where } t' \text{ in} \\ [t+a, t+b] \text{ is the smallest} \\ \text{timestep st. } \text{sem}(\phi)(t') \\ \rho(b) \text{ if no such } t' \text{ exists.} \end{cases} \quad (57)$$

$$\rho(\phi_1 \downarrow \text{U}_{[a,b]}^b \phi_2)(t) = \begin{cases} \rho(\phi_1)(t') \text{ where } t' \text{ in} \\ [t+a, t+b] \text{ is the smallest} \\ \text{timestep st. } \text{sem}(\phi_2)(t') \\ \rho(b) \text{ if no such } t' \text{ exists.} \end{cases} \quad (58)$$

$$\rho(\phi_1 \bar{\text{U}}_{[a,b]} \phi_2)(t) = \begin{cases} (b-t') * \rho(\text{On}_{[t, t']} \text{Forall } \phi_1)(0) \\ \text{where } t' \in [t+a, t+b] \\ \text{is the smallest timestep} \\ \text{st. } \text{sem}(\phi_2)(t') \\ -\infty \text{ if no such } t' \text{ exists.} \end{cases} \quad (59)$$

$$(60)$$

The robust semantics for logic aggregates is defined as follows:

$$\rho(\text{Forall } \phi)([a, b]) = \min_{t \in [a, b]} \rho(\phi)(t) \quad (61)$$

$$\rho(\text{Exists } \phi)([a, b]) = \max_{t \in [a, b]} \rho(\phi)(t) \quad (62)$$

The robust interpretation of terms is just their standard interpretation, except for *timepoint until* and *aggregate until* operators where instead of recursing on the standard interpretation of Boolean formulas, we recurse on their robust interpretation.

5.5. Implementation

We implemented a code generator for the logic, which generates highly efficient python code allowing to compute the standard and robust semantics of STL formulas on traces. Given an STL formula (or term) as input, the code generator produces a Python 3.x class definition which implements the standard and robust semantics evaluation rules for that formula. The class takes a trace as constructor argument (i.e. a piecewise constant function specified a sequence of breakpoints and a default value as defined in section 5.2), and offers an `eval` method allowing to compute the standard or robust semantics of the formula at any time step.

The generated code uses a number of techniques for efficiency:

- Constant folding,
- When translating a specification containing several formulas and terms, the code generator implements common subformulas and subterms sharing between all toplevel formulas.
- We leverage the fact that in practice, a same formula will be evaluated on sequences of strictly increasing timesteps, and use an incremental method for the evaluation of sliding window aggregates: when evaluating a Max aggregate (resp. Min, Forall or Exists) at time t , the aggregate term is evaluated on interval $[t+a, t+b]$ and we cache the result (t_{Max}, x_{Max}) , indicating at which instant the Max value was reached in $[t+a, t+b]$. When the aggregate is evaluated again at $t' > t$, we distinguish the following cases:
 - if $t_r \in [t' + a, t' + b]$, we evaluate the aggregate on $[t+b, t'+b]$ and return $Max(x_{Max}, x_{Max}')$, cache new result,
 - if $t_r \in [t + a, t' + a]$, we evaluate the aggregate on the full window $[t' + a, t' + b]$ and cache the result.
- Last, the generated code uses numpy arrays exclusively and contains Numba annotations for all data structures and classes, allowing to use Numba to JIT the evaluation code using LLVM. This JIT optimization provides a 10x to 20x performance boost over the interpreted python version.

6. Formal performance criteria

Designing a controller for a specific application requires balancing multiple criteria such as rising time, overshoot, steady error, etc. In order to quantify rigorously the performance of the learned controller, we formalized requirements using the logic presented in section 5.

A first set of formulae allows to identify instants when a query signal q becomes stable for T time units, and whether q goes up or down at any instant (with ϵ and d two small constants), and the step size:

$$stable(q) = (On_{[0,T]} Max q) - (On_{[0,T]} Min q) < d \quad (63)$$

$$stableup(q) = (D_{-\epsilon}^\perp \neg stable(q)) \wedge stable(q) \quad (64)$$

$$up(q) = q - (D_{-\epsilon}^0 q) > 0 \quad (65)$$

$$down(q) = q - (D_{-\epsilon}^0 q) \leq 0 \quad (66)$$

$$step(q) = ite(stableup(q), q - D_{-\epsilon}^0 q, 0) \quad (67)$$

We consider an angular rate signal x as acceptable if it does not overshoot a stable query q by more than $\alpha\%$

of the step size on $[0, T_1]$, and does not stray away from a stable query q by more than $\beta\%$ of the step size on $[T_1, T]$:

$$stableup(q) \wedge up(q) \implies On_{[0,T_1]} Max (x - q) < \alpha step(q) \quad (68)$$

$$stableup(q) \wedge down(q) \implies On_{[0,T_1]} Max (q - x) < \alpha step(q) \quad (69)$$

$$stableup(q) \implies On_{[T_1,T]} Max \|x - q\| < \beta step(q) \quad (70)$$

We define the rising time RT as the time it takes for x to first reach q within $\gamma\%$:

$$ite(stableup(q), t - (t U_{[0,T]}^{+\infty} \|(x - q)\| < \gamma q), +\infty) \quad (71)$$

Figure 2 illustrates the formalised notions and parameters.

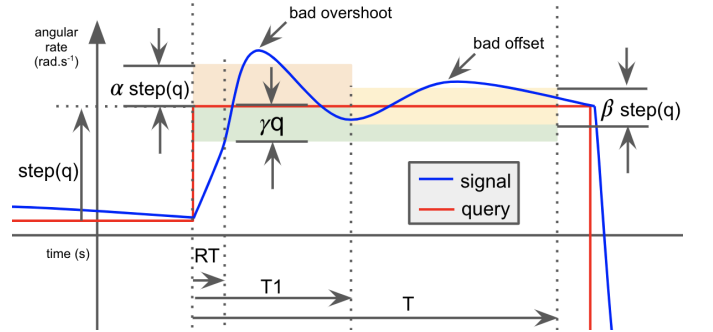


Figure 2: Property parameters T_1 , T , RT , α , β , γ .

Using observers code generated from these specifications, we compute statistics about property violations and associated robustness margins on angular rate signals and queries on pitch, yaw and roll axis of the system, acquired at regular intervals during the training of the controller. For evaluation each property $P(x, q)$ is wrapped in a *globally* modality over the episode length yielding $G_{[0, episode_length]} P(x, q)$. Automating the computation of these behavioral metrics is essential in allowing to scale up the hyper-parameter space exploration and identify the best controller according to objective measurements.

7. Experimental setup

7.1. Implementation

We have developed a platform² with the purpose of running experiments in a reproducible and scalable way, becoming an integration layer between the different moving parts in both training and testing. From a technological standpoint the platform is based on the Stable Baselines 2.7.0 reinforcement learning library [54] itself based

²The full code is available as open source at <https://github.com/uber-research/rl-controller-verification>.

on Tensorflow [55], all of our code is in Python and we used Bazel [56] as build system. We used Tensorboard to monitor losses and the internal dynamics of the neural networks during the training.

One intermediate goal was to explore the large combinatorial hyperparameter space efficiently, to be able to identify the best hyperparameters values with respect to the STL metrics we defined and to get a better understanding of their impact.

With 4 different algorithms, 20 possible configurations for the network architecture and 3 sets of observed states, our hyperparameters space contains a total of 240 points that need to be trained and tested. The corresponding jobs are dispatched on our Kubernetes cluster [57] where they can run in parallel. Disposing of 1 vCPU on the Cascade Lake platform (base frequency of 2.8 GHz), the 3 millions iterations of a single training job take between 3 and 8 hours to complete. The cluster autoscales with the workload and allowed us to run 1,200 hours worth of training in half a day.

The container images that end up running on the cluster are created, uploaded and finally dispatched in a reproducible manner thanks to the Bazel rules of our Research Platform. Those rules are built on top of the Bazel Image Container Rules [58] and the Bazel Kubernetes Rules [59] and specially designed to generate all the experiment jobs of the hyperparameters analysis.

The training and testing results are automatically uploaded on our cloud storage where they can be browsed for quick inspections, or fed as input for the next pipeline stage. We saved 30 checkpoints per experiment (each file containing 100k training iterations weights between 10KB and 100KB). Including the TensorFlow logs, the training results amount to over 100GB of data.

Each of the 30×240 checkpoints was then evaluated on 100 queries computed by the Query Generator, producing the same number of concrete traces representing the commands and the states over the whole episode. Each set of such traces is about 600k hence it yields total of 60MB per checkpoint. Finally each of the $30 \times 240 \times 100$ traces was evaluated with STL properties observer to compute synthetic metrics: aggregating the 100 traces of a single checkpoint produced a 150KB file and required approximately 45 minutes. The checkpoint-specific CSV files were further aggregated in experiment-specific and round-specific checkpoints for final visual inspection.

7.2. Interactive browsing of the experiments database

We want to understand what correlations exist between controller performance and the way it has been trained, and for this, we used Hiplot [60] for browsing through the enormous number of parameters and data generated. We show in Figure 3 how we used Hiplot in an interactive manner for verifying our hypotheses. Each parameter and performance measure is represented by a column in the graph generated by Hiplot from our database. For each

parameter, either fixed or free, choosing intervals of values for each performance measure creates lines that link parameter values to performance values within the chosen intervals. The number of entries in the database (i.e. the number of controllers) that satisfy the constraints is also shown, as well as the table of all their corresponding parameters and performance values.

For instance, we used Hiplot to select the "best" networks, filtering the data set of controllers, only retaining the ones with better success in offset, overshoot and rising times altogether, with respect to the best PIDs. This resulted in two neural nets with much better performances than the PIDs on offset and on rising time, as we will discuss in Section 8.

8. Experimental results

8.1. Performance metrics

Each controller is evaluated on a hundred evaluation episodes using STL observers defined in Equations (68) to (71), where parameters are set to $\alpha = 10\%$, $\beta = 5\%$ and $\gamma = 5\%$, $T = 0.5s$, $T_1 = 0.25s$, $\epsilon = 0.01s$, $d = 0.005$. For each evaluation episode the following statistics are computed over all stable query plateaus:

- average and maximum overshoot percentage relative to the query step size,
- average and maximum offset percentage relative to the query step size,
- average and maximum rising time values in seconds (only for plateaus where the signal actually reaches $\gamma\%$ of the query within $[0, T]$).

For each metric (overshoot, offset, rising time), we compute the *success percentage* % OK, i.e. the percentage of stable plateaus of the episode for which the controller behaviour satisfies the specification.

Then, episode-level statistics are further averaged, yielding results presented in the tables of the following sections, where columns represent:

- avg (resp. max) overshoot: is the per-episode-average of the average (resp. maximum) overshoot values,
- avg (resp. max) offset: is the per-episode-average of the average (resp. maximum) offset values,
- avg (resp. max) rising time: is the per-episode-average of the average (resp. max) rising time,
- % OK offset (resp. overshoot, rising time): is the per-episode-average of the success percentage for the offset (resp overshoot, rising time) metric.

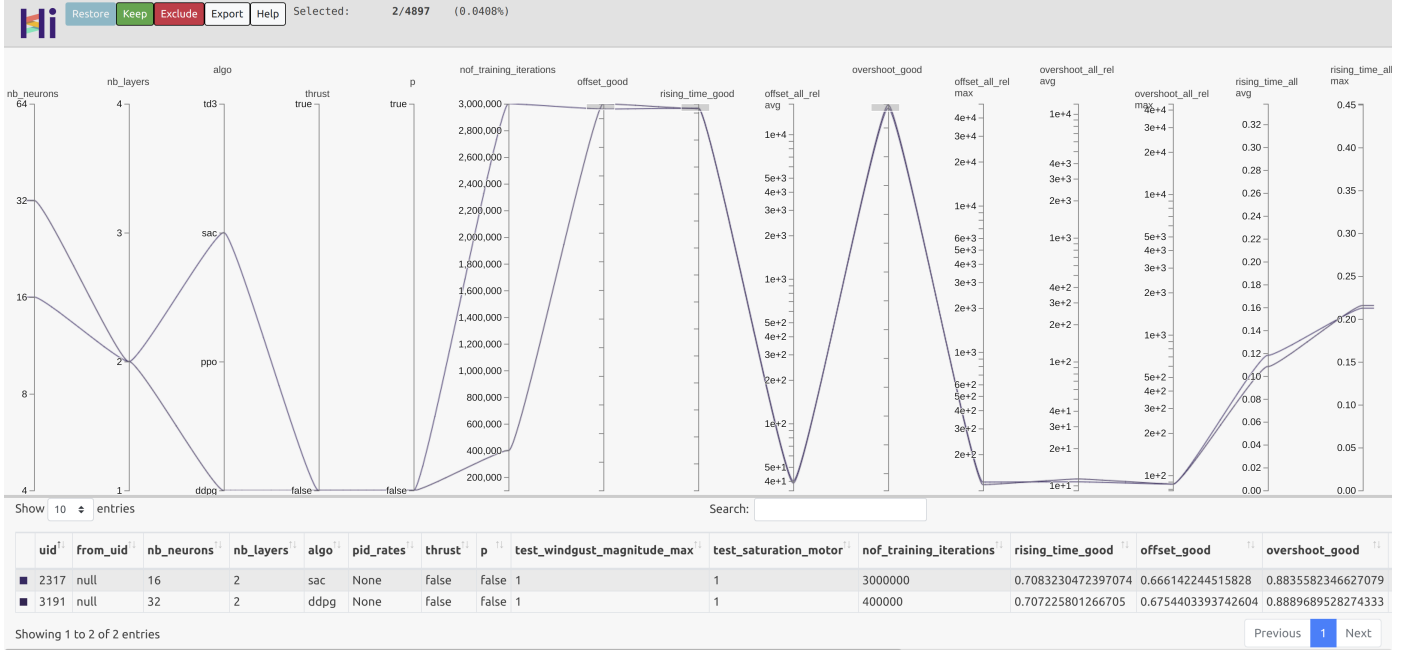


Figure 3: Hiplot interactive session




Figure 4: PID2 controller query tracking

8.2. Performance of nominal-trained networks in nominal test case

8.2.1. Overall best performance comparison

The PID performance metrics in the nominal case are reported in the first two lines of Table 3 to serve as a reference point for neural controller evaluation. Examples of query tracking behavior are given in Section 8.2.1 for reference.

PID2 reaches within 5% of the target state for about 70% of the queries, and is relatively slow with an average rising time of 0.44s. PID1 in comparison reaches within 5% of the target state for only 8% of the queries, with a (very slightly) better rising time. Overshoot success rates are really good for both PIDs (95-100% OK). Offset success rates are bad (1-3% OK), due to their slow convergence.

We will hence use PID2 as a reference for discussing neural controller performance.

The comparison between the best networks and the PIDs is also reported in Table 3.

We see that our neural nets provide much quicker controls, with an average rising time of about a fourth to a fifth of the rising time for the two PIDs, although with a negligible offset. This is at the expense of a slightly less good performance on the maximum overshoot at least for SAC and DDPG trained networks, with respect to PID2 (our neural nets are still much better than PID1). Results are far less good, in particular concerning overshoots, with PPO and TD3 trained networks. This is also visible when comparing signals between Figure 5 and Section 8.2.1. Somehow, neural nets exhibit extreme reactivity as well as good asymptotic convergence, but show some very short-lived "spikes", as in the sample trajectory shown in Figure 5.

When we filter the neural nets meeting or exceeding the performances of PID2, many networks remain, among which the best are:

- DDPG $64 \times 64 \times 64 \times 64$ trained for 1,500,000 iterations (and also DDPG 32×32 , 400,000 iterations) on the three-dimensional observation space ($p - p_{sp}, q - q_{sp}, r - r_{sp}$)
- SAC $32 \times 32 \times 32 \times 32$ (and SAC 32×32 and 16×16 trained for 3,000,000 iterations coming very close) trained for 2,900,000 iterations on the same three-dimensional observation space

algo	OK rising t.	OK off.	OK overshoot	avg rising t.	avg off.	avg overshoot	max rising t.	max off.	max overshoot
PID2	70.52	0.52	100.00	0.41	20.22	0.00	0.48	25.26	0.00
PID1	7.73	3.44	94.88	0.33	58.93	3.91	0.38	138.44	50.56
SAC	97.35	99.54	96.55	0.08	0.12	0.50	0.21	2.44	6.21
DDPG	99.05	98.59	98.21	0.08	0.21	0.23	0.17	3.97	3.91
PPO	96.96	97.26	91.98	0.08	0.43	1.41	0.19	7.08	11.93
TD3	96.70	86.80	88.16	0.09	2.40	2.13	0.22	18.08	20.09

Table 3: PIDs and overall best networks performance (all in % except rising t. in seconds)




Figure 5: Neural controller behaviour (sac, 2 layers, 16 neurons per layer, 3M iterations)

8.2.2. Training algorithm influence

We observe in Table 3 that PPO and TD3 do not show as good performance as SAC (and even DDPG), moderating the conclusion of [15], and the common belief that TD3 should improve performance of neural net control. We have for now no explanation for this, largely because we have not been able (which is also the case in [15]) to get rid of the overshoot spikes, even using SAC which does some amount of regularization, or TD3 which should lead to more stable solutions, potentially at the expense of a slower convergence rate. In terms of optimal control, if the neural net controller were trained with correctness objectives³, these spikes would certainly be much smaller and appear only at the very beginning of plateaus.

8.2.3. Convergence of the training algorithms

We show in Figure 6 the evolution of the three main performance measures, the OK overshoot, OK offset and OK rising time, for one of the best network architecture and training algorithm, SAC 32×32 neurons. The three metrics improve quickly and almost stabilize in the first 1,000,000 iterations.

³Future work to cope with this phenomenon includes improving the reward function using our STL observers, and adding some more regularization during training.




Figure 6: Performance of SAC 32x32 on dim 3 observation space trained neural nets w.r.t. the number of iterations

8.2.4. Observation state influence

Of course, for a given number of iterations, smaller-dimensional observation states yield better quality of the sampling. Still, we observe that using a Markovian state or the simpler three-dimensional state space (err_p, err_q, err_r) does not change significantly the performance of the best neural nets obtained, see Table 4, although the 3-dimensional observation space gives slightly better performance overall. In fact, we even get a worse performance with the 7-dimensional full state, mostly because of the difficulty to sample this higher dimensional space, and identify the subtle second-order effects of some of these states on angular rates.

8.2.5. Neural net architecture influence

First, we observe that almost none of the single-layer neural nets seem to converge to a correct controller (see e.g. Figure 7). At 64 neurons, 1 hidden layer networks seem to exhibit some good behaviour, but still far from any of the e.g. two-layers neural nets.

Still, 3-layers and even 4-layers networks do not seem to exhibit much better behaviour than the "best" 2-layers networks, with 16 or 32 neurons each, although they converge faster. Recently Sinha et al. in [61] empirically observed the performance of SAC have a peak using 2 layers MLP and their explanation for this result relies on the

algo	dim	OK rising t.	OK off.	OK overshoot	avg rising t.	avg off.	avg overshoot	max rising t.	max off.	max overshoot
SAC	3	98.88	98.01	98.58	0.09	0.28	0.18	0.19	5.01	3.28
SAC	6	97.24	98.70	94.82	0.09	0.25	0.81	0.19	4.72	8.07
SAC	7	94.26	94.57	97.79	0.09	0.80	0.33	0.25	11.56	5.75

Table 4: Influence of the observable space dimension (all in % except rising t. in seconds)



Figure 7: OK rising t. for our best SAC network wrt number of training iterations for different architectures

Data Processing Inequality hence the fact that mutual information between layers decreases with depth. This will have to be further investigated in our framework.

8.3. Performance of nominal-trained networks in non-nominal test cases

We now assess the robustness of our PIDs and "best" neural nets (trained in nominal situations as discussed in Section 8.2) to perturbed, non-nominal conditions, without training the neural nets nor changing the gains of PIDs to cope specifically for the new situation. We report the same performance measures as the ones used in the nominal case, in the test cases where a perturbation can happen, at the start of any new plateau along the 20 second episodes that we are observing (which can contain about 30 different target angular states, or plateaus, to reach within a short time). We take maxima and averages of these measures on 100 such queries as before.

8.3.1. Robustness to partial motor failures

We report in Table 5 results where the perturbation is a partial power loss of motor 1, down to 80% of its maximal power.

For this case of partial motor failure, our best SAC trained neural net behaves much better than our two PIDs: it keeps on reaching plateaus within 0.5 seconds for about 94% of the time, whereas even the best PID goes down to less than 60% success rate. Our network is even better when it comes to satisfying offset constraints (82% of the time) whereas the PIDs almost never comply. Performances concerning overshoot are comparable, even though the PIDs are very slightly better, but this only concerns

cases where PIDs actually reach the target state, which is the case much less often. Essentially, the best neural nets that have been trained under nominal conditions show very little degradation of performance when a partial failure occurs.

8.3.2. Robustness to wind gusts

We present in Table 5 results where the perturbation is the occurrence of randomly chosen wind gusts (as described in Section 3.3.1) of magnitude up to 10 m.s^{-1} from any fixed direction in the inertial frame.

The PIDs and the neural nets exhibit the same kind of minor loss of performance, and the nominal trained neural nets are still far superior to the two PIDs.

8.4. Performance of non-nominal-trained networks

8.4.1. Training under partial motor failures

In what follows, we train the attitude controller to sustain partial motor failures adding the magnitude of the power loss (1 extra dimension) to the observation states discussed in Section 4.1. We report the performance measures obtained in the non-nominal case in Table 6. The concern one may have is that, training the neural net in more various conditions (nominal and non-nominal), the resulting controller may exhibit lower performance. We thus report the same performance measures for neural nets trained with potential motor failures, in nominal situations, e.g. when no power loss happens, see Table 7

We see that we still achieve much better performance than PIDs, but that we are only similar and even slightly worse than the neural nets trained in nominal conditions, both in nominal conditions (compare Table 7 to Table 3) and in non-nominal conditions (compare Table 6 to Table 5). Understanding this non intuitive behaviour and improving the training in this case is left for future work.

8.4.2. Training under wind gusts

In what follows, we train the attitude controller to sustain wind gusts up to 10 m.s^{-1} in any direction, adding to the observation states we discussed in Section 4.1 the wind gust magnitude and directions (4 additional dimensions) plus the linear velocities of the quadcopter (u, v

mode	algo	OK rising t.	OK off.	OK overshoot	avg rising t.	avg off.	avg overshoot	max rising t.	max off.	max overshoot
windgust	SAC	98.50	97.40	98.48	0.09	0.41	0.22	0.17	6.78	3.81
windgust	PID2	70.32	0.65	100.00	0.41	20.16	0.00	0.48	24.25	0.00
windgust	PID1	6.55	2.97	94.28	0.33	60.23	4.43	0.38	146.00	57.65
saturation	SAC	94.20	81.58	92.85	0.12	17.63	6.04	0.33	145.91	81.25
saturation	PID2	58.24	2.94	93.98	0.39	34.33	4.66	0.48	129.36	58.46
saturation	PID1	14.50	3.71	92.13	0.30	64.83	5.67	0.40	166.82	70.45

Table 5: Robustness of the best networks and PIDs in case of wind gusts and motor saturation (all in % except rising t. in seconds)

algo	OK rising t.	OK off.	OK overshoot	avg rising t.	avg off.	avg overshoot	max rising t.	max off.	max overshoot
DDPG	89.40	73.82	90.53	0.13	20.61	6.84	0.37	156.87	87.98
SAC	90.93	79.00	90.45	0.12	20.36	7.18	0.35	164.34	95.79

Table 6: Best networks trained for partial motor failures, tested under potential motor failures situations (all in % except rising t. in seconds)

and w , 3 additional dimensions) since they are necessary for determining the relative wind velocity.

We report the performance measures that we get in the non-nominal case in Table 8 and in the nominal case in Table 9.

We see that the SAC and DDPG controller trained with potential wind gusts still behave about as well as the nominal controller (compare Table 9 to Table 3). Surprisingly, the best (SAC) network behaves slightly worse than the nominal-trained SAC network under wind gusts (compare Table 8 to Table 5), where we can see a slight drop of performance in e.g. *OK off.* and *OK overshoot*: it does not seem to be able to learn correctly how to stay close enough to the target plateau, in some cases.

9. Lessons learned

Sampling. First, we observed that we should restrict to a “good” subspace of the (full quadcopter) states that is sufficiently low dimensional for efficient sampling and such that it avoids potentially spurious correlations, while still providing sufficient information for learning. For instance, in the nominal case, the observation space (err_p, err_q, err_r) was found to be the optimal choice. Training depends of course on sampling data, that has to be done on representative data, and on sampling initial states in a large enough space. In order to do this, for better results, we developed a specific query generator, and we sampled initial states in quite large spaces.

Training algorithms. SAC gives very good results as expected. It is most probably more efficient due to entropy regularization that partially cancels spurious correlations, but this has still to be confirmed in more general situations. A lesson for us was that TD3 was not behaving

as well as expected. Our current guess is that TD3 suffers from too much bias on the Q-function estimation at some point in our training environment, or that TD3 needs many more iterations to converge in our case due to bad exploration performance. Recent papers have suggested that action clipping in TD3 can result in poor exploration performance on problems with bounded action spaces (actions on the boundary are too frequently sampled) which has been shown to be remedied by the entropy regularization of SAC or other output scaling and replay buffer sampling approaches that simulate entropy regularization, [62, 63]. Another newly documented [64] undesired behavior of TD3 is to have all Critics converge to a same point in parameter space and degenerate into single-Q-network performance. Without further experiments we cannot say if poor performance is due to action clipping, to critic diversity collapse, or both. Considering SAC works a lot better and also uses dual Q-networks like TD3, it seems more likely that clipping and bad exploration are to blame than diversity collapse.

Quality of deep and shallow controllers. It is actually hard to find good attitude controllers using RL, probably explaining why papers in this area generally only discuss a single neural net controller: we found only 9 out of about 5000 controllers which complied with our specifications. The very last 5% performance seems to be very hard to get because of “spikes” we observed, due to spurious correlations in the fully connected neural net controllers we have been considering. We also note that small and rather shallow (two or three hidden layers) networks were observed to be best trained and to be behaving best for attitude control.

Spurious correlations. Even if the STL metrics gives excellent results for some networks, there are still some spikes, as shown with the behavior of one of our best networks on a simple roll rate query in Figure 8, that we identified to be due to spurious correlations between the errors on one axis and the command on another axis.

algo	OK rising t.	OK off.	OK overshoot	avg rising t.	avg off.	avg overshoot	max rising t.	max off.	max overshoot
SAC	95.35	95.42	97.47	0.09	0.66	0.34	0.22	9.35	5.30
DDPG	94.21	95.17	94.69	0.09	0.83	0.80	0.21	11.29	10.51

Table 7: Performance of best networks trained with potential motor failures, and tested in nominal situations (all in % except rising t. in seconds)

algo	OK rising t.	OK off.	OK overshoot	avg rising t.	avg off.	avg overshoot	max rising t.	max off.	max overshoot
DDPG	99.13	90.04	95.28	0.09	1.75	0.79	0.21	17.46	11.40
SAC	97.67	89.58	92.97	0.10	1.52	1.01	0.28	13.91	11.65

Table 8: Best networks trained for wind gusts conditions, tested under wind gusts conditions (all in % except rising t. in seconds)

algo	OK rising t.	OK off.	OK overshoot	avg rising t.	avg off.	avg overshoot	max rising t.	max off.	max overshoot
DDPG	96.53	91.43	94.83	0.09	1.76	0.96	0.23	17.47	12.41
SAC	95.96	97.09	96.77	0.09	0.42	0.45	0.21	5.57	6.34

Table 9: Best networks trained for wind gusts conditions, tested in nominal conditions (all in % except rising t. in seconds)




Figure 8: Spikes in roll rate control, with one of our best trained neural nets

These spurious correlations can be exposed by training a controller on a single axis, here the roll axis, and showing that they indeed do not appear in that case where correlations cannot possibly be made. This new controller was trained with only the error on the roll rate as input, with the objective of controlling only cmd_ϕ . During training, we have been controlling cmd_ψ and cmd_θ by PIDs. We see in Figure 9 that the controller on roll rate is now almost perfect, showing no spikes.

Another way to expose the spurious correlations is to examine the connections between the neurons of our controller, and in particular to show which ones are above a certain threshold, for a given input. In Figures 10a and 10b, we depicted the case of a 16×16 neural net controller for the roll, pitch and yaw rates, with two different inputs. The red arc is the same connection in the two figures, between some neuron of the second hidden layer and the neuron governing the roll rate output. In both cases it has





Figure 9: Behaviour of a controller trained on roll only.


a high value (weight), but in the first case, Figure 10a, it shows a good correlation with the input that is used, while in the other, Figure 10b, it shows a spurious correlation. The correlation of Figure 10b is deemed “good”, or correct, since the second input linked to the second axis is connected to the second output (on the same axis - connections between neurons are highlighted in the corresponding figures) and the correlation of Figure 10b is deemed spurious because it happens when the third input linked to the third axis is connected to the second input (an error on the pitch axis should not influence an action on the yaw axis).

During the training phase, i.e. gradient descent, the weight of the connection will never converge to something sensible enough: when the network sees the first type of

input, it will increase the importance of this connection while for the second input, it will reduce the importance of the same connection.



(a) A case of good correlation between the input and the roll rate output



(b) A case of spurious correlation between the input and the roll rate output

Figure 10: Two different types of correlations during training

Training for nominal and non-nominal situations. We also observed that there is some amount of robustness built in neural net controllers, suitably trained in nominal conditions, to certain non-nominal situations. We believe this is due to the fact that the controllers which are trained in the nominal case, are actually trained in many different states that appear in non-nominal situations, for the same neural net inputs (e.g. angular rate errors), by using a very wide distribution of initial states during training.

Similar observations on robustness by training from wide initial state distributions were made in [65].

Finally, training neural net controllers to both nominal and non-nominal situations is not an easy endeavor and should be further studied. The difficulty lies in training on sufficiently many non-nominal data, as well as avoiding over-fitting to non-nominal cases: reward distributions can become multi-modal and expectation maximization could be bad in such cases.

For instance, when we saturate a motor, we lose a degree of freedom and we can just hope for, for instance, a good control on the roll and pitch axis, at the expense of some degradation for controlling the yaw rate. Indeed, it is much harder for the quadcopter to generate a moment on the yaw axis than on the roll or on the pitch axis.

When the controller has only been trained in nominal mode (i.e. without any saturation), it can stay for a rather long time far from the query, when used in non-nominal mode (here, with one motor saturated to a portion of its power, as explained in Section 3.2), as shown in Figure 11. When the controller is trained in non-nominal mode, it learns to overcompensate and does not remain far from the query for a long time, see Figure 12. Indeed, when one motor is saturated, the command on one axis will create a moment on another axis that needs to be compensated. This is what the drone successfully learns when trained with a saturated motor.

10. Conclusion

We have presented a complete study of learned attitude controls for a quadcopter using reinforcement learning. In particular we extend previous results by modeling partial motor failure as well as wind gusts, and generating extensive tests of various network architectures, training algorithms and hyperparameters using a flexible and robust experimental platform. We also present a precise evaluation mechanism based on robust signal temporal logic observers, which allows us to characterize the best options for training attitude controllers. Results show that learned controllers exhibit high quality over a range of query signals, and are more robust to perturbations than PID controllers.

The immediate next step will be to start using STL-derived reward signals during training on the most promising architectures, and try to improve training under non-nominal situations.

Finally, because we use an explicit ODE model, we can hope to discuss formal reachability properties of the complete controlled system, using or elaborating on approaches such as [66] and [41].

References

- [1] D. Bertsekas, Reinforcement Learning and Optimal Control, Athena Scientific optimization and computation series, Athena Scientific, 2019.
URL <https://books.google.fr/books?id=Z1BIyQEACAAJ>




Figure 11: Controller trained in nominal mode and tested with motor saturation




Figure 12: Controller trained with motor saturation and tested with motor saturation

- [2] E. Kaufmann, A. Loquercio, R. Ranftl, M. Müller, V. Koltun, D. Scaramuzza, Deep drone acrobatics, CoRR abs/2006.05768 (2020).
URL <https://arxiv.org/abs/2006.05768>
- [3] N. Bernini, M. Bessa, R. Delmas, A. Gold, E. Goubault, R. Penne, S. Putot, F. cois Sillion, A few lessons learned in reinforcement learning for quadcopter attitude control, in: In ACM International Conference on Hybrid Systems: Computation and Control, 2021.
- [4] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, S. Levine, Soft actor-critic algorithms and applications, CoRR abs/1812.05905 (2018).
URL <http://arxiv.org/abs/1812.05905>
- [5] S. Fujimoto, H. van Hoof, D. Meger, et al., Addressing function approximation error in actor-critic methods, Proceedings of Machine Learning Research 80 (2018).
- [6] Förster, Julian, System Identification of the Crazyflie 2.0 Nano Quadcopter, B.S. Thesis, Institute for Dynamic Systems and Control, Swiss Federal Institute of Technology (ETH) Zurich (August 2015).
- [7] R. S. Sutton, A. G. Barto, R. J. Williams, Reinforcement learning is direct adaptive optimal control, IEEE Control Systems Magazine 12 (2) (1992) 19–22.
- [8] N. O. Lambert, D. S. Drew, J. Yaconelli, S. Levine, R. Calandra, K. S. J. Pister, Low-level control of a quadrotor with deep model-based reinforcement learning, IEEE Robotics and Automation Letters 4 (4) (2019) 4224–4230.
- [9] J. Yoo, D. Jang, H. J. Kim, K. H. Johansson, Hybrid reinforcement learning control for a micro quadrotor flight, IEEE Control Systems Letters 5 (2) (2020) 505–510.
- [10] M. Deisenroth, C. Rasmussen, PILCO: A model-based and data-efficient approach to policy search, in: ICML, 2011.
- [11] T. Lillicrap, J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, D. Wierstra, Continuous control with deep reinforcement learning, CoRR abs/1509.02971 (2016).
- [12] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, O. Klimov, Proximal policy optimization algorithms, CoRR abs/1707.06347 (2017).
- [13] M. Gaon, R. I. Brafman, Reinforcement learning with non-markovian rewards, in: The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7–12, 2020, AAAI Press, 2020, pp. 3980–3987.
- [14] A. Nilim, L. E. Ghaoui, Robust markov decision processes with uncertain transition matrices, Ph.D. thesis, USA, aAI3165509 (2004).
- [15] W. Koch, R. W. Renato Mancuso, A. Bestavros, Reinforcement Learning for UAV Attitude Control, ACM Trans. Cyber-Phys. Syst. 3, 2, Article 22 ((February 2019)).
- [16] A. Molchanov, T. Chen, W. Hönig, J. A. Preiss, N. Ayanian, G. S. Sukhatme, Sim-to-(multi)-real: Transfer of low-level robust control policies to multiple quadrotors, CoRR abs/1903.04628 (2019). arXiv:1903.04628.
URL <http://arxiv.org/abs/1903.04628>
- [17] F. Fei, Z. Tu, X. Deng, Learn-to-recover: Retrofitting uavs with reinforcement learning-assisted flight control under cyberphysical attacks, in: ICRA, 2020.
- [18] T. Koning, Developing a self-learning drone (april 2020).
- [19] L. Bjarre, Learning for quadcopter control (Dec. 2019).
- [20] M. Wen, R. Ehlers, U. Topcu, Correct-by-synthesis reinforcement learning with temporal logic constraints, in: 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2015, Hamburg, Germany, September 28 - October 2, 2015, IEEE, 2015, pp. 4983–4990. arXiv:1503.01793, doi:10.1109/IROS.2015.7354078.
URL <https://doi.org/10.1109/IROS.2015.7354078>
- [21] Q. Gao, D. Hajinezhad, Y. Zhang, Y. Kantaros, M. M. Zavlanos, Reduced variance deep reinforcement learning with temporal logic specifications, in: X. Liu, P. Tabuada, M. Pajic, L. Bushnell (Eds.), Proceedings of the 10th ACM/IEEE International Conference on Cyber-Physical Systems, ICCPS 2019, Montreal, QC, Canada, April 16–18, 2019, ACM, 2019, pp. 237–248. doi:10.1145/3302509.3311053.
URL <https://doi.org/10.1145/3302509.3311053>
- [22] M. Hasanbeig, Y. Kantaros, A. Abate, D. Kroening, G. J. Pappas, I. Lee, Reinforcement learning for temporal logic control synthesis with probabilistic satisfaction guarantees, in: 58th IEEE Conference on Decision and Control, CDC 2019, Nice, France, December 11–13, 2019, IEEE, 2019, pp. 5338–5343.

- doi:10.1109/CDC40024.2019.9028919.
 URL <https://doi.org/10.1109/CDC40024.2019.9028919>
- [23] M. Hasanbeig, D. Kroening, A. Abate, Towards verifiable and safe model-free reinforcement learning, in: N. Gigante, F. Mari, A. Orlandini (Eds.), Proceedings of the 1st Workshop on Artificial Intelligence and Formal Verification, Logic, Automata, and Synthesis, co-located with the 18th International Conference of the Italian Association for Artificial Intelligence, OVERLAY@AI*IA 2019, Rende, Italy, November 19-20, 2019, Vol. 2509 of CEUR Workshop Proceedings, CEUR-WS.org, 2019, p. 1.
 URL <http://ceur-ws.org/Vol-2509/invited.pdf>
- [24] M. Alshiekh, R. Bloem, R. Ehlers, B. Könighofer, S. Niekum, U. Topcu, Safe reinforcement learning via shielding, in: S. A. McIlraith, K. Q. Weinberger (Eds.), Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018, AAAI Press, 2018, pp. 2669–2678.
 URL <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/17211>
- [25] W. Zhang, O. Bastani, MAMPS: safe multi-agent reinforcement learning via model predictive shielding, CoRR abs/1910.12639 (2019). arXiv:1910.12639.
 URL <http://arxiv.org/abs/1910.12639>
- [26] O. Bastani, Safe reinforcement learning with nonlinear dynamics via model predictive shielding (2020). arXiv:1905.10691.
- [27] O. Bastani, Safe reinforcement learning via online shielding, CoRR abs/1905.10691 (2019). arXiv:1905.10691.
 URL <http://arxiv.org/abs/1905.10691>
- [28] G. E. Fainekos, G. J. Pappas, Robustness of temporal logic specifications for continuous-time signals, Theor. Comput. Sci. 410 (42) (2009) 4262–4291. doi:10.1016/j.tcs.2009.06.021.
 URL <https://doi.org/10.1016/j.tcs.2009.06.021>
- [29] A. Donzé, On signal temporal logic, in: A. Legay, S. Bensalem (Eds.), Runtime Verification - 4th International Conference, RV 2013, Rennes, France, September 24-27, 2013. Proceedings, Vol. 8174 of Lecture Notes in Computer Science, Springer, 2013, pp. 382–383. doi:10.1007/978-3-642-40787-1_27.
 URL https://doi.org/10.1007/978-3-642-40787-1_27
- [30] L. Brim, P. Dluhos, D. Safránek, T. Vejpustek, Stl: Extending signal temporal logic with signal-value freezing operator, Inf. Comput. 236 (2014) 52–67. doi:10.1016/j.ic.2014.01.012.
 URL <https://doi.org/10.1016/j.ic.2014.01.012>
- [31] T. Akazaki, I. Hasuo, Time robustness in MTL and expressivity in hybrid system falsification, in: D. Kroening, C. S. Pasareanu (Eds.), Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part II, Vol. 9207 of Lecture Notes in Computer Science, Springer, 2015, pp. 356–374. doi:10.1007/978-3-319-21668-3_21.
 URL https://doi.org/10.1007/978-3-319-21668-3_21
- [32] A. Bakhirkin, N. Basset, Specification and efficient monitoring beyond STL, in: T. Vojnar, L. Zhang (Eds.), Tools and Algorithms for the Construction and Analysis of Systems, Vol. 11428 of Lecture Notes in Computer Science, Springer, 2019, pp. 79–97.
- [33] H. Abbas, Y. V. Pant, R. Mangharam, Temporal logic robustness for general signal classes, in: N. Ozay, P. Prabhakar (Eds.), Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control, HSCC 2019, Montreal, QC, Canada, April 16-18, 2019, ACM, 2019, pp. 45–56. doi:10.1145/3302504.3311817.
 URL <https://doi.org/10.1145/3302504.3311817>
- [34] I. Haghighi, N. Mehdipour, E. Bartocci, C. Belta, Control from signal temporal logic specifications with smooth cumulative quantitative semantics, in: 58th IEEE Conference on Decision and Control, CDC 2019, Nice, France, December 11-13, 2019, IEEE, 2019, pp. 4361–4366. doi:10.1109/CDC40024.2019.9029429.
 URL <https://doi.org/10.1109/CDC40024.2019.9029429>
- [35] N. Mehdipour, C. I. Vasile, C. Belta, Arithmetic-geometric mean robustness for control from signal temporal logic specifications, in: 2019 American Control Conference, ACC 2019, Philadelphia, PA, USA, July 10-12, 2019, IEEE, 2019, pp. 1690–1695.
 URL <http://ieeexplore.ieee.org/document/8814487>
- [36] Y. Gilpin, V. Kurtz, H. Lin, A smooth robustness measure of signal temporal logic for symbolic control, IEEE Control. Syst. Lett. 5 (1) (2021) 241–246. doi:10.1109/LCSYS.2020.3001875.
 URL <https://doi.org/10.1109/LCSYS.2020.3001875>
- [37] Z. Zhang, I. Hasuo, P. Arcaini, Multi-armed bandits for boolean connectives in hybrid system falsification, in: I. Dillig, S. Tasiran (Eds.), Computer Aided Verification - 31st International Conference, CAV 2019, New York City, NY, USA, July 15-18, 2019, Proceedings, Part I, Vol. 11561 of Lecture Notes in Computer Science, Springer, 2019, pp. 401–420. doi:10.1007/978-3-030-25540-4_23.
 URL https://doi.org/10.1007/978-3-030-25540-4_23
- [38] D. Aksaray, A. Jones, Z. Kong, M. Schwager, C. Belta, Q-learning for robust satisfaction of signal temporal logic specifications, in: 55th IEEE Conference on Decision and Control, CDC 2016, Las Vegas, NV, USA, December 12-14, 2016, IEEE, 2016, pp. 6565–6570. doi:10.1109/CDC.2016.7799279.
 URL <https://doi.org/10.1109/CDC.2016.7799279>
- [39] X. Li, C. Belta, Temporal logic guided safe reinforcement learning using control barrier functions (2019). arXiv:1903.09885.
- [40] L. Lindemann, D. V. Dimarogonas, Control barrier functions for signal temporal logic tasks, IEEE Control. Syst. Lett. 3 (1) (2019) 96–101. doi:10.1109/LCSYS.2018.2853182.
 URL <https://doi.org/10.1109/LCSYS.2018.2853182>
- [41] E. Goubault, S. Putot, Inner and Outer Reachability for the Verification of Control Systems, HSCC (April 2019).
- [42] C. Luis, J. Le Ny, Design of a Trajectory Tracking Controller for a Nanoquadcopter, Tech. rep., Mobile Robotics and Autonomous Systems Laboratory, Polytechnique Montreal (August 2016).
- [43] Bitcraze, <https://store.bitcraze.io/>.
- [44] M. Bangura, R. Mahony, Nonlinear dynamic modeling for high performance control of a quadrotor, in: Australasian Conference on Robotics and Automation, 2012.
- [45] C. Poussot-Vassal, F. Demourant, A. Lepage, D. Le Bihan, Gust load alleviation: Identification, control, and wind tunnel testing of a 2-d aeroelastic airfoil, IEEE Transactions on Control Systems Technology 25 (5) (2017) 1736–1749. doi:10.1109/TCST.2016.2630505.
- [46] R. Arora, A. Basu, P. Mianjy, A. Mukherjee, Understanding deep neural networks with rectified linear units, CoRR abs/1611.01491 (2016).
 URL <http://arxiv.org/abs/1611.01491>
- [47] A. Bemporad, M. Morari, V. Dua, E. N. Pistikopoulos, The explicit linear quadratic regulator for constrained systems, Automatica 38 (1) (2002) 3–20.
- [48] J. Ferlez, X. Sun, Y. Shoukry, Two-level lattice neural network architectures for control of nonlinear systems, CoRR abs/2004.09628 (2020).
 URL <https://arxiv.org/abs/2004.09628>
- [49] J. Ferlez, Y. Shoukry, Aren: assured relu NN architecture for model predictive control of LTI systems, in: HSCC, ACM, 2020, pp. 6:1–6:11.
- [50] S. Lucia, B. Karg, A deep learning-based approach to robust nonlinear model predictive control, IFAC-PapersOnLine 51 (20) (2018) 511 – 516, 6th IFAC Conference on Nonlinear Model Predictive Control NMPC 2018.
- [51] A. Donze, Monitoring temporal properties of continuous signals, International Conference on Computer Aided Verification (2010) 167–170 https://link.springer.com/chapter/10.1007/978-3-642-14295-6_17.
- [52] A. Bakhirkin, N. Basset, Specification and efficient monitoring beyond stl, in: T. Vojnar, L. Zhang (Eds.), Tools and Algo-

- rithms for the Construction and Analysis of Systems, Springer International Publishing, Cham, 2019, pp. 79–97.
- [53] T. Akazaki, I. Hasuo, Time robustness in mtl and expressivity in hybrid system falsification, International Conference on Computer Aided Verification (2015) 356–374 <https://arxiv.org/pdf/1505.06307.pdf>.
 - [54] A. Hill, A. Raffin, M. Ernestus, A. Gleave, R. Traore, P. Dhariwal, C. Hesse, O. Klimov, A. Nichol, M. Plappert, A. Radford, J. Schulman, S. Sidor, Y. Wu, Stable Baselines, <https://github.com/hill-a/stable-baselines> (2018).
 - [55] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, et al., Tensorflow: A system for large-scale machine learning, in: 12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16), 2016, pp. 265–283.
 - [56] Bazel Documentation, <https://docs.bazel.build/>.
 - [57] Kubernetes Documentation, <https://kubernetes.io/docs/reference/>.
 - [58] Bazel Container Image Rules, https://github.com/bazelbuild/rules_docker.
 - [59] Bazel Kubernetes Rules, https://github.com/bazelbuild/rules_k8s.
 - [60] D. Haziza, J. Rapin, G. Synnaeve, Hiplot, interactive high-dimensionality plots, <https://github.com/facebookresearch/hiplot> (2020).
 - [61] S. Sinha, H. Bharadhwaj, A. Srinivas, A. Garg, D2rl: Deep dense architectures in reinforcement learning (2020). arXiv:2010.09163.
 - [62] C. Wang, Y. Wu, Q. Vuong, K. Ross, Striving for simplicity and performance in off-policy DRL: Output normalization and non-uniform sampling, in: Proceedings of the 37th International Conference on Machine Learning, Vol. 119 of Proceedings of Machine Learning Research, 2020, pp. 10070–10080.
 - [63] N. Rao, E. Aljalbout, A. Sauer, S. Haddadin, How to make deep rl work in practice (2020). arXiv:2010.13083.
 - [64] H. U. Sheikh, L. Bölöni, Reducing overestimation bias by increasing representation dissimilarity in ensemble based deep q-learning (2020). arXiv:2006.13823.
 - [65] D. Reda, T. Tao, M. van de Panne, Learning to locomote: Understanding how environment design matters for deep reinforcement learning, in: Motion, Interaction and Games, MIG, Association for Computing Machinery, 2020.
 - [66] S. Dutta, X. Chen, S. Sankaranarayanan, Reachability analysis for neural feedback systems using regressive polynomial rule inference, in: 22nd ACM International Conference on Hybrid Systems: Computation and Control, HSCC '19, ACM, New York, NY, USA, Montreal, QC, Canada, April 16-18, 2019, <https://doi.org/10.1145/33025043313351>.