

Learning-based vs Model-free Adaptive Control of a MAV under Wind Gust^{*}

Thomas Chaffre^{1,2}, Julien Moras³, Adrien Chan-Hon-Tong³, Julien Marzat³,
Karl Sammut², Gilles Le Chenadec¹, and Benoit Clement^{1,2}

¹ Lab-STICC UMR CNRS 6285, ENSTA Bretagne, Brest, France
{thomas.chaffre,gilles.le.chenadec}@ensta-bretagne.org

² Centre for Maritime Engineering, Flinders University, Adelaide, SA, Australia
{karl.sammut,benoit.clement}@flinders.edu.au

³ DTIS, ONERA, Université Paris-Saclay, 91123 Palaiseau, France
{julien.moras,adrien.chan_hon_tong,julien.marzat}@onera.fr

Abstract. Navigation problems under unknown varying conditions are among the most important and well-studied problems in the control field. Classic model-based adaptive control methods can be applied only when a convenient model of the plant or environment is provided. Recent model-free adaptive control methods aim at removing this dependency by learning the physical characteristics of the plant and/or process directly from sensor feedback. Although there have been prior attempts at improving these techniques, it remains an open question as to whether it is possible to cope with real-world uncertainties in a control system that is fully based on either paradigm. We propose a conceptually simple learning-based approach composed of a full state feedback controller, tuned robustly by a deep reinforcement learning framework based on the Soft Actor-Critic algorithm. We compare it, in realistic simulations, to a model-free controller that uses the same deep reinforcement learning framework for the control of a micro aerial vehicle under wind gust. The results indicate the great potential of learning-based adaptive control methods in modern dynamical systems.

1 Introduction

Aerial vehicles are exposed to a mixture of perturbations that fluctuate vigorously because of the hazardous environment they are evolving in. They operate over a wide range of speeds and altitudes. Their dynamics are nonlinear and can also be of time-varying nature (as they fly, their mass slowly decreases due to fuel consumption and their center of gravity can greatly vary). Micro Aerial Vehicles (MAVs) are challenging in some other ways. They have to constantly compensate for the distribution of forces that act on their small airframes because of

^{*} *This work was supported by ISblue project, Interdisciplinary graduate school for the blue planet (ANR-17-EURE-0015) and co-funded by a grant from the French government under the program "Investissements d'Avenir".*

wind and turbulence. These forces can significantly increase when weather conditions change and are also influenced by the velocity and heading of the vehicle. Because they are small, rigid, light, and move slowly, MAVs are highly sensitive to wind gusts even if mild. In addition to the latter unpredictable (external) forces, we face the difficulty of accurately modeling their (internal) dynamics under these conditions. With the development of flight controllers for hypersonic aircraft by NASA in the early 1950's [1], it was found that conventional linear feedback control strategies (e.g. fixed-gain control [2], robust control [3], sliding mode control [4] or fuzzy logic control [5]) are too limited to handle such entire regimes. It is also true for the MAVs context where for a variety of reasons it is almost impossible to manually re-tune the control parameters, these include:

- The uncertainty level of the wind speed is high and can vary more quickly than can be compensated by these types of controllers.
- The response at some operating points has to be overly conservative in order to satisfy specifications at other operating points.
- The controlled process itself varies significantly during operations (the motors' and propellers' efficiency for example).

The field of adaptive control broadly addresses these challenges by granting the control law some flexibility to modify its action based on the process variations. Despite the great advancements conducted in the theoretical area, the expansion of autonomous agents to the real-world is still limited. This is mostly due to the high dimensional problems that often appear when working under real conditions where it usually requires many degrees of freedom to describe the robot state. The lack of basic knowledge of the various natural phenomena taking place in the robot's environment, e.g. wind gust for MAVs or sea current for autonomous underwater vehicles (AUVs) makes it even more difficult to actively control it [6] with model-based adaptive methods.

For the past few years, model-free adaptive control methods have attracted significant attention. They exploit strong statistical tools that provide control systems the ability to automatically learn and improve from experience without being explicitly told how to. In particular, model-free adaptive methods based on deep reinforcement learning (DRL) have shown promising achievements. This was possible especially thanks to the use of deep neural networks to extract physical insights through sensory data (empowered by progress in data collection with high fidelity simulations, cheaper storage drives, faster computers, etc). However, as stated in [7], the extension of DRL techniques for robotic tasks raises serious questions. Compared to other application fields, robots have to interact with a dynamic environment where relevant information about the system is not always accessible or tractable over time.

In their up-to-date thorough investigations [8], G. Dulac-Arnold *et al.* presented real-world reinforcement learning challenges that are still not resolved, these include *Satisfying Environmental Constraints*, *High-Dimensional Continuous State and Action Spaces* or *Multi-Objective Reward Functions*. In addition, the notions of stability (in terms of Lyapunov stability) in the DRL framework is more deeply investigated [9] but the lack of formalism is still often highlighted

by the control community as a significant concern. The purpose of this study is to explore to what extent model-free control techniques combined with the classical theory of dynamical systems can be, at least, part of the solution to these great challenges. This is the objective of the learning-based control area [10] which aims to combine the advantages of the aforementioned paradigms (model-based and model-free) into one hybrid control scheme. Therefore, we propose a learning-based adaptive control system composed of a state-feedback control law whose parameters (i.e. gains) are tuned by the Soft Actor-Critic DRL algorithm [11] using the pole placement method. It is fairly compared, in a realistic simulation setting, with a model-free adaptive method trained end-to-end with the same DRL framework (as initially applied to robot navigation in [12]). The control objective treated is a hexacopter performing a waypoint rallying mission under unknown wind fields.

The paper is organized as follows. Section 2 gives details on the theoretical bases of the proposed control strategies and related work. In Section 3, we present background material on the simulation and control of MAVs. Section 4 is dedicated to the description of our contribution, with the application of DRL to the design of model-free and learning-based controllers. A description of the simulated training and evaluation sessions is provided in Section 5 with the obtained results. Finally, a thorough analysis of the outcomes and some suggestions to extend the work further are presented in Section 6.

2 Related work

An adaptive controller is essentially a controller that can adjust its own behavior in response to changes in the dynamics of the process and the character of the disturbances [13]. Since ordinary feedback control also seeks to reduce these undesirable effects, one needs to clearly distinguish between the two. In this paper, we will designate as *adaptive*, a control system that provides the desired system performance by changing its parameters and/or structure in order to reduce the uncertainties effect and to improve the knowledge of the desired system. As mentioned by Nikolai M. Filatov [14], the goal of adaptive systems is to switch off the adaptation as soon as the system uncertainty is reduced sufficiently and to then use the adjusted controller with a fixed structure and fixed parameters. Although it is more common in the control community to classify methods based on the nature of the model (e.g. linear vs nonlinear, continuous vs discrete), we classify here adaptive controllers based on their dependence on the mathematical model of the controlled system. This algebraic representation of the dynamics of the system can take several forms such as differential equations, state-space representations, or frequency domain representations. We can thus identify three types of adaptive control methods: totally model-based, totally model-free, and learning-based. Since the approach presented here consists in merging together some model-based controllers and model-free algorithms, in this chapter we will present the basis of the aforesaid techniques and recent applications to the problem of MAVs under unknown wind fields. Then, we will

see how reinforcement learning can be used to strengthen control laws, which will lead to the introduction of the proposed learning-based adaptive controller.

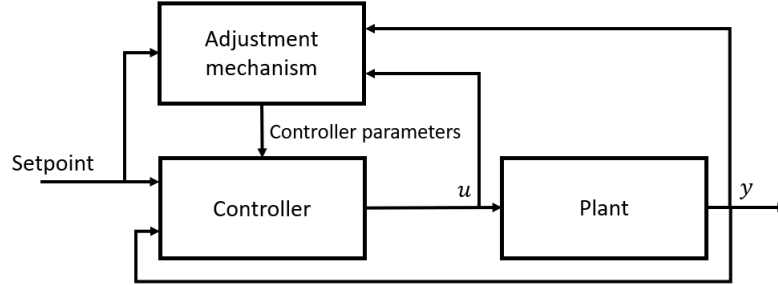


Fig. 1: General block diagram of an adaptive control system. It is composed of two loops: a normal feedback loop with the plant and the controller and a second loop with the parameter adjustment mechanism (which is often slower than the first one).

Formulation of the adaptive control problem:

Consider the system described by the following continuous-time state equations:

$$\begin{aligned} \dot{x}(t) &= f(x(t), u(t), p(t), t), \\ y(t) &= h(x(t), u(t), t), \end{aligned} \quad (1)$$

where $x \in \mathbb{R}^n$ is the system state²; $u \in \mathbb{R}^{n_u}$ is the control vector; $p(t) \in \mathbb{R}^p$ is the constant or time-varying vector composed of the unknown or uncertain parameters of the model; $y \in \mathbb{R}^m$ is the output vector of interest. The probability density of the initial values $\mathbb{P}[x(0), p(0)]$ is assumed to be known. We can define the set of outputs and control inputs available at time t as:

$$\Gamma_t = \{y(t), \dots, y(0), u(t-1), \dots, u(0)\}, \quad t = 0, 1, \dots, T-1, \quad \Gamma_0 = \{y(0)\} \quad (2)$$

The control performance index can have the following form:

$$J(t) = \frac{1}{t} \int_0^t \bar{e}^2(\tau) d\tau, \quad \bar{e}(t) = w(t) - y(t) \quad (3)$$

where $w(t)$ is the targeted set point. The general problem of adaptive control consists in finding the control policy $u(t) = u_t(\Gamma_t) \in \bar{\Omega}_t$ that minimizes the performance index (3) for the system described by (1) with $\bar{\Omega}_t$ the domain in the space \mathbb{R}^{n_u} where the admissible control values are defined. Backward recursion

² In this paper, x refers to state in terms of state space representation and s in regard to the context of machine learning because although sharing the same title, these entities do not hold the same nature.

of the following stochastic continuous-time dynamic programming equations can give the optimal control of the above problem:

$$J_{t-1}(\Gamma_{t-1}) = \min_{u(t-1) \in \Omega_{t-1}} \left[\frac{1}{t} \int_0^t \bar{e}^2(\tau) d\tau | \Gamma_{t-1} \right], \quad (4)$$

$$J_t(\Gamma_t) = \min_{u(t) \in \Omega_t} \left[\frac{1}{t} \int_0^t \bar{e}^2(\tau) d\tau + J_{t+1}(\Gamma_{t+1}) | \Gamma_t \right], \quad (5)$$

for $t = t-2, t-3, \dots, 0$.

This optimal solution is usually impractical to derive analytically from (4) and (5) because of the dimension of the underlying spaces, even for simple cases. Near-optimal solutions can be obtained with model-based adaptive methods.

2.1 Model-based adaptive control

The design of a controller usually begins with a mathematical model of the system to be controlled. For mechanical systems of moderate dimension, it is possible to write down such a model (e.g. based on the Newtonian, Lagrangian, or Hamiltonian formalism) and eventually linearize its dynamics around a fixed point or periodic orbit. Model-based adaptive control refers to adaptive control methods that are completely reliant on this type of model. This means that the feedback control law will adjust the controller parameters online to compensate for model uncertainty. This adaptation can be performed in different manners, which gives the classification of model-based methods in two groups: direct and indirect. In direct schemes, designers attempt to estimate the control parameters. The adjustment rule tells directly how the controller's parameters should be updated. It is done without intermediate calculations involving plant parameter estimates. This is possible because in many cases, there exist measurable variables that correlate well with changes in the process dynamics. The *Gain Scheduling* [15] and *Model-Reference Adaptive Systems* [16] algorithms are examples of direct schemes. On the other hand, if the estimates of the process parameters are updated and the controller parameters are obtained from the solution of a design problem using the estimated parameters, we obtain what is called an indirect approach. This kind of controller can be seen as automation of process modeling and design, in which the process model and the control design are updated at each sampling period. The *Self-Tuning Regulators* [17] and *Dual Controllers* [18] are well-known indirect methods. Instead of further extending the theory basis of model-based control (since it is not the main focus of this study), we describe here recent applications to the control of quadrotor MAVs under wind disturbances. An output controller was proposed in [19] to cancel wind disturbances where the nonlinear dynamical mathematical model of a quadcopter was decomposed into two parts: a static MIMO transformation and few SISO channels. It was assumed that the unknown wind force acts on each channel of the quadcopter as a constant signal that has to be canceled. They proposed to first design a virtual control law for each SISO channel and

then to inverse the MIMO transformation to get the control laws for the initial considered system. Their approach showed, under simulation, that it was able to efficiently reach steady-state under unknown wind perturbations. Later on, to reduce the undesirable effects of wind on a similar quadcopter, A. Razinkova et al. [20] chose another strategy. They assumed the wind to be an unknown but uniform and time-varying force acting on the vehicle body in the X and Y -axes. Therefore, they proposed to augment a PD controller with additional terms adapting to the aforesaid forces. The resulting PD control law according to an axis i takes the form: $u_i = K_p \times e(t) + K_d \times e(t)/d_t - \gamma \tilde{i}$, where $e(t)$ is the position error and $\gamma \tilde{i}$ the adaption term. Such an adaptation converges [20] if and only if the adaptation rate γ remains positive. Later on, in [21] a \mathcal{L}_1 adaptive controller [22,23] (which is an advanced version of MRAC [16]) was proposed to control a quadcopter that is performing wind turbine inspection. They proved the robustness of their controller with respect to wind under simulated and real flights by comparing it to a basic Linear Quadratic Regulator (LQR). More recently, in [24] the gains of a PID controller have been tuned especially in order to handle wind gusts. They tuned this classical feedback controller in the \mathcal{H}_2 optimal control framework and compared it with the existing LQR-based tuning method [25]. Simulated experiments proved that the classical PID controller combined with a dedicated parameter adjustment mechanism is better at rejecting wind disturbances than the classical LQR method.

These approaches have in common the use of the *Certainty Equivalence* (CE) approach, which means that the estimation uncertainty is not taken into consideration. The wind estimates are used in the control law as if they were the real values of the unknown wind field. The CE paradigm has been used for a long time and makes it possible to generate a wide class of model-based adaptive controllers by combining different on-line parameter estimation algorithms (e.g. sensitivity methods, positivity and Lyapunov design, gradient and least-squares methods, etc) with different control laws (e.g. PID, LQG, LQR, Pole placement, etc). However, it was proved in [14] that control systems based on the CE approach are not always optimal and can be far from so. A straightforward improvement would be to apply reinforcement on the uncertainty estimates over time, in a sort of *learning process*, as a mean to optimize control performance. This goal, combined with the development of complex systems, where precise modelization is extremely difficult, has motivated the development of the model-free adaptive control field.

2.2 Model-free adaptive control

The *model-free* wording refers here to the fact that these families of controllers do not rely on any mathematical model of the controlled system. They aim at describing complex systems from observational data collected directly by embedded sensors rather than first-principles modeling. As shown in Figure 2, model-free algorithms can be seen as an optimization problem where the goal is to minimize a cost function without closed-form knowledge of the function or its gradient. A classic approach widely used in the model-free control framework

is the *Extremum-Seeking* (ES) methods [26,27]. We present now a simple ES algorithm by considering the same equations of state (1). The goal of the control is again to optimize the performance of the system described by the cost function (3). We model it as a smooth function $J(x, u) : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}$. We will denote it simply by $J(u)$ because the state vector x is driven by u . In order to write convergence results we need the following assumptions:

1. There exists a smooth function $l : \mathbb{R} \rightarrow \mathbb{R}^n$ such that:

$$f(x(t), u(t), p(t), t) = 0, \text{ if and only if } x(t) = l(u(t)) \quad (6)$$

2. For each $u \in \mathbb{R}$, the equilibrium $x = l(u)$ is locally exponentially stable.
3. There exists (a maximum) $u^* \in \mathbb{R}$ such that:

$$\begin{aligned} (J \circ l)^{(1)}(u^*) &= 0 \\ (J \circ l)^{(2)}(u^*) &< 0 \end{aligned} \quad (7)$$

Based on these assumptions, one can easily design a simple extremum seeker with proven convergence bounds. One of the simplest way to maximize J is to use a gradient-based ES control law as:

$$\dot{u} = k \frac{dJ}{du}, \quad k > 0 \quad (8)$$

The convergence of (8) can be analysed with the Lyapunov function:

$$V = J(u^*) - J(u) > 0, \quad \text{for } u \neq u^* \quad (9)$$

The derivative of V gives:

$$\dot{V} = \frac{dJ}{du} \dot{u} = -k \left(\frac{dJ}{du} \right)^2 \leq 0 \quad (10)$$

This proves that the algorithm drives u to the invariant set $\frac{dJ}{du} = 0$, which is equivalent to $u = u^*$. However, as simple as it seems, this approach requires the knowledge of the gradient of J . An ES-based model-free approach was proposed in [28] for the optimal control of a quadcopter carrying different payloads. They proved, with experimental flights, that their model-free controller is able to find the speed that maximizes the flight time (endurance) or flight distance (range) of the vehicle when transporting an unknown payload. They designed specific cost functions for each of these goals. Then, the ES algorithm was applied to estimate an optimal velocity which is used to transform the desired path into a trajectory. They proved that the ES scheme is able to find unknown, time-varying, operating points that minimize these cost functions directly from sensorial measurements and without any model of the MAV power consumption.

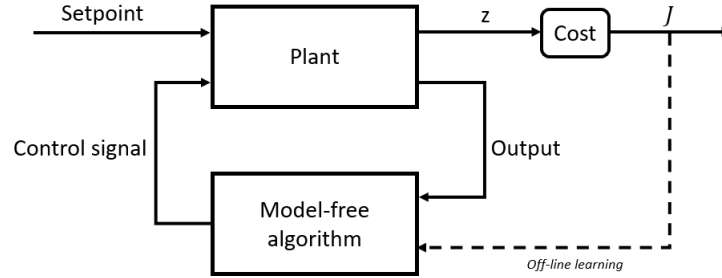


Fig. 2: General block diagram of a classical model-free controller. The control objective is to minimize a well-defined cost function J within the space of possible control laws. The off-line learning loop provides experiential data to train the controller. The vector z is composed of all the information that may factor into the cost.

Model-free control is a blossoming field where a handful of strong techniques are currently being developed and applied to minimize or maximize specific cost functions by using the system outputs as shown in Figure 2. Reinforcement learning (RL) is currently the leading technique used in this research area. It is a class of machine learning methods in which an autonomous agent in an environment has to learn how to take actions in order to maximize the notion of cumulative reward [29]. More formally, the objective of RL is for a learning agent to find an optimal strategy behavior (called a policy and denoted by π) from experimental trials and errors. Those repeated interactions are under the form of the execution of an action $a_t \in A$ from state $s_t \in S$ which makes the agent transit to a new state $s_{t+1} \in S$. This transition produces a reward signal denoted by $r(s_t, a_t) \in \mathbb{R}$, which quantitatively transcribes how well the agent is doing in the environment. The policy is a rule used by the agent to decide what actions to take and is mathematically defined as a function:

$$\pi : S \rightarrow A \quad (11)$$

This type of procedure may be framed as a Markov Decision Process [30] along a trajectory (a sequence of T actions and $T + 1$ states in the environment) $\tau = (s_0, a_0, s_1, a_1, \dots, a_{T-1}, s_T)$. Hence, the probability of a trajectory for a π policy is:

$$P(\tau|\pi) = \rho_0(s_0) \prod_{t=0}^{T-1} P(s_{t+1}|s_t, a_t)\pi(a_t|s_t) \quad (12)$$

where $\rho_0(s_0)$ means that the starting state s_0 is randomly sampled from the distribution ρ_0 and $\pi(a_t|s_t)$ means that the action a_t is randomly sampled by the policy π in the state s_t . The overall optimization problem considered in RL is to determine the policy π^* which maximizes the expected return when the agent decides to take actions according to this policy:

$$\pi_* = \arg \max J_{RL}(\pi) \quad (13)$$

where $J_{RL}(\pi)$ denotes the expected return

$$J_{RL}(\pi) = \int_{\tau} \sum_{t=0}^{T-1} r_t \cdot P(\tau|\pi) = \underset{\tau \sim \pi}{E} \left[\sum_{t=0}^{T-1} r_t \right] \quad (14)$$

Both RL and model-free paradigms are based on the same optimization formalism. We can easily see the similarity between this cost function (14) and the one from the ES algorithm (8). Note that the terms model-based and model-free can also be found in the RL theory but do not correspond to the ones used in the control field. In RL, they refer to whether or not a model of the environment (a function which predicts state transitions and rewards) is given beforehand to the RL algorithm to improve the policy whereas, in control, it is the model of the controlled dynamical system that might be provided.

Various methods exist to tackle the RL optimization problem. Among them, *Policy Gradient* techniques have enabled to use RL in real-world robotic contexts (see [31] for an extensive definition of these approaches and [32,33,34] for successful real-world robotic applications). These techniques are based on the use of a stochastic policy formally denoted by π_{θ} , where θ is a vector of parameters. The vector θ is estimated iteratively by gradient ascent. The expression of the gradient of $J(\pi_{\theta})$ with respect to θ is given by:

$$\nabla_{\theta} J_{RL}(\pi_{\theta}) = \underset{\tau \sim \pi_{\theta}}{E} \left[\sum_{t=0}^T \nabla_{\theta} \log \pi_{\theta}(a_t|s_t) (Q^{\pi_{\theta}}(s_t, a_t) - V^{\pi_{\theta}}(s_t)) \right] \quad (15)$$

where $V^{\pi_{\theta}}(s_t)$ is V -value function quantifying the expected return if the trajectory starts in a state s_t taking actions upon the policy π_{θ} :

$$V^{\pi_{\theta}}(s_t) = \underset{\tau \sim \pi_{\theta}}{E} \left[\sum_{t=0}^T r_t | s_0 = s_t \right] \quad (16)$$

and $Q^{\pi_{\theta}}(s_t, a_t)$ is the Q -value function quantifying the expected return if the trajectory starts in a state s_t , takes an action a_t and then takes actions upon the policy π_{θ} :

$$Q^{\pi_{\theta}}(s_t, a_t) = \underset{\tau \sim \pi_{\theta}}{E} \left[\sum_{t=0}^T r_t | (s_0 = s_t, a_0 = a_t) \right] \quad (17)$$

2.3 Learning-based adaptive control

Learning-based controllers can be seen as hybrid control schemes. They are used when we only have access to an imperfect physics-based model of the system. The idea is to strengthen it with some model-free algorithm to compensate for the uncertain or the missing parts of the model. This compensation can either be done directly by *learning* the uncertain parts or indirectly by *tuning* the

controller parameters to cope with the uncertainty. To illustrate this idea, we consider the system (1) with a specific structure that can be written as:

$$\begin{aligned}\dot{x}(t) &= f_1(x(t), u(t), t) + f_2(x(t), p(t), t), \\ y(t) &= h(x(t), u(t), t)\end{aligned}\tag{18}$$

In this model, f_1 represents the known part of the plant that can be efficiently driven by a classical model-based control architecture, while f_2 represents the unknown part of the model that can be compensated by some model-free learning algorithm. The goal of learning-based controllers is to use a model-free step to optimize an unknown performance function and then use a model-based control law to guide the system's dynamics toward the optimal performance. The idea of combining both techniques is indeed attractive. One designer could take advantage of the model-based design, with its stability characteristics, and add to it the advantages of model-free learning, with its fast convergence and robustness to uncertainties. Overall, in Neural-Network (NN) learning-based control design, the idea is to write the model of the plant as a combination of a known and unknown part (i.e. the disturbances). The NNs are then used to estimate the unknown part of the model. As a result, a controller based on the known part (model-based) and the NN estimates of the uncertainties (model-free) is determined to realize some desired regulation or tracking performance. To illustrate this concept we formalize it below, in a similar spirit as the formalism from [10] where a state-space model of (1) under the Brunovsky form is considered as:

$$\begin{cases} \dot{x}_1 = x_2, \\ \dot{x}_2 = f(\cdot) + u, \end{cases}\tag{19}$$

where $f(\cdot) = f_1(x(t), u(t), t) + f_2(x(t), p(t), t)$, with f_1 known and f_2 an unknown smooth function of the state variables $x = (x_1, x_2)^T$ and u the control signal. The unknown part of the model, namely f_2 , is estimated by a NN as:

$$\hat{f}_2 = \hat{W}^T S(x(t)),\tag{20}$$

where $\hat{W} = (\hat{w}_1, \dots, \hat{w}_N)^T \in \mathbb{R}^N$ is the estimated vector of synapse weights of the neural network node and $S(x) = (s_1(x), \dots, s_n(x))^T$ is the regressor vector, with $s_i, i = 1, \dots, N$. Consider the reference model:

$$\begin{cases} \dot{x}_{ref1} = x_{ref2}, \\ \dot{x}_{ref2} = f_{ref}(x), \end{cases}\tag{21}$$

where f_{ref} is a known nonlinear smooth function of the desired trajectories $x_{ref} = (x_{ref1}, x_{ref2})^T$. A basic learning-based controller can be defined as:

$$u = -e_1 - c_1 e_2 - \hat{W}^T S(e) + \dot{v},\tag{22}$$

with:

$$\begin{aligned}
e_1 &= x_1 - x_{ref1}, \\
e_2 &= x_2 - v, \\
v &= -c_2 e_1 + x_{ref2} \\
\dot{v} &= -c_2(-c_2 e_1 + e_2) + f_{ref}(x_{ref}), \quad (c_1, c_2 > 0) \\
\dot{W} &= \Gamma(S(e)e_2 - \sigma \hat{W}), \quad (\sigma > 0 \text{ and } \Gamma^T > 0)
\end{aligned} \tag{23}$$

It can be observed that u is now a function of the Brunovsky form (19) (i.e., model-based information), and the remaining part is based on the NN estimates of f_2 (model-free estimation). The unknown part is here estimated by a neural network but other strategies can be used in the learning-based scheme.

This concept of using NNs to estimate the unknown parameters (i.e. the wind field for our use case) seems appealing. In fact, as stated by the universal approximation theorem [35] for any function $f(x)$, regardless of its complexity, there exists a neural network such that for every possible input x , the value $f(x)$ is a feasible output from this network. If the external disturbances are well estimated, efficient control laws can be designed accordingly. This idea was investigated in [36,37] where a *Geometric Adaptive Controller* was proposed based on NNs for a quadcopter in wind fields. Two control laws were defined with adaptive control terms denoted as \bar{V}_i to mitigate the effects of the unknown disturbance. These adaptive control terms were computed with neural networks, exactly like in (22), as: $\bar{V}_i = \bar{W}_i^T \zeta(\bar{z}_i)$ with $\bar{z}_i = \bar{V}_i^T x_{nn_i}$ (\bar{V}_i being the current estimate of the ideal weighting parameters and x_{nn_i} the inputs of the neural network). They proved with real experiments that the learning-based controller succeeded to complete the considered backflip maneuver followed by a stable hovering flight whereas it was not possible without the disturbance rejection terms. Another strategy was proposed in [38], where a learning-based safety-preserving cascaded QP controller (SPQC) using *Gaussian Processes* (GP) [39] has been proposed for safe trajectory tracking by a quadcopter in a cluttered environment. More precisely, the cascade controller is composed of two QP controllers: a position and an attitude level QP controller. The first one generates the desired thrust while the second makes use of it, together with the high confidence uncertainty interval obtained via GPs, to compute the desired body rotational rates. They evaluated their approach under numerical simulation whose results proved that the proposed learning-based controller is able to perform a trajectory tracking task with obstacle avoidance capacity under changing wind fields.

Compared to the approaches presented in this related work Section, we propose to only treat the parameters adjustment task. We design a learning-based controller that uses a DRL-based model-free algorithm to perform online-tuning of the parameters of a state-feedback controller. We compare this method to the model-free strategy initially applied in [12] but this time for the application of a waypoint rallying mission by a MAV under unknown wind gusts.

3 MAV simulation, modeling and control

This Section presents the ROS package used to simulate the aerial vehicle and wind perturbations and the corresponding model of the Firefly platform, which is the hexacopter considered in this study. The principal elements necessary to understand the control designs proposed in Section 4 are also derived.

3.1 RotorS package

Testing algorithms on physical platforms can be very time-consuming and dangerous for the robot. It can be even riskier when one designer wants to perform the training of machine learning algorithms directly on the real robot (because of their initial hazardous behavior). There exist many robotic simulation tools to reduce field testing time and make debugging easier. In this work, we relied on the ubiquitous Gazebo simulator [40] which is connected to the Robot Operating System framework (ROS [41]). ROS is an open-source meta-operating system for robots that facilitates the reuse of code.

The Gazebo-based package called RotorS [42] is a high-fidelity simulation framework for MAVs (developed by the Autonomous Systems Lab team from ETH Zurich) that does not require any additional components to simulate high-level tasks (e.g. path-planning, collision avoidance, or vision-based problems). A complete model of the Firefly MAV is directly included in the package along with various world models and a plugin to simulate wind fields. We relied on the wind plugin provided by RotorS to generate wind fields in the environment. This plugin allows to define the wind as a 3D field sampled over a regular grid and each point specifies a wind velocity vector (in $m.s^{-1}$). More specifically, we employed the environment named *hemicyl* (see Figure 3) and its pre-configured wind field, the field used is composed of 6282 vertices and the wind velocity vary here between $-5m.s^{-1}$ and $+10m.s^{-1}$.

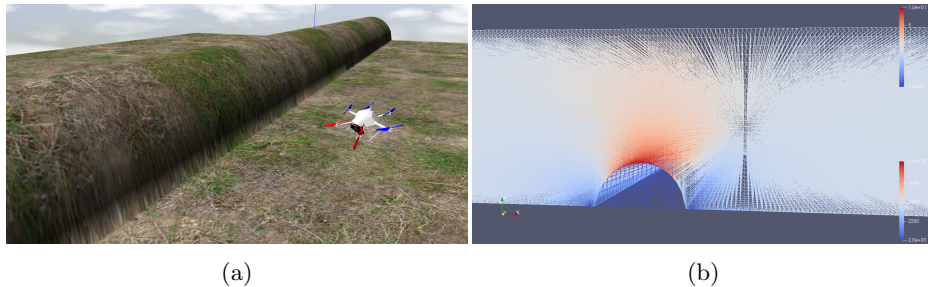


Fig. 3: Visualization of the simulated environment in Gazebo (a) and the wind field in Paraview (b). The complete details of the wind field we used can be find on the RotorS-based [extended wind plugin](#) page.

3.2 MAV model

The parametrization of the Firefly platform considers two reference frames: the world-fixed inertial frame \mathcal{R}_W and the body reference frame \mathcal{R}_B which is attached to the center of mass of the hexacopter. Coordinates in the world frame are denoted as $[x_W, y_W, z_W]^T$ while they are denoted as $[x_B, y_B, z_B]^T$ in the body frame. The pose of the hexacopter is given by its position $\zeta = [x_W, y_W, z_W]^T$ and orientation $\eta = [\phi, \theta, \psi]^T$ in the three Euler angles (respectively roll, pitch and yaw). For the sake of clarity, $\sin(\cdot)$ and $\cos(\cdot)$ are abbreviated as $s\cdot$ and $c\cdot$ in the next equation. The transformation from the world frame \mathcal{R}_W to the body frame \mathcal{R}_B is given by:

$$\begin{bmatrix} x_B \\ y_B \\ z_B \end{bmatrix} = \begin{bmatrix} c\theta c\psi & c\theta s\psi & -s\theta \\ s\phi s\theta c\psi - c\phi s\psi & s\phi s\theta s\psi + c\phi c\psi & s\phi c\theta \\ c\phi s\theta c\psi + s\phi s\psi & c\theta s\theta s\psi - s\phi c\psi & c\phi c\theta \end{bmatrix} \begin{bmatrix} x_W \\ y_W \\ z_W \end{bmatrix} \quad (24)$$

The main forces acting on the vehicle come from gravity and the thrust of the rotors. Rotors drag and air friction are neglected here to simplify the model. It is classically assumed that the hexacopter is a rigid body with a symmetrical structure, and tensions in all directions are proportional to the square of the propeller speed. The equations of motion follow as:

$$\begin{aligned} \dot{\zeta} &= v \\ \dot{v} &= -ge_3 + \mathbf{R} \left(\frac{b}{m} \sum \Omega_i^2 \right) \\ \dot{\mathbf{R}} &= \mathbf{R}\hat{\omega} \\ \mathbf{I}\dot{\omega} &= -\omega \times \mathbf{I}\omega - \sum J_r(\omega \times e_3)\Omega_i\tau \end{aligned} \quad (25)$$

where \mathbf{R} the rotation matrix from \mathcal{R}_B to \mathcal{R}_W ; $\hat{\omega}$ is the skew symmetric matrix; Ω_i is the i -th rotor speed; \mathbf{I} the body inertia; J_r the rotor inertia; b is the thrust factor and τ is the torque applied to the body frame due to the rotors. A classical cascade control structure is adopted [43], where the built-in low-level controller of the RotorS package is used to track a reference in roll ϕ_r , pitch θ_r and thrust \mathcal{T} . The yaw angle ψ is kept constant without loss of generality. Under the small-angle assumption on ϕ and θ , the guidance model reduces to the double-integrator model:

$$\begin{aligned} \dot{\zeta} &= v \\ \dot{v} &= u = [u_x, u_y, u_z]^T \end{aligned} \quad (26)$$

where the computed accelerations are converted into low-level control inputs as:

$$\begin{aligned} \mathcal{T} &= m(u_z + g) \\ \theta_r &= \frac{m}{\mathcal{J}}(c\psi u_x + s\psi u_y) \\ \phi_r &= \frac{m}{\mathcal{J}}(s\psi u_x - c\psi u_y) \end{aligned} \quad (27)$$

3.3 Controller parametrization

For future use by the learning-based control approach (Section 4.4), a specific parametrization is adopted to control the system (26). The control objective is to stabilize the robot at a waypoint ζ_{ref} with a velocity $v = 0$. Defining $e = \zeta - \zeta_{\text{ref}}$ and taking into account unmodeled disturbances (e.g. wind) by considering an additional steady-state error variable $z = \int_0^t e(\tau)d\tau$, the augmented model with state vector $X = [z, e, v]$ becomes:

$$\begin{bmatrix} \dot{z} \\ \dot{e} \\ \dot{v} \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} z \\ e \\ v \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} u \quad (28)$$

The corresponding state-feedback controller, equivalent to a PID, is such that:

$$u = -k_i z - k_p e - k_d v \quad (29)$$

The poles of the closed-loop system are solution to the following equation:

$$\lambda^3 + \lambda^2 k_d + \lambda k_p + k_i = 0 \quad (30)$$

The controller has been re-parametrized using pole placement so as to guarantee convergence to steady-state without oscillations. This way, the action space for learning purposes is limited to desired solutions in the real part of the pole map, which prevents from sampling unnecessary solutions in the space of the control gains. The desired constants $\tau_1 > 0, \tau_2 > 0, \tau_3 > 0$ are then defined as:

$$\lambda_1 = \frac{-1}{\tau_1}; \lambda_2 = \frac{-1}{\tau_2}; \lambda_3 = -\frac{1}{\tau_3} \quad (31)$$

Since each one is solution to (30), it follows that:

$$\begin{bmatrix} 1 & \frac{-1}{\tau_1} & \frac{1}{\tau_1^2} \\ 1 & \frac{-1}{\tau_2} & \frac{1}{\tau_2^2} \\ 1 & \frac{-1}{\tau_3} & \frac{1}{\tau_3^2} \end{bmatrix} \begin{bmatrix} k_i \\ k_p \\ k_d \end{bmatrix} = \begin{bmatrix} \frac{1}{\tau_1^3} \\ \frac{1}{\tau_2^3} \\ \frac{1}{\tau_3^3} \end{bmatrix} \Leftrightarrow MK^T = N \quad (32)$$

Finally, the gains of the controller (29) are obtained as $K^T = M^{-1}N$:

$$\begin{aligned} k_i &= \frac{1}{\tau_1 \tau_2 \tau_3} \\ k_p &= \frac{\tau_1 + \tau_2 + \tau_3}{\tau_1 \tau_2 \tau_3} \\ k_d &= \frac{\tau_1 \tau_2 + \tau_1 \tau_3 + \tau_2 \tau_3}{\tau_1 \tau_2 \tau_3} \end{aligned} \quad (33)$$

4 Methodology

In this Section, we present how to build a model-free and learning-based controller with a DRL framework composed of two components: the Soft Actor-Critic algorithm and the Experience Replay technique. Later, we present how to apply it to the model-free and learning-based control schemes.

4.1 Soft Actor-Critic

In this study, we used the same policy gradient algorithm as in our previous paper [12] to either maximize the RL-based cost function of a model-free controller (see Section 2.2) or to estimate the controller parameters of a learning-based controller (as described in Section 2.3). The *Soft Actor-Critic* (SAC) algorithm proposed by T. Haarnoja et al. in [11] is based on three concepts which we formally define thereafter.

First, the SAC algorithm uses an actor-critic architecture which concurrently learns both State value and Q-value functions and the policy $\pi_\theta(a_t|s_t)$ as well. During the training stage, the critic part updates parameters of the *State* or Q-value functions while the actor network updates the policy parameters in the direction suggested by the critic. Note that a recent version of the SAC algorithm [44] allows to only estimate the Q-value function, while we used in this work the original version of the algorithm [11] where the State value function is also estimated to help stabilizing the overall training process. Precisely, our implementation of the SAC algorithm aims at iteratively learning three functions modeled by three NNs:

- a policy function $\pi_\theta(a_t|s_t)$ parameterized by the NN weights θ ,
- a soft Q-value function $Q_w(s_t, a_t)$ parameterized by the NN weights w ,
- and a soft State value function $V_\Psi(s_t)$ parameterized by the NN weights Ψ .

Beyond their modelization, these soft Q-value and State value functions are induced by the policy π_θ and defined as follows:

$$Q_w^{\pi_\theta}(s_t, a_t) = r_t(s_t, a_t) + \gamma \mathbb{E}_{s_{t+1} \sim \rho_\pi(s)} [V_\Psi^{\pi_\theta}(s_{t+1})] \quad (34)$$

$$V_\Psi^{\pi_\theta}(s_t) = \mathbb{E}_{a_t \sim \pi_\theta} [Q_w^{\pi_\theta}(s_t, a_t) - \alpha \log \pi_\theta(a_t|s_t)] \quad (35)$$

The terms $\rho_\pi(s)$ and $\rho_\pi(s, a)$ denote the state and the state-action marginals of the state distribution induced by the policy $\pi_\theta(a|s)$. The NNs parameters w and Ψ of the Q-value and State value functions respectively, are the solutions minimizing the soft Bellman residual $J_Q(w)$ and the mean squared error $J_V(\Psi)$:

$$J_Q(w) = \mathbb{E}_{(s_t, a_t) \sim D} \left[\frac{1}{2} \left(Q_w^{\pi_\theta}(s_t, a_t) - \left(r(s_t, a_t) + \gamma \mathbb{E}_{s_{t+1} \sim \rho_\pi(s)} [V_\Psi^{\pi_\theta}(s_{t+1})] \right) \right)^2 \right] \quad (36)$$

$$J_V(\Psi) = \mathbb{E}_{s_t \sim D} \left[\frac{1}{2} (V_\Psi(s_t) - \mathbb{E}[Q_w^{\pi_\theta}(s_t, a_t) - \log \pi_\theta(a_t, s_t)])^2 \right] \quad (37)$$

with $\bar{\Psi}$ the target value function. The NN parameters θ of the policy are updated in order to minimize the expected Kullback-Leibler divergence:

$$J_\pi(\theta) = \mathbb{E}_{s_t \sim D} \left[D_{KL} \left(\pi_\theta(\cdot | s_t) \left\| \frac{\exp(Q_w^{\pi_\theta}(s_t, \cdot))}{Z_w^{\pi_\theta}(s_t)} \right\| \right) \right] \quad (38)$$

with $Z_w^{\pi_\theta}$ a partition function used to normalize the distribution. The $(\cdot) \sim D$ means that the expected values are computed using the pair (s_t, a_t) and (s_t) sampled from a replay buffer D (see Section 4.2) in which the agent experience $e_t = (s_t, a_t, r_t, s_{t+1})$ is stored at each time-step. Moreover, the authors use an exponentially moving average, with a smoothing constant $\tau = 5e^{-3}$, to update the target value network weights w (the parameters Ψ are also update in the same way since they are directly related to w). Thus these weights are constrained to change slowly from one iteration to another. The name *Soft Actor-Critic* is based on this soft update procedure. Secondly, the SAC algorithm optimizes a stochastic policy in an off-policy manner [45]. This means that the policy used to explore the environment is different from the one that is being evaluated and improved. Finally, the policy is trained to maximize simultaneously the expected return and its entropy, which is, in some sense, a measure of its randomness:

$$J_{SAC}(\pi) = \sum_{t=1}^T \mathbb{E}_{(s_t, a_t) \sim \rho_\pi} [r(s_t, a_t) + \alpha H(\pi_\theta(\cdot | s_t))] \quad (39)$$

where:

$$H(\pi_\theta(\cdot | s)) = - \sum_{a \in A} \pi_\theta(a) \log \pi_\theta(a | s) \quad (40)$$

The term $H(\pi_\theta)$ is the entropy measure of policy π_θ and α is a fixed temperature parameter that determines the relative importance of the entropy term (an automatic adjustment mechanism for α was proposed in [44] by considering the constrained optimization problem of maximizing the expected return while satisfying a minimum entropy constraint but was not used in our implementation). The authors state that entropy maximization leads to policies that have better exploration capabilities. Note that this RL problem (39) is different from the initial RL problem (13). The resulting Q-value and V-value functions will therefore also include this entropy term. Two additional Q-functions are used to reduce positive bias in the policy improvement step, which is known to degrade performance of value based methods [46,47]. They are trained independently and the minimum of the two Q-functions is used as proposed in [47].

The SAC algorithm is part of the family of deep reinforcement learning approaches which try to exploit the strong representation capabilities offered by NNs to represent the Q-value, State-value and policy functions. The NNs structure of our SAC implementation is shown in Figure 4.

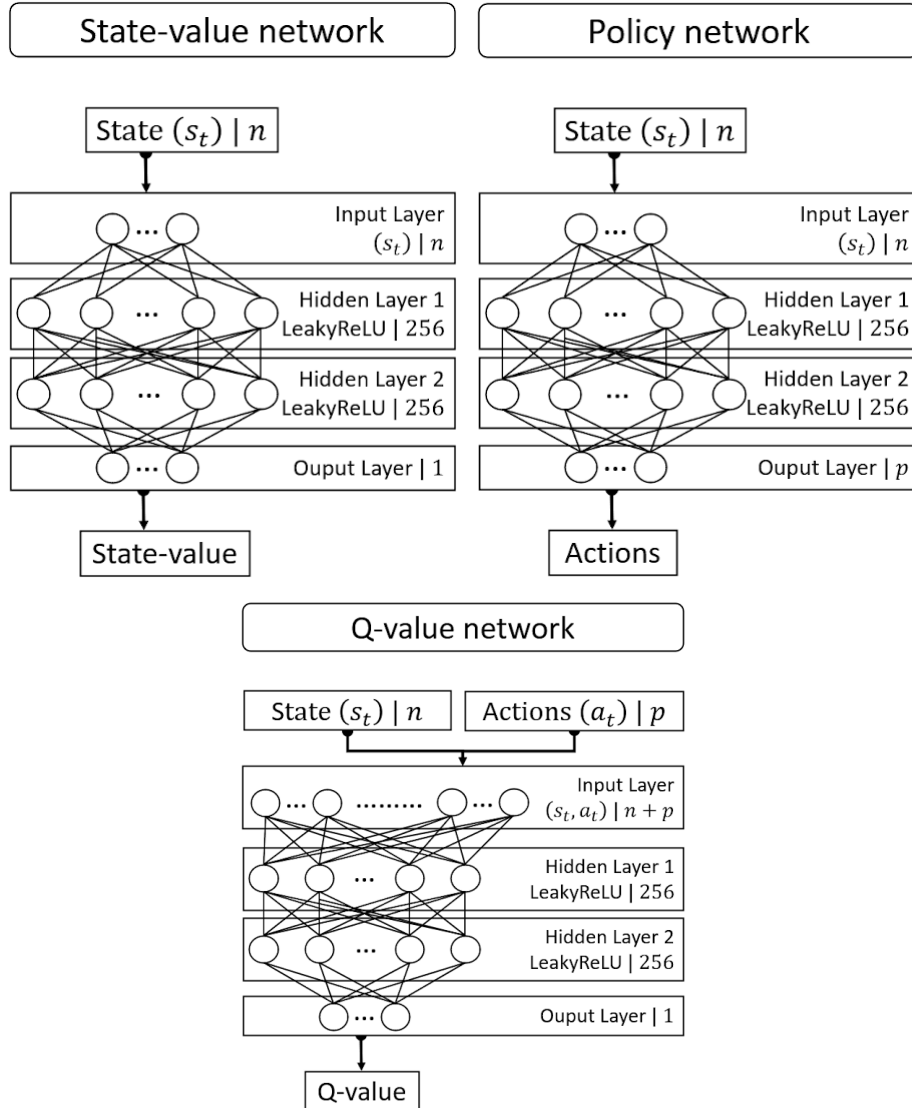


Fig. 4: The neural networks structure for our implementation of the SAC. The networks are composed of dense layers only and of two hidden layers for each network. Each layer is a fully-connected layer represented by its type, output size and activation function. The networks uses the same optimizer (Adam [48]), activation function (Leaky Relu [49]) and learning rate $l_r = 3e^{-4}$. Other parameters are a discount factor $\gamma = 0.99$ for (34), a number of 256 hidden units per hidden layer and one gradient update is performed at each time-step.

4.2 Experience Replay

To improve its performance, one needs to learn from its past actions. In [52], L. Lin showed that past experience of the learning agent should and could be used in an effective way by using what he called *Experience Replay* (ER). An *experience* is the result of a transition to a new state which can be framed as a quadruplet $e_t = (s_t, a_t, r_t, s_{t+1})$. The ER technique consists in storing at each time-step the experience e_t of the agent in a data-set (a replay buffer) $D = e_1, \dots, e_n$ of fixed and substantial size. During training, the RL algorithm updates are then applied to mini-batches of experiences randomly pooled over the stored samples to reinforce the functions estimates.

This intuitive concept can be refined by distinct means. It seems natural that when learning a new task, some experience might be more valuable than the others and should be used more often in the update process. This relevance is in general not directly accessible, therefore a specific criterion has to be chosen to quantify it. In [53], T. Schaul proposed the *Prioritized Experience Replay* (PER) technique which consists in sampling more often from the replay buffer the transitions with high expected learning progress, as measured by the magnitude of their temporal-difference (TD) error. By doing so, they were able to obtain state-of-the-art performance on the Atari 2600 benchmark suite.

Nevertheless, depending on the task, PER can be very computationally expensive. Another parameter that has been considered is the replay capacity that is the total number of transitions possibly stored in the replay buffer. In [54], S. Zhang showed that both large and small size replay buffers can significantly damage the performance. They proposed Combined Experience Replay (CER) to cope with this problem. It consists of adding the latest transition to the mini-batch pooled over the replay buffer which only requires $O(1)$ extra computation compares to PER to reduce the negative effect of the buffer size. The difference with PER is that by using CER, the latest transitions will undoubtedly be sampled. We decided to use the CER technique and to keep a mini-batch composed of randomly sample experiences $e_{rd,n}$ and the last one as:

$$\text{mini-batch} = [e_{rd,1} ; e_{rd,2} ; \dots ; e_{rd,n-1} ; e_{-1}] \quad (41)$$

with $n = 256$, the size of our mini-batch. We kept a replay capacity of 10^6 as recommended in the original paper of the SAC [11].

4.3 Application to the model-free controller

We will now present the use of the previous deep reinforcement learning framework to build a model-free controller. The methodology follows the same line of thought as our previous work [12]. The objective of the learning agent is to build a predictive model, using the SAC algorithm, that directly maps the actions (27) that control the MAV from the current state:

$$\pi_\theta : s_t \rightarrow [\phi_r ; \theta_r ; \mathcal{T}] \quad (42)$$

where $(\phi_r, \theta_r) \in [-\frac{\pi}{6}; +\frac{\pi}{6}]$ and $\mathcal{T} \in [m \times g; m \times (g + 3.0)]$ (with $m = 1.544\text{Kg}$, $g = 9.81\text{m.s}^{-1}$). In order to achieve the target rallying mission, we need to provide relevant data to the agent. Therefore, we defined the following vector o_t as the observation at the time-step t of the environment:

$$o_t = [a_{t-1} ; \phi ; \theta ; \psi ; v_x ; v_y ; v_z ; \omega_\phi ; \omega_\theta ; \omega_\psi ; \zeta_t ; e_t ; d_t] \quad (43)$$

where a_{t-1} are the last actions performed ; $[\phi; \theta; \psi]$ are the Euler angles of the robot (roll, pitch and yaw) ; $[v_x; v_y; v_z]$ and $[\omega_\phi; \omega_\theta; \omega_\psi]$ are its linear and angular velocities respectively in \mathcal{R}_W and \mathcal{R}_B ; ζ_t represents its position in \mathcal{R}_W (obtained from its embedded odometry sensor), $e_t = [e_x; e_y; e_z]$ are the current errors on the set-points in terms of Euclidean distance and d_t is the current Euclidean distance between the hexacopter and the target in \mathcal{R}_W . All the variables involved are assumed to be measured, their estimation is out of the scope of this work. We denote the target position by $\Lambda = [\Lambda_x; \Lambda_y; \Lambda_z]^T \in \mathcal{R}_W$. The dimension of this observation vector is 19 and it has been standardized to have zero mean and a variance of 1. In order to provide a higher time horizon to the agent, we constructed the state vector out of the current and past two observations vectors $[o_t; o_{t-1}; o_{t-2}]$. For the purpose of providing the agent a sense of “velocity” in the evolution of the state, we consider the two by two difference of these vectors such as $vel_t = (o_t - o_{t-1})$ and $vel_{t-1} = (o_{t-1} - o_{t-2})$. We went even further and tried to add a sense of “acceleration” in the evolution of the state by including the difference between the latter vectors $acc_t = (vel_t - vel_{t-1})$. The resulting state vector is therefore defined as:

$$s_t = [o_t ; o_{t-1} ; o_{t-2} ; vel_t ; vel_{t-1} ; acc_t] \quad (44)$$

The dimension of this state vector is 114. The reward function r_t has been designed in order to teach the agent how to complete the mission of target rallying:

$$r_t = \begin{cases} r_{receded} & \mathbf{if} \quad d_{rate} \leq 0 \\ r_{forward} & \mathbf{if} \quad d_{rate} > 0 \\ r_{reached} & \mathbf{if} \quad d_t \leq d_{reached} \\ r_{failed} & \mathbf{if} \quad z_w \notin [0.25; 20] \end{cases} \quad (45)$$

where r_t is the reward of the agent at time t ; d_{rate} is the distance rate to the target performed between the last two steps such as $d_{rate} = d_t - d_{t-1}$; $d_{reached}$ is the limit value beneath which we consider the target to be reached ; $z_w \in \mathcal{R}_W$ is the MAV altitude; both $r_{reached}$ and r_{failed} are terminal rewards determined at the end of the ongoing episode. Each of these terms represent the specific features of the desired behavior of the MAV:

- $r_{receded}$ is a constant negative reward equal to -20 that is sent to the agent whenever the robot is getting away from the target (or staying immobile).
- $r_{forward}$ is a positive reward sent when the relative distance to the target is decreasing. It is defined as:

$$r_{forward} = C_1 \times e \left(- \left[\left(\frac{d_t}{1+d_{rate}} \right) \times \frac{1}{C_2} \right]^2 \right)$$

With this design, we encourage the robot to move toward the target as fast as possible. We chose the value of the constants C_1 in order to scale the positive reward signal. This is particularly important because as mentioned in [11], the SAC algorithm is particularly sensitive to the scaling of the reward signal which is the magnitude of the reward value. We followed the recommendation prescribed in [11] and chose to set $C_1 = 20$ in order to obtain a positive reward scale of 20 (which we found in practice to be the reward scale that gave us the best performance for this problem). The constant C_2 represents how sparse is $r_{forward}$ based on the distance to the target. We chose the value $C_2 = 20$ empirically. Therefore, the positive reward is equal to:

$$r_{forward} = 20 \times e \left(- \left[\left(\frac{d_t}{1+d_{rate}} \right) \times \frac{1}{20} \right]^2 \right) \quad (46)$$

- A constant positive reward is sent to the agent when it succeeded to complete the mission, meaning $d_t \leq d_{reached}$. This generates $r_{reached} = +1000$.
- If the MAV altitude z_w exceed a predefined threshold, the constant negative reward $r_{failed} = -550$ is generated.

We define as *sampling rate*, the rate at which a state is sampled from the environment after the execution of the actions. A good practice consist in synchronizing it with the slowest sensor of the system which ensure no potential loss of information in the state vector and the fastest sampling rate. We synchronized it with the embedded odometry sensor which led to a *sampling rate* of about 20Hz.

4.4 Application to the learning-based controller

As stated in Section 2.3, learning-based methods consist in exploiting standard model-based control architectures that are either tuned or redesigned by a model-free algorithm in order to compensate for the unknown part of the model. The MAV is subject to an additive but unknown wind perturbation. Therefore, we combined the PID controller (29) under parametrization (33) within the DRL framework (Section 4.1) to iteratively auto-tune the feedback gains parameters (i.e. the constants τ_1 , τ_2 and τ_3), with the ambition of optimizing online a desired performance cost function. For behavior stability purpose, instead of directly trying to estimate the values of τ_i , the SAC algorithm estimate at each time-step small increments ∇_i that are added to the latter in order to cope with the wind disturbances. By doing so, we avoid jumping from a configuration of gains to a totally different one, which usually give rise to undesired oscillations. The gains (33) are then computed using $\tau_i = \tau_i + \nabla_i$ with $\nabla_i \in [-0.01; +0.01]$ (while making sure that $3 \geq \tau_1 \geq \tau_2 \geq \tau_3 \geq 5 \times 10^{-3}$) and the resulting control law (29) is finally applied. The τ_i parameters are initialized to the nominal configuration $\tau_1 = 1, \tau_2 = 2.5, \tau_3 = 0.875$. The objective of the learning agent is now to build a predictive model that directly maps ∇_i from the current state to adjust the gains starting from this initial configuration:

$$\begin{aligned} \pi_\theta : s_t \rightarrow [& \nabla_{\tau_{1_{roll}}} ; \nabla_{\tau_{2_{roll}}} ; \nabla_{\tau_{3_{roll}}} ; \nabla_{\tau_{1_{pitch}}} ; \\ & \nabla_{\tau_{2_{pitch}}} ; \nabla_{\tau_{3_{pitch}}} ; \nabla_{\tau_{1_{thrust}}} ; \nabla_{\tau_{2_{thrust}}} ; \nabla_{\tau_{3_{thrust}}}] \end{aligned} \quad (47)$$

We used a slightly different observation vector from the model-free scheme (43). Indeed, we propose to add the resulting parameters and the PID controller outputs (29) in the observation vector:

$$o_t = [a_{t-1} ; \tau_{rpt} ; pid_{rpt} ; k_{p_{roll}} ; k_{i_{roll}} ; k_{d_{roll}} ; k_{p_{pitch}} ; k_{i_{pitch}} ; k_{d_{pitch}} ; k_{p_{thrust}} ; k_{i_{thrust}} ; k_{d_{thrust}} ; \phi ; \theta ; \psi ; v_x ; v_y ; v_z ; \omega_\phi ; \omega_\theta ; \omega_\psi ; \zeta_t ; e_t ; d_t] \quad (48)$$

where $dim(\tau_{rpt}) = 9$ and $pid_{rpt} = [\phi_r ; \theta_r ; \mathcal{T}]$. The dimension of this observation vector is 37. We constructed the state vector out of this observation vector as earlier (44), resulting to a state vector of dimension 222. We used the exact same reward function and parameters as described earlier (45) to train this controller.

5 Results

A training episode is defined as follows: at the beginning of the episode, the MAV is set at the center of the environment at an altitude of 3 meters with roll, pitch and yaw angles equal to 0. A target is then initialized at a fixed and uniformly random position $\Lambda = [\Lambda_x ; \Lambda_y ; \Lambda_z]^T$ with $[\Lambda_x ; \Lambda_y]^T \in [-20; -5 \cup 5; 20]$ and $\Lambda_z \in [2; 20]$. The mission then begins and is considered as a success if the relative distance to the target is inferior to a predefined threshold $d_{reached}$ and as a failure if an error signal is generated, both cases ending the episode. Otherwise, the episode is ended if the number of iteration steps reach the maximum value allowed per episode that is set at 300. The training for both controllers consisted in performing 1 000 000 iterations in a environment with a varying wind field (as described in Section 3.1). To help the agent, $d_{reached}$ is reduced during training as follows: at first $d_{reached} = 3\text{m}$, from iteration 250 000th we set $d_{reached} = 2\text{m}$ and from iteration 500 000th we set $d_{reached} = 1\text{m}$. The PyTorch framework [50] was used to carry out the numerical experiments, along with the CUDA toolkit [51] and an RTX 2060 GPU card, allowing us to perform the training of one controller in approximately 10 hours. It can be seen in Figure 5 that during training, both the learning-based (LB) and model-free (MF) controllers were able to reach a high success rate under unknown wind gust disturbances, which shows the applicability of the SAC DRL procedure for this type of aerial navigation problems. The LB strategy presents a much higher convergence speed to a significant success rate than the MF strategy, which shows the great potential of combining model-based classical controllers with learning procedures. This can be explained by the model-based part of the LB controller which allows it to choose relatively good actions despite being at the early stage of the training session. Therefore, from the beginning of the training the LB controller is able to explore a much higher part of the reward space than the one of the MF controller. We believe this significantly helps the DRL algorithm to find better overall strategies. The evaluation consisted in performing the target rallying mission in areas of the same environment that had never been explored by neither controllers during training (i.e. the wind field in these areas were

totally unknown to the neural networks) with $d_{reached} = 1\text{m}$. The evaluation is composed of a total of 500 episodes, different from each other in terms of target position and with a max step size per episode of 1000. The targets during evaluation were uniformly distributed in the space defined by $\Lambda = [\Lambda_x; \Lambda_y; \Lambda_z]^T$ with $[\Lambda_x; \Lambda_y]^T \in [-50; -20 \cup] 20; 50]$ and $\Lambda_z \in [2; 20]$. The same set of evaluation episodes was used for each controller. We also evaluated a fixed control strategy which consisted in a PID controller with the fixed nominal poles configuration (4.4). The outcomes of this evaluation are provided in Table 1. We can see that the LB controller achieves a better performance for each metric. Noted that it requires, on average, less actions to achieve the task with the LB controller despite sharing the same *sampling rate*. The mean reward per step of the LB controller is more than 2 times higher than the MF one. On the other hand, the fixed control strategy is showing a much lower success rate, close to 50%. We observed that with this strategy, failures mostly consisted in cases where the MAV is close to the steady-state but is being deviated by the wind gust. The vehicle is then not able to recover with fixed-pole parametrization. It shows the benefits of adaptive control methods when facing unknown disturbances.

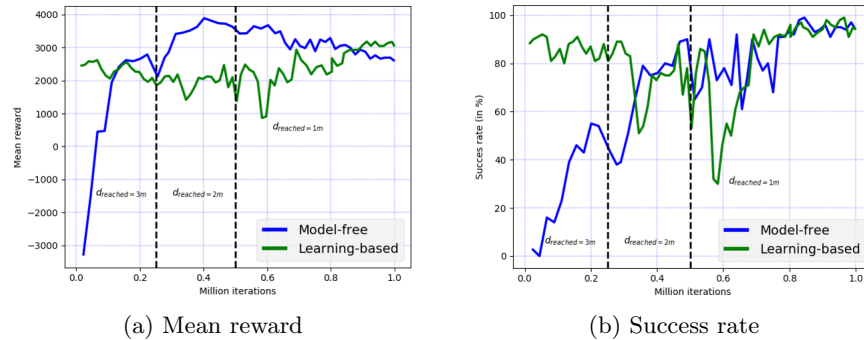


Fig. 5: Training curves showing the mean reward and success rate computed per 100 episodes over a moving window of 100 episodes. Note that with SAC, the policy is trained to maximize also the entropy, therefore the mean action does not always correspond to the optimal action for the maximum return objective.

Controller type	Mean step number	Mean total reward	Mean reward per step	Success rate	Positive reward rate
Fixed Poles PID	488	710.797	1.454	50.6%	61.197%
Model-free	357	2238.970	6.256	74.2%	86.056%
Learning-based	281	4034.434	14.346	91.6%	89.2%

Table 1: Evaluation results.

6 Conclusions

We present a novel application of the SAC procedure for learning both model-free and learning-based adaptive controllers applied to the autonomous navigation of a MAV under wind gust conditions. These two strategies have been trained and compared in the same ROS-Gazebo reference simulation. Both strategies were able to reach high levels of performances under unknown uncertainty conditions, but the learning-based scheme, with a judicious parametrization of its model-based control part, exhibited a much faster convergence rate. Our results suggest that learning-based adaptive methods can be much more efficient than model-free ones and allow a better stability analysis in the DRL scheme.

Acknowledgements The authors thank Dr. Estelle Chauveau from Naval Group for the help provided. This work was supported by SENI, the research laboratory between Naval Group and ENSTA Bretagne.

References

1. Z. T. Dydek and A. M. Annaswamy and E. Lavretsky: *Adaptive Control and the NASA X-15-3 Flight Revisited*. In: IEEE Control Systems Magazine, Vol. 30, nb. 3, pp. 32-48, 2010.
2. William J. Terrell: *Some fundamental control Theory vol. 2: Feedback linearization of single input nonlinear systems*. In: American Mathematical Monthly, 1999.
3. John C. Doyle: *Robust and optimal control*. In: IEEE Conference on Decision and Control, 1995.
4. Vadim I. Utkin: *Sliding Modes in Control and Optimization*. In: Communication and Control Engineering Series, 1992.
5. R. E. Bellman and L. A. Zadeh: *Decision-making in a fuzzy environment*. In: Management Science, 1970.
6. B. D. Anderson: *Failures of adaptive control theory and their resolution*. In: Communications in Information and Systems, Vol. 5, No. 1, pp. 1-20, 2005.
7. Niko Sünderhauf and O. Brock and W. Scheirer and Raia Hadsell and D. Fox and J. Leitner and B. Uprocft and P. Abbeel and W. Burgard and Michael Milford and P. Corke: *The limits and potentials of deep learning for robotics*. In: The International Journal of Robotics Research, 2018.
8. Gabriel Dulac-Arnold and N. Levine and Daniel J. Mankowitz and J. Li and Cosmin Paduraru and Sven Gowal and T. Hester: *An empirical investigation of the challenges of real-world reinforcement learning*, 2020.
9. J. Garcia and F. Fernández: *A comprehensive survey on safe reinforcement learning*. In: Journal of Machine Learning Research, 2015.
10. Mouhacine Benosman: *Learning-based adaptive control for nonlinear systems*, 2016.
11. T. Haarnoja, Aurick Zhou, P. Abbeel, S. Levine: *Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor*. In: International Conference on Machine Learning (ICML), 2018.
12. T. Chaffre and J. Moras and A. Chan-Hon-Tong and J. Marzat: *Sim-to-Real Transfer with Incremental Environment Complexity for Reinforcement Learning of Depth-Based Robot Navigation*. In: Proceedings of the 17th International Conference on Informatics in Control, Automation and Robotics, 2020.
13. Karl J. Astrom and Bjorn Wittenmark: *Adaptive Control, second edition*. In: Library of Congress Cataloging-in-Publication Data, published by DOVER PUBLICATIONS, INC, 2008.

14. Nikolai M. Filatov and H. Unbehauen: *Adaptive Dual Control Theory and Applications*. In: Lecture Notes in Control and Information Sciences, Vol. 302, published by Springer, 2004.
15. G. Stein: *Adaptive flight control: a pragmatic view*. In: Applications of Adaptive Control, eds. K. S. Narendra and R. V. Monopoli. New York: Academic Press, 1980.
16. Osburn, P. V. and H. P. Whitaker and A. Kezer: *New developements in the design of adaptive control systems*. In: Institute of Aeronautical Sciences, 1961.
17. U. Borisson: *Self-tuning regulators for a class of multivariable systems*. In: Automatica, 1979.
18. Sternby, J.: *A simple dual control problem with an analytical solution*. In: IEEE Transactions on Automatic Control, 1976.
19. A. Pyrkin and A. Bobtsov and S. Kolyubin and O. Borisov and Vladislav Gromov and S. Aranovskiy: *Output controller for quadcopters with wind disturbance cancellation*. In: IEEE Conference on Control Applications (CCA), 2014.
20. A. Razinkova and I. Gaponov and H.-C. Cho: *Adaptive control over quadcopter UAV under disturbances*. In : International Conference on Control, Automation and Systems (ICCAS), 2014.
21. R. S. Fernández and Sergio Domínguez and P. Campoy: *L1 adaptive control for Wind gust rejection in quad-rotor UAV wind turbine inspection*. In: International Conference on Unmanned Aircraft Systems (ICUAS), 2017.
22. C. Cao and N. Hovakimyan: *L1 Adaptive Output Feedback Controller to Systems of Unknown Dimension*. In: American Control Conference, 2007.
23. C. Cao and N. Hovakimyan: *Adaptive controller for systems with unknown time-varying parameters and disturbances in the presence of non-zero trajectory initialization error*. In: International Journal of Control, 2008.
24. S. Kim and Vedang M. Deshpande and R. Bhattacharya: *\mathcal{H}_2 Optimized PID Control of Quad-Copter Platform with Wind Disturbance*. In: International Conference on Unmanned Aircraft Systems (ICUAS), 2020.
25. Lucas M. Argentim and Willian C. Rezende and Paulo E. Santos and Renato A. Aguiar: *PID, LQR and LQR-PID on a quadcopter platform*. In: International Conference on Informatics, Electronics and Vision (ICIEV), 2013.
26. K. Ariyur and M. Krsti: *Real time optimization by Extremum Seeking control*, 2003.
27. Chunlei Zhang and R. Ord: *Extremum-Seeking Control and Applications: A Numerical Optimization-Based Approach*, 2011.
28. Andrea Tagliabue and Xiangyu Wu and M. Mueller: *Model-free Online Motion Adaptation for Optimal Range and Endurance of Multicopters*. In: International Conference on Robotics and Automation (ICRA), 2019.
29. R. Sutton and A. Barto: *Reinforcement Learning: An Introduction*. In: IEEE Transactions on Neural Networks, 2005.
30. Howard A. Ronald : *Dynamic Programming and Markov Processes*. In: The M.I.T. Press., 1960.
31. J. Schulman and P. Moritz and S. Levine and Michael I. Jordan and P. Abbeel: *High-Dimensional Continuous Control Using Generalized Advantage Estimation*. In: Computing Research Repository, 2016.
32. Antonio Loquercio and Ana I. Maqueda and Carlos R. del-Blanco and D. Scaramuzza: *DroNet: Learning to Fly by Driving*. In: IEEE Robotics and Automation Letters, 2018.
33. G. Kahn and A. Villafior and B. Ding and P. Abbeel and S. Levine: *Self-Supervised Deep Reinforcement Learning with Generalized Computation Graphs for Robot Navigation*. In: IEEE International Conference on Robotics and Automation, 2018.

34. Nihal Soans and Ehsan Asali and Yi Hong and P. Doshi: *SA-Net: Robust State-Action Recognition for Learning from Observations*. In: IEEE International Conference on Robotics and Automation, 2020.
35. F. W. Lewis, S. Jagannathan, and A. Yesildirak: *Neural Network Control Of Robot Manipulators And Non-Linear Systems*. In: CRC Press, 1998.
36. M. Bisheban and T. Lee: *Geometric Adaptive Control for a Quadrotor UAV with Wind Disturbance Rejection*. In: IEEE Conference on Decision and Control, 2018.
37. M. Bisheban and T. Lee: *Geometric Adaptive Control with Neural Networks for a Quadrotor UAV in Wind fields*. In: IEEE Transactions on Control Systems Technology, 2019.
38. L. Zheng and Rui Yang and Jiesen Pan and H. Cheng and H. Hu: *Safe Online Learning Tracking Control for Quadrotors under Wind Disturbances*, 2020.
39. C. Rasmussen and H. Nickisch: *Gaussian Processes for Machine Learning (GPML) Toolbox*. In: Journal of Machine Learning Research, 2010.
40. Nathan P. Koenig and A. Howard: *Design and use paradigms for Gazebo, an open-source multi-robot simulator*. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, 2004.
41. M. Quigley: *ROS: an open-source Robot Operating System*. In: IEEE International Conference on Robotics and Automation, 2009.
42. F. Furrer and M. Burri and M. Achtelik and R. Siegwart: *RotorS—A Modular Gazebo MAV Simulator Framework*. In: Robot Operating System (ROS): The Complete Reference (Volume 1), published by Springer International Publishing, 2016.
43. S. Bertrand and N. Guénard and T. Hamel and H. Piet-Lahanier and L. Eck: *A hierarchical controller for miniature VTOL UAVs: Design and stability analysis using singular perturbation theory*. In: Control Engineering Practice, Vol. 10, Nb. 10, Pp. 1099–1108, 2011.
44. T. Haarnoja and Aurick Zhou and Kristian Hartikainen and G. Tucker and Sehoon Ha and J. Tan and V. Kumar and H. Zhu and A. Gupta and P. Abbeel and S. Levine: *Soft Actor-Critic Algorithms and Applications*, 2018.
45. T. Degris and Martha White and R. Sutton: *Off-Policy Actor-Critic*, 2012.
46. H. V. Hasselt: *Double Q-learning*. In: Conference on Neural Information Processing Systems, 2010.
47. Scott Fujimoto and H. V. Hoof and David Meger: *Addressing Function Approximation Error in Actor-Critic Methods*. In: IEEE ICML, 2018.
48. Diederik P. Kingma and J. Ba: *Adam: A Method for Stochastic Optimization*. In: Computing Research Repository, 2015.
49. B. Xu and Naiyan Wang and T. Chen and Mu Li: *Empirical Evaluation of Rectified Activations in Convolutional Network*, 2015.
50. P. Adam and G. Sam and M. Francisco and L. Adam and B. James and C. Gregory and K. Trevor and L. Zeming and G. Natalia and A. Luca and D. Alban and K. Andreas and Y. Edward and D. Zachary and R. Martin and T. Alykhan and C. Sasank and S. Benoit and F. Lu and B. Junjie and C. Soumith: *PyTorch: An Imperative Style, High-Performance Deep Learning Library*. In: Advances in Neural Information Processing Systems 32, 2019.
51. John Nickolls and I. Buck and M. Garland and K. Skadron: *Scalable parallel programming with CUDA*. In: IEEE Hot Chips 20 Symposium, 2008.
52. L. Lin: *Reinforcement learning for robots using neural networks*. In: Technical report, DTIC Document, 1993.
53. T. Schaul and J. Quan and I. Antonoglou and D. Silver: *Prioritized Experience Replay*. In: Computing Research Repository, 2016.
54. S. Zhang and R. Sutton: *A Deeper Look at Experience Replay*, 2017.