

Deep reinforcement learning of airfoil pitch control in a highly disturbed environment using partial observations

Diederik Beckers *

Graduate Aerospace Laboratories, *California Institute of Technology*, Pasadena, California 91125, USA

Jeff D. Eldredge 

Mechanical and Aerospace Engineering, *University of California, Los Angeles*, California 90095-1597, USA



(Received 4 April 2024; accepted 28 August 2024; published 12 September 2024)

This study explores the application of deep reinforcement learning (RL) to design an airfoil pitch controller capable of minimizing lift variations in randomly disturbed flows. The controller, treated as an agent in a partially observable Markov decision process, receives non-Markovian observations from the environment, simulating practical constraints where flow information is limited to force and pressure sensors. Deep RL, particularly the TD3 algorithm, is used to approximate an optimal control policy under such conditions. Testing is conducted for a flat plate airfoil in two environments: a classical unsteady environment with vertical acceleration disturbances (i.e., a Wagner setup) and a viscous flow model with pulsed point force disturbances. In both cases, augmenting observations of the lift, pitch angle, and angular velocity with extra wake information (e.g., from pressure sensors) and retaining memory of past observations enhances RL control performance. Results demonstrate the capability of RL control to match or exceed standard linear controllers in minimizing lift variations. Special attention is given to the choice of training data and the generalization to unseen disturbances.

DOI: [10.1103/PhysRevFluids.9.093902](https://doi.org/10.1103/PhysRevFluids.9.093902)

I. INTRODUCTION

Large flow disturbances can drastically alter the aerodynamic forces on an airfoil and even induce flow separation, possibly compromising the safety, performance, and lifetime of aerodynamic vehicles and energy-conversion systems. Flow disturbances can appear as localized flow features, or gusts, that are incident on the airfoil, either as discrete instances or as continuous variations, or they can manifest as unsteady ambient flow, for instance, a varying freestream flow. The interaction of large-amplitude flow disturbances with an airfoil and the flow response to actuation is generally characterized by the nonlinear formation and shedding of leading- and trailing-edge vortices and is not always fully understood for every type of actuation or disturbance. Furthermore, the flow response during a rapid succession of such actuations and disturbances cannot be simply superposed [1] and their large parameter spaces [2] complicate the design of a comprehensive control strategy. This is further complicated by the limited availability of sensor information in practical applications.

Given the control objectives (e.g., lift enhancement, drag reduction, stall delay, etc.), the actuation method (e.g., pitching, control surface actuation, synthetic jets, plasma actuators, etc.), and the available measurements (e.g., pressure, velocity, force, torque, etc.), the control of unsteady flows over airfoils requires the selection of a specific control strategy and design approach. Active control strategies can be predetermined, feedforward, or feedback control, and the design approach can be

*Contact author: beckers@caltech.edu

signal-based, frequency-based, or based on numerical optimization [3]. In the domain of unsteady aerodynamics, common control systems target the reduction of lift variations due to unstable states or disturbances using feedback control. Signal-based approaches to such control systems assume that the unstable behavior or the shape or time profile of the disturbance is known and examples include the use of a linear quadratic regulator to suppress vortex shedding using point forcing near the trailing edge [4] and the lift regularization for a flat plate encountering a transverse gust through pitching based on heuristics [5]. Frequency-based control aims first to shape the open-loop or closed-loop transfer functions to achieve specific robustness or performance objectives and then design a controller whose transfer functions approximate those shapes. Examples of such control for lift mitigation in unsteady aerodynamics include controller designs based on \mathcal{H}_∞ synthesis to reject vertical acceleration disturbances for a flat plate through pitching [6–8] and to reject longitudinal gusts for semicircular wings using pulsed blowing [9].

The previous examples employed linear controllers that are designed using a linearized model of the flow. However, for large flow disturbances or a rapid succession of disturbances and actuation, modeling and controlling nonlinear effects become increasingly important. Simultaneously, the extensive parameter spaces associated with disturbances and actuation impede the practical application of linearized models. Nonlinear controllers can be designed using the numerical optimization approach, which is witnessing rapid advancement due to contributions emerging from the machine learning and data science domains. Notable flow control examples include model predictive control of cylinder rotation to control the wake [10], iterative maneuver optimization to rapidly design pre-determined nonlinear pitch control for lift regulation during gust encounters [11], and reinforcement learning, which we discuss next.

Reinforcement learning solves control problems formulated as Markov decision processes by iteratively learning optimal control through trial-and-error interactions with an environment. In its early applications, reinforcement learning (RL) relied on tabular representations for the value of every possible state in a Markov decision process to optimize a control law, or policy, which is only tractable for modest-dimensional, discrete problems. In the past decade, however, the successful integration of deep learning into RL (forming deep RL) enabled the approximation of policies and value functions over high-dimensional continuous spaces using multilayered artificial neural networks, leading to many breakthroughs in artificial intelligence, especially in the fields of autonomous systems and controls [12]. Furthermore, the use of policy and value function approximations extends RL to solving partially observable Markov decision processes (POMDP) [13,14], i.e., Markov decision processes where the full state of the system cannot be directly measured and whose exact solutions would be intractable except for the smallest problems. Thanks to these features, deep RL has been successfully applied to the control or modeling of high-dimensional fluid flows with uncertainty about the underlying state, e.g., navigating through wakes [15–17] or vortex shedding suppression [18–20] with limited sensor information or uncertainty about model parameters, e.g., from turbulence models [21,22]. Similar to the objective of the previously introduced signal- and frequency-based control problems, Renn and Gharib [23] apply deep RL to learn airfoil flap control to minimize lift variations on an airfoil positioned downstream of a vortex-shedding cylinder using only lift measurements.

In this work, we apply deep RL for learning airfoil pitch control to minimize lift variations due to flow disturbances in two types of flow environments that model the flow around the airfoil, demonstrated in Fig. 1: one type that models the flow purely as a potential flow with small disturbances, thus allowing us to simplify the system to a classical aerodynamics setup with relatively few states, and another type that simulates the viscous flow at a low Reynolds number with large disturbances, necessitating the use of a high-dimensional model of this nonlinear system. We use these environments to focus on two aspects of this reinforcement learning problem: the partial observability, or how the performance of the control depends on the type and availability of measurements of the state of the system, and the generalizability of the trained controller, or how a policy trained on one type of disturbance performs when faced with different types of disturbances. To apply reinforcement learning to the control problem, we first formulate it as a (partially observable) Markov decision

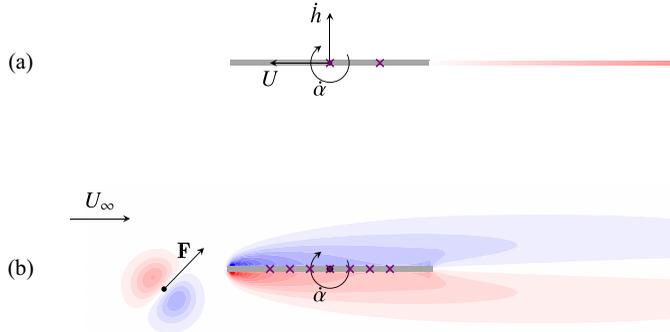


FIG. 1. (a) Diagram of the airfoil’s motion and wake for the classical unsteady aerodynamics environment. (b) Diagram of the airfoil, wake, freestream, and forcing for the viscous flow environment. The crosses indicate the positions of the pressure sensors, when used.

process and select an appropriate learning algorithm, described in Sec. II. We then analyze the performance of the RL pitch control using different training setups in an environment with the classical unsteady aerodynamics model in Sec. III and compare it with a linear-quadratic-Gaussian controller for smooth and impulsive vertical acceleration disturbances. In Sec. IV, we perform a similar analysis of the RL performance in an environment with the viscous flow model for flow disturbances introduced upstream of the airfoil with a pulsed point force.

II. PITCH CONTROL LEARNING FRAMEWORK

Consider a flat plate airfoil oriented at an angle α with the negative x axis of an inertial reference frame (positive in the clockwise direction). We describe the kinematics of the airfoil in terms of the velocity $(-U, \dot{h})$ of a pivot point on the chord and the angular velocity $\dot{\alpha}$ about this pivot point. We restrict the horizontal motion of the pivot point to a constant velocity with magnitude U in the negative x direction. Equivalently for a flat plate, we can also analyze the airfoil from the perspective of a reference frame with the same orientation as the inertial reference frame, centered at and moving with the pivot point. In this specific moving reference frame, which we will refer to as the wind tunnel reference frame, the flow around the airfoil from its horizontal motion is modeled with a constant horizontal freestream velocity $U_\infty = U$. In this work, we focus on the control of the airfoil’s pitch by varying the angular acceleration $\ddot{\alpha}$ about the pivot point while the airfoil encounters disturbances that affect the aerodynamic forces that act on it. The specific disturbances we will study consist of randomized uniform vertical flow accelerations \dot{h} or vortex structures generated by a randomized forcing \mathbf{F} in the flow upstream of the airfoil. A set of measurements is fed back to the controller. These can consist of the angle α , angular velocity $\dot{\alpha}$, lift f_y , pressure p , and the states encoding the wake evolution in the Wagner model, with the exact combination of measurements depending on the case.

The control problem is to design the discrete-time feedback controller \mathbf{K} in Fig. 2 that pitches the airfoil such that lift variations about a prescribed reference are minimized and α , $\dot{\alpha}$, and $\ddot{\alpha}$ are limited to prescribed ranges. In this work, we set the reference lift to zero, but the so-called servo problem of tracking other reference lifts can be explored by straightforward extension.

A. Pitch control problem formulated as a POMDP

We can formulate the control problem as a POMDP framework [13,24] following the agent-environment structure depicted in Fig. 3. Here, the agent assumes the role of the controller and the environment contains all the processes that the agent cannot control arbitrarily, i.e., the aerodynamics model, the disturbance generator, and the reward and observation functions. If we

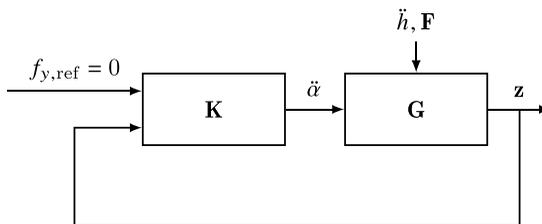


FIG. 2. Control diagram for the pitch control problem.

were to consider a nonzero reference lift, the process that generates this signal would also have to be included in the environment. The aerodynamics model and the disturbance and reference lift generators together define the probabilistic state transition function that, during each time step k , advances the environment state S_k to S_{k+1} as a result of the latest action $A_k := \ddot{\alpha}_k$ of the agent. This action remains piecewise constant over the time between two agent-environment interactions. The agent then receives a partial observation O_{k+1} of the environment and a reward R_{k+1} resulting from that action, which it maps to its next action A_{k+1} . The POMDP formalizes the sequential decision making process consisting of the sequence $O_1, A_1, \dots, O_k, A_k, R_{k+1}, O_{k+1}, \dots$. For simplicity, we only study sequences of finite length, with each such sequence called an *episode*. Solving the POMDP amounts to finding the optimal policy that maximizes the expected sum of rewards in one episode, also called the episode *return*, thus finding a decision-making strategy that selects actions accounting for potential long-term rewards but also uncertainty about the environment.

The state S_k of the environment is composed of the state of the aerodynamics model at time step k , which in our work comprises the kinematic state of the airfoil, the state of the flow around it, which can be stochastically altered by the disturbance generator, the previous action A_{k-1} for use in the reward function, and any variables needed to construct the elements of the measured variables in \mathbf{z} . The initial state of the aerodynamic model at the beginning of every episode is the steady-state motion of the airfoil at zero angle of attack and the corresponding steady-state flow solution. We will impose bounds on the environment’s state for training and stability purposes. If these bounds are violated at any point during the episode, the episode is terminated early and reset to the initial conditions.

The observation function maps the state and action to an observation for the agent. In our model, we assume that this is done without additional noise. For the cases in this work, O_k always contains $\alpha_k, \dot{\alpha}_k$, and the latest available lift f_y . Since the lift generally depends on the action at the same instant (through the inertial reaction of the fluid), we assume that the agent only has knowledge of

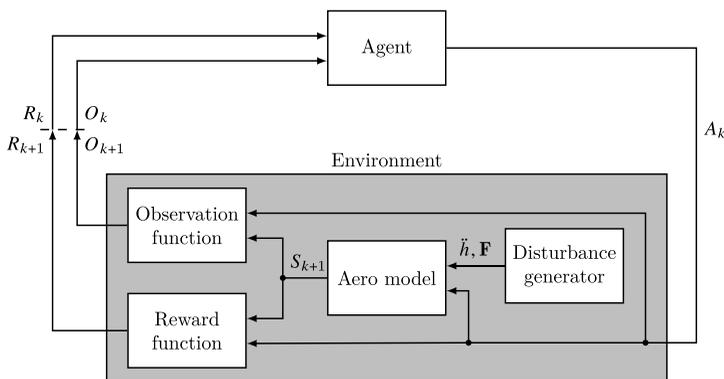


FIG. 3. Diagram of the POMDP learning framework for the pitch control problem.

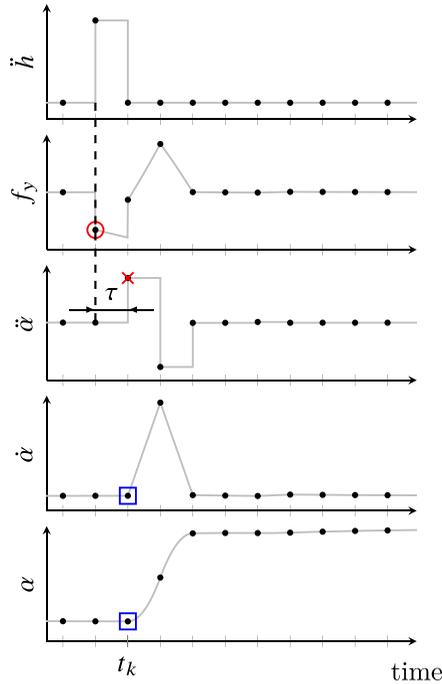


FIG. 4. Behavior of the lift f_y , control input $\ddot{\alpha}$, and angular states $\dot{\alpha}$ and α during the flow response to an impulsive \ddot{h} disturbance with output feedback control. The control input is piecewise constant between time steps and the pitching is applied at the midchord point. The lift output observed by the controller is delayed, in this case by the same time as the action time step, and the feedback controller’s action at time step k (\times) is therefore based on the lift from the previous time step (\circ). In contrast, for the control decision at that same time step, all the RL agents in this work additionally observe the angular states at k (\square) and potentially the pressure or other wake information (not shown).

the lift from an earlier time instant, creating a delay in the response to a disturbance, as illustrated in Fig. 4. We will assume that this output delay is also present in other control strategies that we use for comparison in this paper. Though this delay limits the controller’s performance, it is likely to arise in practice.¹ Depending on the studied case, the observation vector may also include the latest available pressure values at the airfoil or, in one case, the full state of the flow, all from the same time instant as the lift.

The final element of the environment is the reward function, which maps the new state S_{k+1} and the last action A_k to a reward signal R_{k+1} , used in the objective function to optimize the policy. For this reason, the reward function should introduce the control objective into the POMDP framework.

B. Agent and the learning algorithm

The agent’s task is to find the optimal policy for its actions, leveraging the rewards and observations it receives following each action for learning this policy. In the online reinforcement learning approach, it considers the environment as a black box and learns about its behavior

¹In a realistic setup, the measurements at one time instant inform the action at the next. Our setup is similar, but by providing the controller with the angular information at the time instant of the next control input, these angular states are assumed to be predicted one time step ahead, which is possible since they only rely on previous control inputs.

while simultaneously interacting with it. To balance between exploiting the learned policy to maximize rewards and exploring new actions to discover potentially better strategies, the agent often incorporates stochastic behavior during its interactions with the environment.

In a realistic setting, the state of the environment is continuous-valued and usually highly dimensional due to the flow dynamics. The observations rarely encompass the entire flow field, rendering the environment's state only partially observable and resulting in the agent receiving non-Markovian signals. In the POMDP formalism, the agent addresses this uncertainty by maintaining a belief state—a probability distribution reflecting the true state of the environment [24]. This belief state informs optimal decision making, with the agent employing a policy that maps these belief states to actions. Alternatively, effective strategies can take a more straightforward approach, partially restoring the Markov property by utilizing a fixed-width finite history window of recent observations fed to a neural network or a recurrent neural network [25,26]. These memory-based approaches to POMDPs have been applied in recent years with great success using model-free deep reinforcement learning techniques [27–29] and we adopt here the deep reinforcement learning approach of using a fixed-width finite history window of recent observations.

Model-free approaches require sampling many interactions with the environment, which can be computationally expensive for high-fidelity flow simulators. This warrants the use of an off-policy learning algorithm, which makes more efficient use of each sampled interaction than on-policy algorithms. Because we also desire a control law for a continuous-valued action, the training uses the twin delayed deep deterministic policy gradient (TD3) algorithm [30], which we found in previous work [31] to perform better in this problem than alternative state-of-the-art, off-policy, continuous-action algorithms, such as soft actor-critic [32]. TD3 employs a deterministic policy and stochastically explores actions and states during training by adding normally distributed *action noise* to the actions chosen by the latest policy. Details of the algorithm's implementation in this work can be found in Appendix A.

III. LEARNING IN A CLASSICAL UNSTEADY AERODYNAMICS ENVIRONMENT

In this first example of the control analysis, we consider an inviscid, two-dimensional incompressible flow and disturbances that are either random vertical accelerations \ddot{h} of the plate's pivot point or, equivalently for a flat plate, random accelerations of a vertical freestream flow. If we assume the resulting velocities from the disturbance and pitch control are small compared to the forward velocity, i.e., $\dot{h} \ll U$ and $\dot{\alpha} c \ll U$, then we can assume that the wake's vorticity is concentrated in an infinitesimally thin vortex sheet that convects away from the plate's trailing edge with the relative velocity U and remains parallel to the plate, illustrated in the top panel of Fig. 1. Under these assumptions, the theory of classical unsteady aerodynamics as analyzed by Wagner [33] and Theodorsen [34] is valid and the following linear dependence of the lift coefficient $c_L = 2f_y/(\rho U^2 c)$ on the disturbance \ddot{h} and control input $\ddot{\alpha}$ can be derived (see Appendix B):

$$\begin{bmatrix} \dot{\alpha}_{\text{eff}} \\ \ddot{\alpha} \\ \dot{\hat{\mathbf{x}}} \end{bmatrix} = \begin{bmatrix} 0 & 1 & \mathbf{0} \\ 0 & 0 & \mathbf{0} \\ 2\pi\hat{\mathbf{B}} & \mathbf{0} & \hat{\mathbf{A}} \end{bmatrix} \begin{bmatrix} \alpha_{\text{eff}} \\ \dot{\alpha} \\ \hat{\mathbf{x}} \end{bmatrix} + \begin{bmatrix} -\frac{1}{U} & \frac{1}{U}(c/4 - d) \\ 0 & 1 \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \ddot{h} \\ \ddot{\alpha} \end{bmatrix}, \quad (1)$$

$$c_L = [2\pi\hat{\mathbf{D}} \quad \frac{\pi c}{2U} \quad \hat{\mathbf{C}}] \begin{bmatrix} \alpha_{\text{eff}} \\ \dot{\alpha} \\ \hat{\mathbf{x}} \end{bmatrix} + \begin{bmatrix} -\frac{\pi c}{2U^2} & -\frac{\pi cd}{2U^2} \end{bmatrix} \begin{bmatrix} \ddot{h} \\ \ddot{\alpha} \end{bmatrix}. \quad (2)$$

Here, $\alpha_{\text{eff}} = -\dot{h}/U + \alpha + (c/4 - d)\dot{\alpha}/U$ is the effective angle of attack and $(\hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}, \hat{\mathbf{D}})$ represents the state-space system that models the wake dynamics using the two-element state vector $\hat{\mathbf{x}}$. We will only analyze cases where the pivot point is the midchord point, so $d = 0$. We note that, in this configuration, the angular acceleration has no direct contribution to the lift and the lift value could in principle inform the action value at the same time step. However, for generalizability purposes

and to create a more challenging and realistic control problem, we let the agent observe the lift with a time delay τ equal to the agent-interaction time step.

A. Setup

From now on, we nondimensionalize variables using U , c , and ρ , and the reader should assume that any variable is dimensionless. The time step τ between every agent-environment interaction corresponds to 0.1 convective times and we simulate episodes of at most 200 time steps with the discrete-time form of the state-space system (1) and (2). For later purposes, we define a scale $\check{h}_{\text{scale}} = 0.01$ for the disturbance and $c_{L,\text{scale}} = 0.02$ for the lift coefficient. We use a value $\check{\alpha}_{\text{max}} = 0.1$ to limit the control input to the bounds $[-\check{\alpha}_{\text{max}}, \check{\alpha}_{\text{max}}]$, which is large enough to avoid saturation while achieving the control objective, except for during the early stages of the training. However, if the action space were smaller, the nonlinear nature of deep reinforcement learning is expected to deal adequately with the nonlinear saturation, which might not always be possible with linear control theory.

We study the response to two types of \check{h} variations, which we will label respectively as impulsive and smooth and which do not occur simultaneously in the same episode. The value of \check{h} over the period between two action time steps is always piecewise constant. The smooth disturbances are constructed from a Fourier sine series containing modes 5 through 50, with 2.5 to 25 cycles in one episode, respectively. The n th mode has an amplitude D/n with D sampled from a standard normal distribution for each mode. This choice ensures that the variations in \check{h} are neither too slow, which would lead to high $\dot{\check{h}}$, nor too fast for the controller to react. The signal is then scaled such that the maximum amplitude is equal to \check{h}_{scale} . The impulsive disturbances consist of a random number, uniformly sampled from $[0,20]$, of events, occurring at unique random discrete times between time steps 0 and 199. An event has equal likelihood of being either a step change or a discrete (finite-amplitude) impulse. The amplitude of each event is randomly chosen from a uniform distribution with a zero lower bound and an upper bound that is limited such that $|\dot{\check{h}}| < 0.01$ and $|\check{h}| < 0.7\check{h}_{\text{scale}}$, chosen to approximately match the uncontrolled c_L magnitudes from the smooth disturbances.

During training, we terminate the episode early if $|c_L| > c_{L,\text{scale}}$ to expedite the learning process; this reduces the computational time spent on simulating high-lift states that are reached after poor pitch control. We do not restrict $\check{\alpha}$ or α , although this could also be used to optimize the training, as we will do in the viscous problem in the next section. To train the agent to minimize the lift variations, the reward after every time step consists of

$$R_{k+1} = 1 - \frac{|c_{L,k}|}{c_{L,\text{scale}}}. \quad (3)$$

The value 1 assigns a reward for continuing successfully during that time step within the allowed threshold; the negative part penalizes a lift that is not zero and this penalty increases linearly as the lift magnitude increases. Even though the positive part of the reward is not required to learn to minimize lift variations, we found that the training improved if we included it. With a maximum episode length of 200 time steps, the maximum possible episode return is 200, occurring when the lift is zero throughout the episode.

In a full viscous environment, even if it were possible to observe the entire flow state, it would be intractable to train an agent to act on that observation except with some specialized policy architectures. The classical unsteady aerodynamics environment does not suffer from this same limitation because it can be accurately modeled using only two states that represent the wake. It, therefore, allows us to try out and compare the reinforcement learning framework with three types of agents that observe different degrees of the full flow state. All observe the previous lift $c_{L,k-1}$ and the current angular states α_k and $\dot{\alpha}_k$ but observe different amounts of information about the wake: one that observes no wake information, one that observes pressure information at the airfoil, and one that observes all the system states of the state-space system (1), including the wake states (which we will refer to as “full wake”). For a fair comparison, all the observed values, except

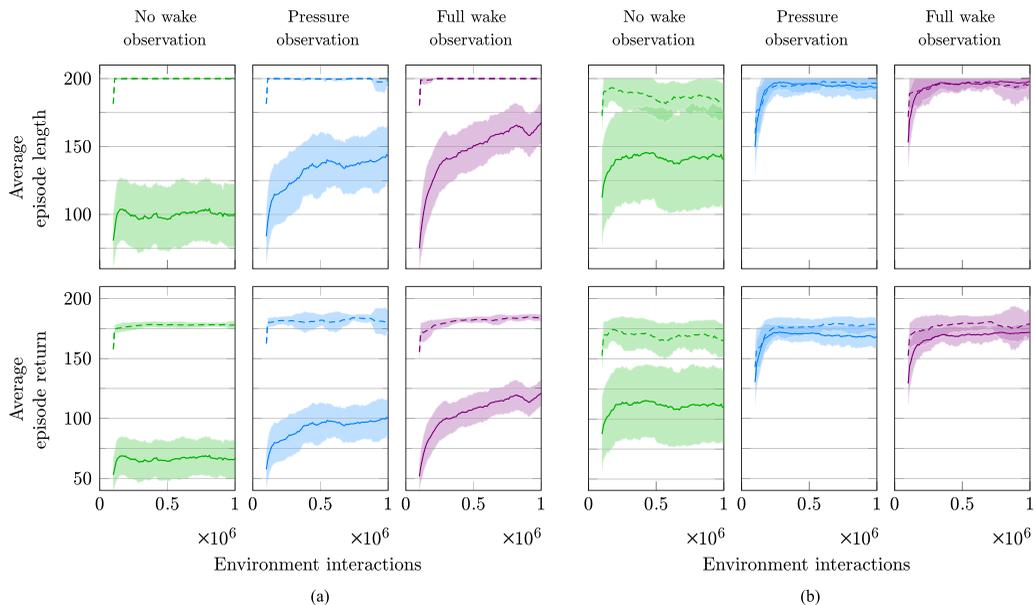


FIG. 5. Mean and standard deviation (shade) of the average episode length and return for agents observing no wake information (green), observing pressure (blue), and observing the full wake (violet), with access to the last one (solid lines) or two (dashed lines) observations. Panels (a) and (b) correspond to agents trained and evaluated on episodes with randomized smooth and impulsive \tilde{h} variations, respectively.

for α and $\dot{\alpha}$, are subject to the same output delay τ . The pressure-observing agent observes the pressure at two points on the plate. This pressure is the sum of the added mass pressure and the circulatory pressure, computed using Eqs. (B4) and (B5) from Appendix B. It follows from those equations that sensor measurements from two unique positions anywhere along the plate allows the agent to distinguish the two contributions and that any additional pressure observations do not introduce additional independent information. To study the effect of memory on the agent’s performance in this POMDP, we will compare two variants of the previous three types of agents: one that only uses the most recent observation and one that retains a memory of the two most recent observations.

B. Training and evaluation

To mitigate variance stemming from initial random actions and policy exploration, we train each agent variant on both types of disturbance across ten environments initialized with random seeds, resulting in ten policies for each agent. Each policy is trained with one million environment interactions. We evaluate the agents at checkpoints that occur after every 10 000 interactions of training. To evaluate an agent, we simulate ten random episodes of one disturbance type with each of its ten policies without applying action noise. We average the 100 ($= 10 \times 10$) episode lengths and returns for each evaluation checkpoint and then create a smoothed *training curve* by applying a moving average over ten evaluation checkpoints. These training curves are shown for each possible combination of agent and disturbance type in Fig. 5 to compare the training progression.

As anticipated, both the extra wake information and memory help to partially restore the Markovian property and increase the average episode return, with the memory having the largest effect. In this simplified flow problem, the pressure provides sufficient additional information to approximately match the performance of an agent that has knowledge of the entire flow at the previous time step, especially when the agent has a memory of multiple observations, in which case

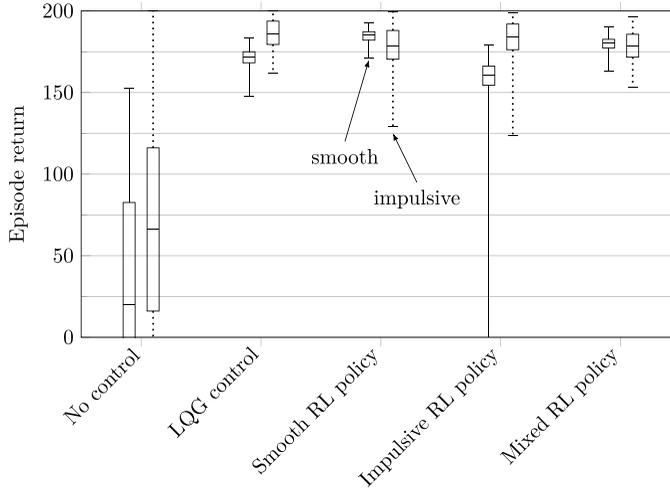


FIG. 6. Comparison of the episode return for different agents for 1000 random episodes each of smooth (solid) and impulsive (dotted) \tilde{h} variations in the classical unsteady aerodynamics environment. The lower and upper whiskers represent the minimum and maximum values of the sampled episodes, respectively.

their performances become indistinguishable for the chosen environment parameters. This confirms that the learning algorithm is able to exploit the pressure information for improved performance, which we will seek to replicate in a viscous environment in the next section, when full-wake observations are unavailable. Additionally, we note that the difference in performance between the agents that do retain memory and those that do not is much smaller for the impulsive disturbances than it is for the smooth disturbances, indicating that there is a fundamental difference for the RL training between both types of disturbances.

As a benchmark for the performance of the RL policies, we tune a discrete linear-quadratic-Gaussian (LQG) controller for the discrete-time system of (1) and (2) using output feedback, subject to the same output delay and assuming no measurement noise, by employing the automated LQG synthesis tuning feature of MATLAB's Control System Designer application [35], resulting in a z -domain transfer function

$$K(z) = \frac{1.1836(z + 0.01069)(z^2 - 1.826z + 0.8353)}{(z + 1.237)(z - 1)(z - 0.8592)}. \quad (4)$$

Figure 6 compares the cases with no control, LQG control, and RL control using boxplots that indicate the minimum, first quartile, median, third quartile, and maximum episode returns for a set of episodes. Each boxplot represents the data of 1000 episodes with randomized disturbances of one type. For the RL cases, these 1000 episodes are collected by simulating 100 episodes with each of the ten policies of the agents that observe the full wake with memory. The figure indicates the ability of the RL control to match the performance of this specific LQG controller for impulsive disturbances and exceed it for smooth disturbances. This is also apparent from the plots in Fig. 7 of the lift variation during an actual episode with those disturbances. Interestingly, the impulsive episode results in Fig. 7 indicate that the RL control learns not to react too strongly to impulsive changes in favor of a longer-term response.

It is important to stress that there are likely to be linear controllers that achieve better and more robust performance than the one devised here for comparison. The main advantage of the RL control here is the ability to treat the plant as a black box. It is also important to acknowledge that the disturbances in this work are functions with random parameters, meaning their values are not stochastically independent from one time step to the next. Although this should ideally be reflected by additional states in the environment's state vector, we have chosen not to model this aspect in our

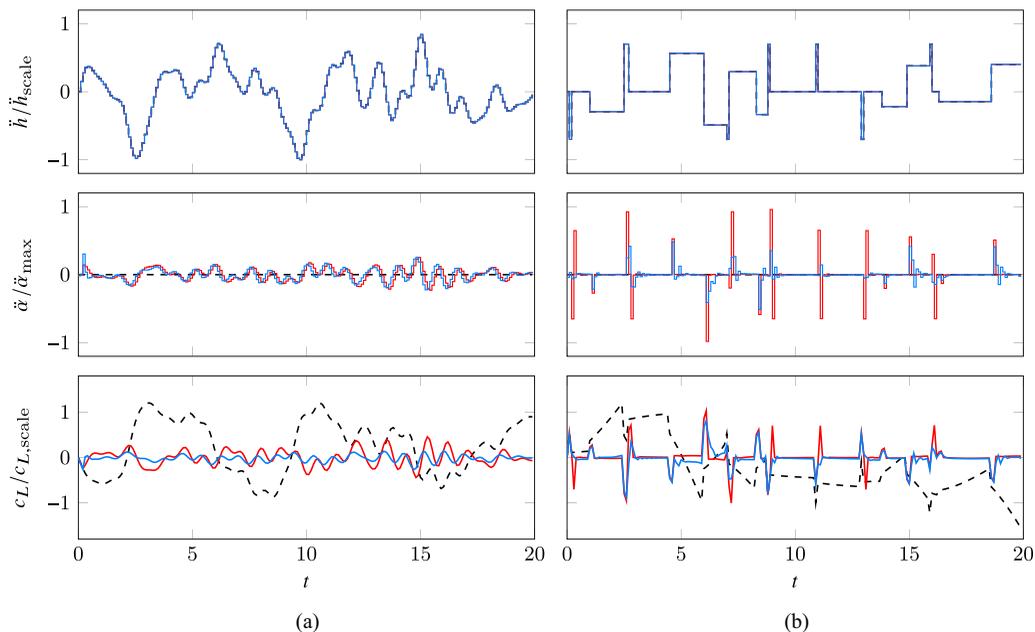


FIG. 7. Vertical acceleration (disturbance), angular acceleration (control input), and lift coefficient for a single episode with (a) smooth disturbances and (b) impulsive disturbances using different controllers: no control (---), LQG control (—), and the best available RL agent (—).

current work, as it remains internal to the environment. Nevertheless, an effective RL agent could potentially exploit this behavior, as might be the case here if the policy trained on one disturbance type learned to anticipate that disturbance’s behavior over the next few time steps.

Lastly, Fig. 6 indicates that, while the RL agents perform well on the disturbances they were trained on, this behavior does not automatically transfer to the other type of disturbance. In one case, this evaluation led to one of the policies giving unstable results, indicated by the extreme value of its boxplot’s lower whisker. This points to the important fact that the nonlinear controller at the heart of the trained deep RL agent is not guaranteed to generalize in the same way that robust linear controllers do. Therefore, care has to be taken when choosing the training data for a controller that should be robust to all kinds of disturbances. In line with this guidance, we trained ten additional policies in an environment with mixed episodes, meaning that the disturbance type for each episode is randomly chosen to be either smooth or impulsive. Their performance on evaluation episodes of one type is slightly lower than the policies trained on exclusively that type of disturbance. However, their worst performance is considerably better than that of the policies trained on one type and evaluated on the other type.

IV. LEARNING IN A VISCOUS FLOW ENVIRONMENT

In this example, we apply airfoil pitch control in an incompressible viscous flow environment at $Re = 200$, which poses a greater challenge to the agent due to its high dimensionality and nonlinear dynamics. In particular, we study if pressure measurements can improve the agent’s performance as it did in a classical unsteady aerodynamics environment. The airfoil is again a flat plate that pitches about the midchord point. Instead of disturbances in the form of \hat{h} variations, we focus on local flow disturbances upstream of the airfoil introduced by a localized, transient forcing, resulting in vortex dipoles that move past the airfoil. This environment is illustrated in the bottom panel of Fig. 1.

A. Setup

The time step between every agent-environment interaction is reduced to 0.03 convective times and we simulate episodes of at most 100 time steps. The flow disturbance is a forcing field that varies spatially and temporally in a Gaussian manner and is described in the wind tunnel reference frame coordinates as

$$\mathbf{F}(x, y, t) = \rho U^2 c \frac{(D_x, D_y)}{\pi^{3/2} \sigma_x \sigma_y \sigma_t} \exp\left(-\frac{(x-x_0)^2}{\sigma_x^2} - \frac{(y-y_0)^2}{\sigma_y^2} - \frac{(t-t_0)^2}{\sigma_t^2}\right), \quad (5)$$

where $t_0 = 0.5$, $\sigma_t = 0.2$, and $x_0 = -0.5$; thus the forcing is applied at a constant horizontal distance upstream of the airfoil. The other parameters are sampled from uniform continuous distributions as follows. The spatial widths σ_x and σ_y are chosen from the interval $[0.04, 0.1]$ and the vertical position y_0 from the interval $[-0.25, 0.25]$. Since the temporal and spatial Gaussian widths are relatively small, the forcing appears as pulsed point forces, which is how we will refer to them in the rest of this work. Such a force produces a vortex dipole, which self-propels in the direction in which the force was applied (in addition to being convected by the background flow). The horizontal amplitude D_x is sampled from the interval $[0, 3]$ and the vertical amplitude D_y from the interval $[-3(y_0 - 0.1)/0.35, -3(y_0 + 0.1)/0.35]$. These values are chosen so that the vortex dipoles are aimed towards the airfoil and are strong enough to induce a nonlinear lift response. We will also investigate the response to two or three of these point force pulses in the same episode, in which case the forcing field is a superposition of multiple instances of (5), using independently sampled parameters. To spread the pulses in time, t_0 of a new pulse is delayed with respect to that of the previous pulse with a delay uniformly sampled from an interval $[0, 1.5]$; the first pulse remains at $t_0 = 0.5$.

The flow simulation is performed using a viscous incompressible flow solver employing an immersed boundary projection method [36] in a body fixed reference frame with a grid cell size $\Delta x = 3c/\text{Re}$ and a time step size $\Delta t = 0.4\Delta x/U_\infty$ on a domain that spans $[-1.25c, 1.5c] \times [-0.5c, 0.5c]$ with free-space boundary conditions. The flow solver is vorticity-based and utilizes a lattice Green’s function for solutions of the Poisson equation, enabling a more compact domain than typical computational aerodynamics studies. The vortex disturbances remain between the upper and lower boundaries and freely convect through the downstream boundary. The flow solver time step in this case equals 0.006 convective times, but we set the agent-environment interaction time step to 0.03 convective times. This limits the controller’s bandwidth, creating a generally more challenging control problem, but also affects the temporal information available in the observations. The latter could positively or adversely affect the RL training depending on the characteristic timescale of the fluid problem and the number of previous observations retained in the agent’s memory, which we do not explore in this work.

We again define a scale $c_{L,\text{scale}} = 0.1$ and a maximum control input magnitude $\ddot{\alpha}_{\text{max}} = 10$. In addition, we also impose the constraints $|\alpha| < 60\pi/180$ to avoid unrealistic configurations and $|\dot{\alpha}| < 1.8$ to maintain a reasonable viscous flow solver time step and we always terminate the episode if these are violated. During training only (and not when evaluating a policy), we terminate an episode if $|c_L|$ exceeds 0.8.

We modify the reward function from the previous section with two additions. We add a negative reward for high variations of $\ddot{\alpha}$ between two time steps to prevent rapid oscillatory behavior and another negative reward when any of the constraints are violated:

$$R_{k+1} = 1 - \frac{|c_{L,k}|}{c_{L,\text{scale}}} - 2 \frac{|\ddot{\alpha}_k - \ddot{\alpha}_{k-1}|}{\ddot{\alpha}_{\text{max}}} - \begin{cases} 100, & S_k \notin \mathcal{S}, \\ 0, & \text{otherwise,} \end{cases} \quad (6)$$

where \mathcal{S} is the space of valid values of the environment’s state, when α , $\dot{\alpha}$, and c_L are within their prescribed ranges. With a maximum episode length of 100 time steps, the maximum possible episode return is 100, again occurring when the lift is zero throughout the episode.

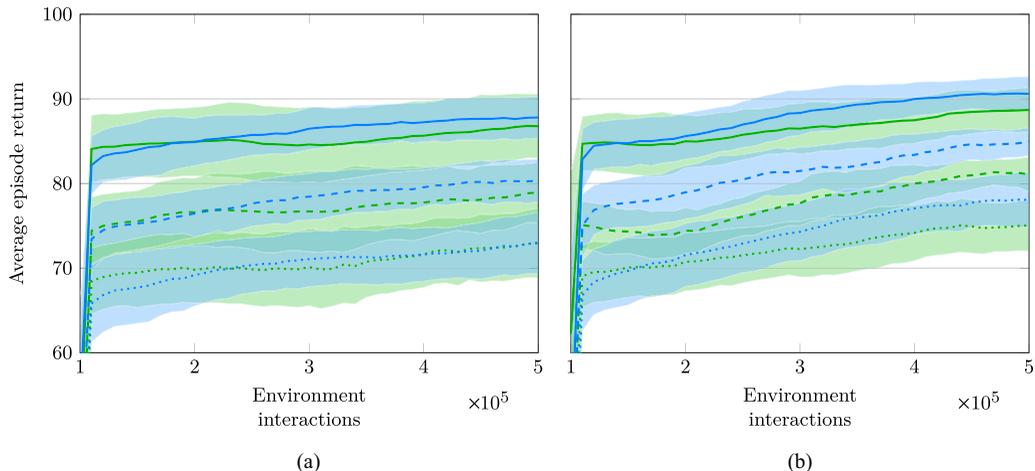


FIG. 8. Mean and standard deviation (shade) of the average episode return for agents observing pressure (blue) and not observing pressure (green), evaluated on episodes consisting of one (solid lines), two (dashed lines), and three (dotted lines) random point force pulses. Panels (a) and (b) correspond to agents trained on episodes with one and two random point force pulses, respectively.

While it would be possible in a simulation to observe the entire vorticity field (and thus the state of the flow) and feed this to a convolutional layer in the policy network, no good results could be obtained with this approach with the TD3 algorithm in our setup. We limit ourselves, therefore, to the following two types of agents: one that observes the current $\alpha_k, \dot{\alpha}_k$, and the latest c_L computed by the flow solver (with an output delay of one flow solver time step) and one that additionally observes the latest pressure jump over the plate at seven positions spread equidistantly between 30% and 70% chord. Both types of agents are implemented with a memory of the last four observation vectors, stacked together to form the input to their policy. We found that increasing the memory to retain more than four observations did not significantly increase the performance of the controller for the size of the neural network used in this work. Lastly, even though the previous action $\ddot{\alpha}_{k-1}$ appears in the reward function, we found that including its value in the observations did not increase the performance.

B. Training and evaluation

We train the policies of the two types of agents over 500 000 interactions on two types of training episodes: one type consisting of a single force pulse disturbance and one consisting of two force pulse disturbances. Each force pulse has new random parameters for every new episode. Each case is run with ten environments initialized with random seeds, resulting in ten policies for each of the four possible combinations of agents and training data. To evaluate the policies and study their generalization to unseen cases, we also use a third type of episode, consisting of three random force pulses, which results in higher disturbed flows with increased nonlinear vortex interactions.

Figure 8 compares the training curves of the different agents, using the same evaluation and averaging procedures as in the classical unsteady aerodynamics environment. The average episode length during training is not shown but rapidly increases to its maximum possible value after approximately 20 000 interactions. This indicates that agents quickly learn to avoid the large negative rewards associated with the early episode termination when permissible bounds are violated. After this initial stage, the reward gradually increases as more experience is obtained and processed in the training.

Training an agent with two-pulse episodes instead of one-pulse episodes provides richer training data, including experiences with disturbances that interact with each other. The training curves show

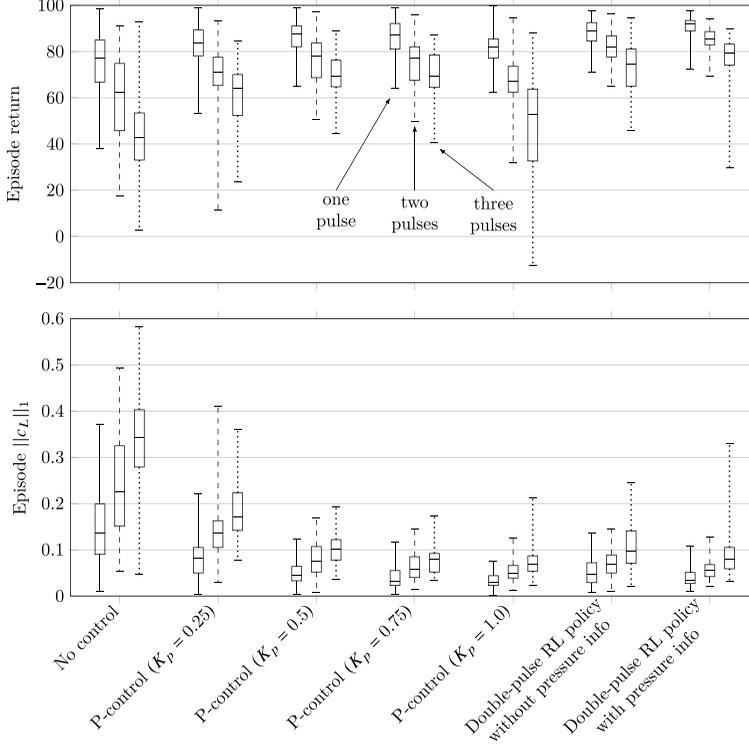


FIG. 9. Comparison of the episode return and lift norm for different agents for 100 random episodes each of one (solid), two (dashed), and three (dotted) point force pulses. The lower and upper whisker represent the minimum and maximum values of the sampled episodes, respectively.

that this only has a small effect on the performance of the agents that do not observe pressure. However, it significantly increases the performance of the pressure-based agents, which then, on average, consistently outperform the agents that do not observe pressure. This increased performance also results in a higher average episode return for single-pulse episodes compared to the agents trained exclusively on single-pulse episodes, for the number of interactions in this training.

To evaluate the performance of the agents after the training, we compare them to proportional control ($\ddot{\alpha}/\ddot{\alpha}_{\max} = -K_p c_L/c_{L,\text{scale}}$, subject to the same output delay), first using episodes with randomized force pulses and then with two deterministic episodes. The results of the former are presented in Fig. 9, which displays (as boxplots) the distributions of the episode return and the 1-norm of the lift coefficient over 100 episodes of the three kinds of random evaluation episodes. For the RL agents, these 100 episodes are divided equally over their ten policies. For the comparison with deterministic episodes, we define a representative episode, consisting of two pulses, each with $y_0 = -0.125$ and all their other parameters set at the average value from their distributions, visualized on the left side of Fig. 10. We also define an extreme episode with three pulses, visualized on the right side of the figure, which creates a much higher disturbed flow than any of the training episodes and can be used to assess how well an agent generalizes to such a flow. These pulses appear at t_0 equal to 0.5, 0.8, and 1.1 and y_0 equal to -0.25 , 0.25 , and -0.25 , respectively. Their amplitudes are set to the strongest values from their distributions and their spatial widths to the lowest value of 0.04. Figure 11 displays the values of the control input, pitch, and lift coefficient during these two episodes, while Table I lists the return, the 1-norms of the lift coefficient and control input, and the average value of the power coefficient $c_P = 2m\dot{\alpha}/(\rho U^3 c)$, where m is the moment about the midchord point.

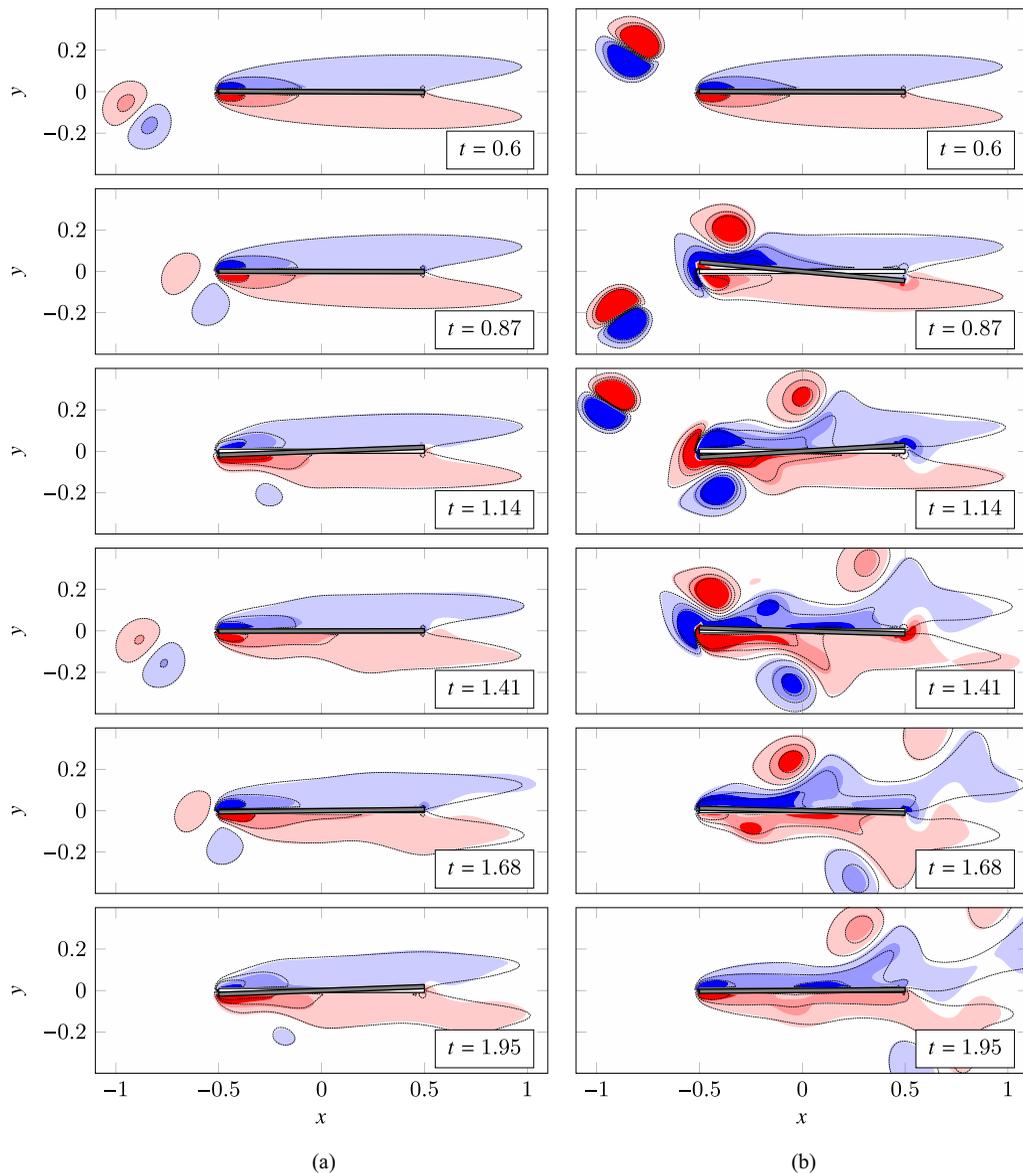


FIG. 10. Snapshots of the vorticity field contours for the best available RL agent (colored contours) and without control (.....) at different time instances of (a) a representative episode with two average point force pulses and (b) an extreme episode with three strong force pulses. The uncontrolled airfoil is displayed in white and the controlled airfoil in dark gray. The vorticity contours are drawn in both cases for the values $-200, -15, -9, -3, 3, 9, 15, 200$.

For the proportional control, Fig. 9 indicates that a higher gain results in higher episode returns and lower lift norms. However, once the gain is approximately one or higher, the control input displays rapid oscillations (visible in Fig. 11) that are eventually stabilized by the control input bounds. This results in lower episode returns because the reward penalizes fast control input variations. Additionally, we found that proportional control augmented with integral and derivative control (PID control) performs only marginally better than proportional control when averaged over

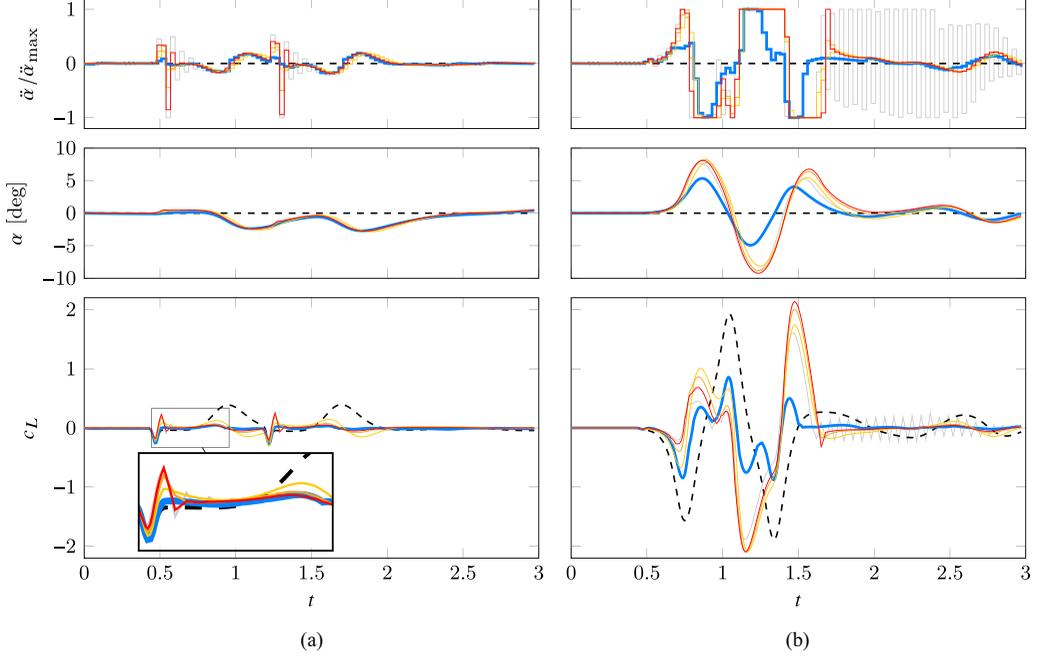


FIG. 11. Angular acceleration, angle of attack, and lift coefficient for (a) a representative episode with two average point force pulses and (b) an extreme episode with three strong force pulses using different controllers: no control (- - -), proportional control with $K_p = 0.25$ (—), 0.5 (—), 0.75 (—), and 1.0 (—), and the best available RL agent (—).

randomized cases. For cases with stronger gusts, like the extreme case defined above, PID control could sometimes perform better than proportional control, depending on the exact strength and directions of the gusts, but requires retuning for each case. For these reasons, we did not include PID control in the results.

For the RL control, Fig. 9 and Table I illustrate that the episode returns for the random and representative deterministic episodes are on average considerably higher than those achieved without control or with proportional control. However, the lift norm is approximately the same magnitude,

TABLE I. Comparison of different agents for a representative episode with two average point force pulses and an extreme episode with three strong point force pulses.

		No control	P-control ($K_p = 0.25$)	P-control ($K_p = 0.5$)	P-control ($K_p = 0.75$)	P-control ($K_p = 1.0$)	Double-pulse RL pressure info policy without	Double-pulse RL pressure info policy with
Return		58.8	70.6	76.5	73.8	63.1	79.4 ± 3.4	84.5 ± 1.3
Two average force pulses	$\ c_L\ _1$	0.247	0.144	0.082	0.062	0.057	0.077 ± 0.021	0.056 ± 0.008
	$\ \ddot{\alpha}_k\ _1$	0.00	3.64	4.22	4.96	6.20	3.97 ± 0.24	3.50 ± 0.27
	$\overline{c_P} (\times 10^{-3})$	0.00	4.15	4.11	4.26	4.43	3.56 ± 0.14	3.53 ± 0.16
Return		-77.2	-90.9	-82.8	-79.9	-209.0	0.1 ± 10.5	-26.3 ± 11.1
Three strong force pulses	$\ c_L\ _1$	1.063	1.032	0.965	0.925	0.787	0.489 ± 0.055	0.637 ± 0.067
	$\ \ddot{\alpha}_k\ _1$	0.00	18.65	20.60	21.34	38.27	14.42 ± 1.33	16.04 ± 1.01
	$\overline{c_P} (\times 10^{-3})$	0.00	102.52	116.04	118.22	115.14	74.21 ± 6.34	86.71 ± 6.71

indicating that RL can achieve similar lift mitigation performance with less control input variation. On the other hand, the extreme deterministic episode shows the superior performance of the RL control compared to the proportional control in lift mitigation and control expenditure. Figure 10 shows that the strong vortex dipoles in this episode interact significantly with the leading edge of the plate, the wake of the airfoil, and the previous vortex dipoles. Such strong flow interactions exhibit substantial nonlinear behavior, further compounded by their nonlinear dependence on the pitch variations. This complexity severely compromises the effectiveness of proportional control, leading to inferior performance compared to the uncontrolled scenario. In contrast, the RL control performs markedly better than the uncontrolled case, underscoring its capability of controlling the nonlinear dynamics. For both deterministic episodes, Table I shows that the RL control requires less power for better lift mitigation, even though the reward function does not consider energy expenditure.

Surprisingly, the data presented in Table I shows that the performance of the agent that does not observe pressure during the extreme episode is better than that of the pressure-observing agent. Similarly, in Fig. 9, the increased range of its boxplot for the three-pulse episodes suggests poorer performance by the pressure-observing agent, despite its higher median performance. This phenomenon may be attributed to a faster convergence during training for the agent that does not observe pressure, as opposed to the pressure-observing agent, whose training may not have fully converged. This discrepancy could primarily affect performance during generalization and extreme cases.

When analyzing these results, one should remember that the RL agents are conditioned through the reward function to find a trade-off between low lift magnitude and small control input variations. This trade-off is further affected by a conservative policy near state thresholds that would result in high negative rewards if exceeded. Designing reward functions with different weights for these contributions (or including different terms) can significantly impact the policy, which we do not explore in this work.

V. CONCLUSIONS

In this work, we studied the use of deep reinforcement learning and the role of partial observability in designing an airfoil pitch controller to minimize lift variations in randomly disturbed flows. In most scenarios, a practical controller would only be able to sense the flow through force and pressure sensors. This information is often not enough to observe the entire state of the flow, which, together with the disturbances, determines the exact lift response of the airfoil. From the perspective of reinforcement learning, the pitch controller is an agent that receives non-Markovian observations from the environment and the control problem can be viewed as a partially observable Markov decision process (POMDP). Deep reinforcement learning has proven to be a viable approach for approximating an optimal control policy for a given reward function in such conditions. We specifically apply deep RL using the TD3 algorithm to learn pitch control of a flat plate and study agents with and without a memory of earlier observations and with different degrees of environment observations.

We tested our reinforcement learning setup in two types of environments. The first environment modeled the flow and disturbances in a Wagner/Theodorsen setup: the flow is a potential flow with a flat vortex sheet wake and the disturbances are random vertical accelerations of the plate. When analyzing the results in this classical unsteady aerodynamics environment, we confirmed that (1) augmenting the observed information with pressure and wake states or (2) retaining a memory of the two most recent observations both increase the performance of our RL control. This was expected because both aim to restore the Markovian property of the POMDP. We found that our RL control can generally match or exceed the performance of an LQG controller when minimizing the lift variations. Both the RL control and LQG can be further fine-tuned for specific performance metrics, which we did not consider in this work. Finally, we found that training on smooth disturbances does not guarantee adequate performance on impulsive disturbances and vice versa,

even though the underlying problem is linear. This can be explained by the nonlinearity of the RL policy and should be taken into account when designing a robust controller.

The second environment was a viscous flow model where flow disturbances are introduced by pulsed point forces upstream of the airfoil. When analyzing the results in this environment, we found that the RL setup can again take advantage of pressure information in the observations to obtain better policies for lift mitigation. The performance of the RL control improves when two force pulse disturbances instead of one are included in each training episode, in which case the RL control also generalizes better to episodes with three force pulse disturbances. This can be explained by the fact that these two-pulse episodes contain much richer training examples that include the interactions between disturbances. Experience with those interactions is valuable also when acting in situations with more than two disturbances. When comparing the performance with proportional control, we found that the RL control can achieve similar performance with less control input variation for moderate episodes and significantly outperform the proportional control for an episode with very strong disturbances, in which case the flow behavior is strongly nonlinear.

The problem setup in this work can be further explored in multiple ways. First, all of the examples in this work can straightforwardly be extended to the servo problem where the lift has to track a nonzero reference value. Secondly, the two-dimensional viscous simulations in this work are performed at a relatively low Reynolds number. Environments with a higher Reynolds number or three-dimensional flows could become too computationally expensive in which to train a TD3 agent. In these cases, a more sample-efficient algorithm, such as a model-based one, could be a more viable RL approach. Finally, while the exact specification of the reward function has a critical effect on the performance of the RL control, we have not yet explored its connection with the optimality or robustness of the trained control policies.

All the computational results can be reproduced with the code available at [37] and [38].

ACKNOWLEDGMENTS

Support for this work by the National Science Foundation under Award No. 2247005 is gratefully acknowledged.

APPENDIX A: REINFORCEMENT LEARNING ALGORITHM DETAILS

The reinforcement learning implementation used in this work is the Stable Baselines3 [39] implementation of the twin-delayed deep deterministic policy gradient algorithm (TD3) [30]. The TD3 algorithm utilizes six neural networks: one policy, or actor, networks, two state-action value (Q-function) approximation, or critic, network, and their target networks. These networks all have the same feedforward architecture: an input layer whose size depends on the observations, two fully connected hidden layers with 400 and 300 rectified linear units (ReLU), respectively, and an output layer representing the action in the case of an actor network or the state-action value in the case of a critic network. The target networks are copies of the actor and critic networks with weights that lag behind their counterparts' weights through polyak averaging as $\phi_{\text{target}} = 0.005\phi_{\text{target}} + 0.995\phi$. Their purpose is to provide a fixed target when updating their counterpart networks, which prevents training instabilities due to the recursiveness of the underlying Bellman equation.

The training algorithm stores agent-environment interactions in a standard first-in-first-out replay buffer, which can store up to 1 000 000 interactions. Each interaction consists of the current observation, the chosen action, the resulting subsequent observation, and the corresponding reward. Once 10 000 interactions have been generated with a random policy, the training starts and the algorithm switches repeatedly between performing a training step and simulating a new episode with the latest trained policy to further update the replay buffer. During each training step, the networks are updated multiple times using stochastic gradient descent combined with the Adam optimization algorithm using a learning rate of 0.001. Each stochastic gradient descent update uses a batch of

256 randomly sampled interactions from the replay buffer. The critic networks and their targets are updated as many times as there were steps in the last generated episode and the actor network and its target are updated during every second update of the critic networks using the same batch of interactions. The training uses a discount factor of 0.98 for future rewards and applies target policy smoothing using zero-mean Gaussian noise with a standard deviation of 0.1. This noise is added to the action generated by the target policy network and regularizes the Q-value approximations of the critic networks. Lastly, the algorithm adds zero-mean Gaussian noise with a standard deviation of 0.1 to the actions of the policy when collecting training data to promote the exploration of the state and action spaces.

APPENDIX B: CLASSICAL UNSTEADY AERODYNAMICS MODEL

Wagner [33] and Theodorsen [34] analyzed the aerodynamic response of an infinitesimally thin flat plate airfoil to specific flow disturbances under the following assumptions. The flat plate has chord length c and is at a small angle of attack α , where $|\alpha| \ll 1$. Its main motion is a constant velocity U in the negative x direction. Any other motions are perturbations of the angle of attack with an angular velocity $\dot{\alpha}$ applied at a pivot point at a coordinate $(d, 0)$ in the body-fixed reference frame positioned at the center of the airfoil with its \tilde{x} axis along the chord and the vertical motion \dot{h} of that pivot point. These perturbations are considered small compared to the forward velocity of the plate, i.e., $|\dot{\alpha}| \ll U/c$ and $|\dot{h}| \ll U$. The wake is assumed to lie along the \tilde{x} axis of the body and the wake vorticity convects with the velocity U away from the airfoil.

Classical unsteady aerodynamics theory [40] provides an expression for the bound vortex sheet strength when the Kutta condition is applied at the trailing edge under the previously mentioned assumptions:

$$\int_{c/2}^{\infty} \frac{(\tilde{x} + c/2)^{1/2}}{(\tilde{x} - c/2)^{1/2}} \gamma_S(\tilde{x}) d\tilde{x} = -\Gamma^{\text{qs}}, \quad (\text{B1})$$

where the quasisteady circulation $\Gamma^{\text{qs}} := -\pi c U \alpha_{\text{eff}}$ depends on the effective angle of attack

$$\alpha_{\text{eff}} := -\frac{\dot{h}}{U} + \alpha + \left(\frac{c}{4} - d\right) \frac{\dot{\alpha}}{U}. \quad (\text{B2})$$

Assuming that the wake convects with the velocity U and Eq. (B1) has to be satisfied at any instant of time, the wake strength can be determined. Furthermore, the classical theory also provides a general expression for the lift on the plate:

$$f_y = \underbrace{-\frac{\pi}{4} \rho c^2 (\ddot{h} + d\ddot{\alpha} - U\dot{\alpha})}_{\text{added-mass lift } f_y^{\text{am}}} - \underbrace{\rho U \Gamma^{\text{qs}}}_{\text{quasisteady lift } f_y^{\text{qs}}} - \underbrace{\frac{1}{2} \rho U c \int_{c/2}^{\infty} \frac{\gamma_S(\tilde{x}_S)}{(\tilde{x}_S^2 - c^2/4)^{1/2}} d\tilde{x}_S}_{\text{wake lift } f_y^{\text{S}}}. \quad (\text{B3})$$

The last two terms (the quasisteady and wake lift) together form the circulatory lift, or $f_y^{\text{circ}} = f_y^{\text{qs}} + f_y^{\text{S}}$.

The distributions of the pressure jump across the plate corresponding to the added-mass and circulatory lift were first derived by Neumark [41] and are for rigid-body motion given by

$$[p^{\text{am}}]_{\pm}^+(\tilde{x}_P) = -2\rho(\ddot{h} + d\ddot{\alpha} - U\dot{\alpha})(c^2/4 - \tilde{x}_P^2)^{1/2}, \quad (\text{B4})$$

$$[p^{\text{circ}}]_{\pm}^+(\tilde{x}_P) = \left(-2\rho U^2 \alpha_{\text{eff}} + \frac{\rho U}{\pi} \int_{c/2}^{\infty} \frac{\gamma_S(\tilde{x}_S)}{(\tilde{x}_S^2 - c^2/4)^{1/2}} d\tilde{x}_S \right) \frac{(c/2 - \tilde{x}_P)^{1/2}}{(c/2 + \tilde{x}_P)^{1/2}}. \quad (\text{B5})$$

The full pressure is then $[p]_{\pm}^+ = [p^{\text{am}}]_{\pm}^+ + [p^{\text{circ}}]_{\pm}^+$. Note that the first factor in parentheses in Eq. (B5) is equal to $-2f_y^{\text{circ}}/\pi c$ and thus the circulatory pressure at any point along the airfoil linearly scales with the circulatory lift.

For an arbitrary, smoothly varying α_{eff} , a more practical expression can be obtained for the circulatory lift if we formulate it as the convolution of the rate of change of Γ^{qs} with the nondimensional lift response to a step change in Γ^{qs} , known as Wagner's function Φ [33], plus the contribution of the initial value of Γ^{qs} times Wagner's function evaluated at the current time:

$$f_y = f_y^{\text{am}} - \rho U \Gamma^{\text{qs}}(0) \Phi(t) - \rho U \int_0^t \dot{\Gamma}^{\text{qs}}(\tau) \Phi(t - \tau) d\tau. \quad (\text{B6})$$

There is no analytical form of Wagner's function available, but several approximations exist, including one proposed by Jones [42] that is expressed as

$$\Phi(t) \approx 1 - 0.165e^{-0.091t} - 0.335e^{-0.6t}. \quad (\text{B7})$$

The circulatory lift in (B6) can be formulated in the Laplace domain as a multiplication of a transfer function with the quasisteady lift:

$$\mathcal{L}\{f_y^{\text{circ}}\} = -\rho U \Gamma^{\text{qs}}(0) \mathcal{L}\{\Phi\} - \rho U \mathcal{L}\{\dot{\Gamma}^{\text{qs}}\} \mathcal{L}\{\Phi\} \quad (\text{B8})$$

$$= -\rho U s \mathcal{L}\{\Phi\} \mathcal{L}\{\Gamma^{\text{qs}}\} \quad (\text{B9})$$

$$= s \mathcal{L}\{\Phi\} \mathcal{L}\{f_y^{\text{qs}}\}, \quad (\text{B10})$$

where s is the Laplace variable. The transfer function $C(s) = s \mathcal{L}\{\Phi\}$ based on the Wagner function approximation (B7) was also derived by Jones [42] and can be expressed as

$$C(s) = s \int_0^\infty \Phi(\tau) e^{-s\tau} d\tau \quad (\text{B11})$$

$$\approx \frac{0.5s^2 + 0.5616s + 0.0546}{s^2 + 0.691s + 0.0546}. \quad (\text{B12})$$

Furthermore, Jones [42] noted, like Garrick [43] before him, that C evaluated on the imaginary axis, i.e., its steady-state behavior for a purely oscillatory, quasisteady circulation, is equivalent to Theodorsen's function [34].

To design a controller or RL policy based on (partial) state feedback, we seek the state-space representation of (B11) and express the lift in terms of the lift coefficient $c_L = 2f_y/(\rho U^2 c)$. The controllable canonical state-space representation for $C(s)$ is

$$\begin{bmatrix} \dot{\hat{x}}_1 \\ \dot{\hat{x}}_2 \end{bmatrix} = \underbrace{\begin{bmatrix} -0.691 & -0.0546 \\ 1 & 0 \end{bmatrix}}_{\hat{\mathbf{A}}} \begin{bmatrix} \hat{x}_1 \\ \hat{x}_2 \end{bmatrix} + \underbrace{\begin{bmatrix} 1 \\ 0 \end{bmatrix}}_{\hat{\mathbf{B}}} c_L^{\text{qs}}, \quad (\text{B13})$$

$$c_L^{\text{circ}} = \underbrace{\begin{bmatrix} 0.2161 & 0.0273 \end{bmatrix}}_{\hat{\mathbf{C}}} \begin{bmatrix} \hat{x}_1 \\ \hat{x}_2 \end{bmatrix} + \underbrace{[0.5]}_{\hat{\mathbf{D}}} c_L^{\text{qs}}. \quad (\text{B14})$$

The response of this system to a unit step increase of c_L^{qs} is shown in Fig. 12.

Following Brunton and Rowley [6], we can use (B11) to create a multi-input, single-output transfer function for the complete lift expression by applying the Laplace transform to (B6) for each input:

$$\mathbf{G}(s) = [G_{\tilde{h}}(s) \quad G_{\tilde{\alpha}}(s)] \quad (\text{B15})$$

$$= \left[-\frac{\pi c}{2U^2} - \frac{2\pi}{sU} C(s) \quad \frac{\pi c}{2U} \left(\frac{1}{s} + \frac{d}{U} \right) + 2\pi \left(\frac{1}{s^2} + \frac{1}{sU} \left(\frac{c}{4} - d \right) \right) C(s) \right], \quad (\text{B16})$$

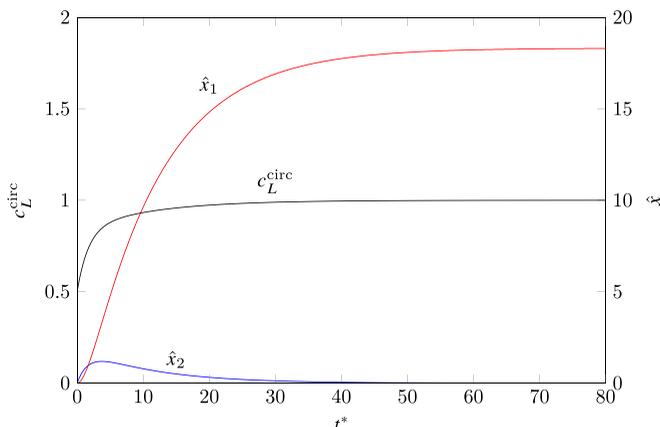


FIG. 12. Unit step response of the circulatory lift coefficient and the states of the controllable canonical state-space representation for the Wagner transfer function by Jones [42].

where $G_{\ddot{h}}$ and $G_{\ddot{\alpha}}$ are the transfer functions that give the response of the output c_L to the input signals \ddot{h} and $\ddot{\alpha}$, respectively. We can also cast this result into a minimal state-space realization:

$$\begin{bmatrix} \dot{\alpha}_{\text{eff}} \\ \dot{\alpha} \\ \dot{\hat{\mathbf{x}}} \end{bmatrix} = \underbrace{\begin{bmatrix} 0 & 1 & \mathbf{0} \\ 0 & 0 & \mathbf{0} \\ 2\pi\hat{\mathbf{B}} & \mathbf{0} & \hat{\mathbf{A}} \end{bmatrix}}_{\mathbf{A}} \begin{bmatrix} \alpha_{\text{eff}} \\ \alpha \\ \hat{\mathbf{x}} \end{bmatrix} + \underbrace{\begin{bmatrix} -\frac{1}{U} & \frac{1}{U}(\frac{c}{4} - d) \\ 0 & 1 \\ \mathbf{0} & \mathbf{0} \end{bmatrix}}_{\mathbf{B}} \begin{bmatrix} \ddot{h} \\ \ddot{\alpha} \end{bmatrix}, \quad (\text{B17})$$

$$c_L = \underbrace{\begin{bmatrix} 2\pi\hat{\mathbf{D}} & \frac{\pi c}{2U} & \hat{\mathbf{C}} \end{bmatrix}}_{\mathbf{C}} \begin{bmatrix} \alpha_{\text{eff}} \\ \alpha \\ \hat{\mathbf{x}} \end{bmatrix} + \underbrace{\begin{bmatrix} -\frac{\pi c}{2U^2} & -\frac{\pi cd}{2U^2} \end{bmatrix}}_{\mathbf{D}} \begin{bmatrix} \ddot{h} \\ \ddot{\alpha} \end{bmatrix}. \quad (\text{B18})$$

-
- [1] X. An, D. R. Williams, J. Eldredge, and T. Colonius, Modeling dynamic lift response to actuation, in *54th AIAA Aerospace Sciences Meeting* (AIAA, Reston, VA, 2016).
 - [2] A. R. Jones, Gust encounters of rigid wings: Taming the parameter space, *Phys. Rev. Fluids* **5**, 110513 (2020).
 - [3] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control: Analysis and Design* (John Wiley, Hoboken, US-NJ, 2005).
 - [4] S. Ahuja and C. W. Rowley, Feedback control of unstable steady states of flow past a flat plate using reduced-order estimators, *J. Fluid Mech.* **645**, 447 (2010).
 - [5] G. Sedky, A. Gementzopoulos, I. Andreu-Angulo, F. D. Lagor, and A. R. Jones, Physics of gust response mitigation in open-loop pitching manoeuvres, *J. Fluid Mech.* **944**, A38 (2022).
 - [6] S. L. Brunton and C. W. Rowley, Empirical state-space representations for Theodorsen's lift model, *J. Fluids Struct.* **38**, 174 (2013).
 - [7] S. L. Brunton, S. T. M. Dawson, and C. W. Rowley, State-space model identification and feedback control of unsteady aerodynamic forces, *J. Fluids Struct.* **50**, 253 (2014).
 - [8] G. Sedky, F. D. Lagor, and A. R. Jones, Unsteady aerodynamics of lift regulation during a transverse gust encounter, *Phys. Rev. Fluids* **5**, 074701 (2020).

- [9] W. Kerstens, J. Pfeiffer, D. R. Williams, R. King, and T. Colonius, Closed-loop control of lift for longitudinal gust suppression at low Reynolds numbers, *AIAA J.* **49**, 1721 (2011).
- [10] K. Bieker, S. Peitz, S. L. Brunton, J. N. Kutz, and M. Dellnitz, Deep model predictive flow control with limited sensor data and online learning, *Theor. Comput. Fluid Dyn.* **34**, 577 (2020).
- [11] X. Xu, A. Gementzopoulos, G. Sedky, A. R. Jones, and F. D. Lagor, Design of optimal wing maneuvers in a transverse gust encounter through iterated simulation or experiment, *Theor. Comput. Fluid Dyn.* **37**, 465 (2023).
- [12] K. Arulkumaran, M. P. Deisenroth, M. Brundage, and A. A. Bharath, Deep reinforcement learning: A brief survey, *IEEE Signal Process. Mag.* **34**, 26 (2017).
- [13] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction* (MIT Press, Cambridge, MA, 2018).
- [14] H. Kurniawati, Partially observable Markov decision processes and robotics, *Annu. Rev. Control Robot. Auton. Syst.* **5**, 253 (2022).
- [15] S. Verma, G. Novati, and P. Koumoutsakos, Efficient collective swimming by harnessing vortices through deep reinforcement learning, *Proc. Natl. Acad. Sci. USA* **115**, 5849 (2018).
- [16] G. Novati, L. Mahadevan, and P. Koumoutsakos, Controlled gliding and perching through deep-reinforcement-learning, *Phys. Rev. Fluids* **4**, 093902 (2019).
- [17] P. Gunnarson, I. Mandralis, G. Novati, P. Koumoutsakos, and J. O. Dabiri, Learning efficient navigation in vortical flow fields, *Nat. Commun.* **12**, 7143 (2021).
- [18] D. Fan, L. Yang, Z. Wang, M. S. Triantafyllou, and G. E. Karniadakis, Reinforcement learning for bluff body active flow control in experiments and simulations, *Proc. Natl. Acad. Sci. USA* **117**, 26091 (2020).
- [19] N. J. Nair and A. Goza, Bio-inspired variable-stiffness flaps for hybrid flow control, tuned via reinforcement learning, *J. Fluid Mech.* **956**, R4 (2023).
- [20] C. Xia, J. Zhang, E. C. Kerrigan, and G. Rigas, Active flow control for bluff body drag reduction using reinforcement learning with partial measurements, *J. Fluid Mech.* **981**, A17 (2024).
- [21] G. Novati, H. L. de Laroussilhe, and P. Koumoutsakos, Automating turbulence modelling by multi-agent reinforcement learning, *Nat. Machine Intell.* **3**, 87 (2021).
- [22] H. J. Bae and P. Koumoutsakos, Scientific multi-agent reinforcement learning for wall-models of turbulent flows, *Nat. Commun.* **13**, 1443 (2022).
- [23] P. I. Renn and M. Gharib, Machine learning for flow-informed aerodynamic control in turbulent wind conditions, *Commun. Eng.* **1**, 45 (2022).
- [24] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, Planning and acting in partially observable stochastic domains, *Artif. Intell.* **101**, 99 (1998).
- [25] L. P. Kaelbling, M. L. Littman, and A. W. Moore, Reinforcement learning: A survey, *J. Artif. Intell. Res.* **4**, 237 (1996).
- [26] L. Lin and T. M. Mitchell, *Memory Approaches to Reinforcement Learning in Non-Markovian Domains* (Carnegie Mellon University, Pittsburgh, PA, 1992).
- [27] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and A. Riedmiller, Playing Atari with deep reinforcement learning, [arXiv:1312.5602](https://arxiv.org/abs/1312.5602).
- [28] M. Hausknecht and P. Stone, Deep recurrent Q-learning for partially observable MDPs, in *2015 AAAI Fall Symposium Series* (AAAI Press, Washington, DC, 2015).
- [29] L. Meng, R. Gorbet, and D. Kulić, Memory-based deep reinforcement learning for POMDPs, in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, New York, 2021), pp. 5619–5626.
- [30] S. Fujimoto, H. van Hoof, and D. Meger, Addressing function approximation error in actor-critic methods, in *Proceedings of the 35th International Conference on Machine Learning*, edited by J. Dy and A. Krause, Proceedings of Machine Learning Research (PMLR, 2018), Vol. 80, pp. 1587–1596.
- [31] D. Beckers, Fast models and reinforcement learning control of unsteady aerodynamics, Ph.D. thesis, University of California, Los Angeles, 2023.
- [32] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor, in *Proceedings of the 35th International*

- Conference on Machine Learning*, edited by J. Dy and A. Krause, Proceedings of Machine Learning Research (PMLR, 2018), Vol. 80, pp. 1861–1870.
- [33] H. Wagner, Über die Entstehung des dynamischen Auftriebes von Tragflügeln, *Z. Angew. Math. Mech.* **5**, 17 (1925).
- [34] T. Theodorsen, General theory of aerodynamic instability and the mechanism of flutter, Tech. Rep. NACA-TR-496 (NASA, 1949), <https://ntrs.nasa.gov/citations/19930090935>.
- [35] The MathWorks Inc., Control system toolbox version: 23.2 (r2023b) (2023).
- [36] J. D. Eldredge, A method of immersed layers on Cartesian grids, with application to incompressible flows, *J. Comput. Phys.* **448**, 110716 (2022).
- [37] <https://github.com/diederikb/AeroGym>.
- [38] https://github.com/diederikb/aero_gym_SB3.
- [39] A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dormann, Stable-Baselines3: Reliable reinforcement learning implementations, *J. Mach. Learn. Res.* **22**, 1 (2021).
- [40] T. von Kármán and W. R. Sears, Airfoil theory for non-uniform motion, *J. Aeronaut. Sci.* **5**, 379 (1938).
- [41] S. Neumark, Pressure distribution on an airfoil in nonuniform motion, *J. Aeronaut. Sci.* **19**, 214 (1952).
- [42] R. T. Jones, Operational treatment of the nonuniform-lift theory in airplane dynamics, Tech. Rep. NACA-TN-667, NACA, 1938 (NASA, 1938), <https://ntrs.nasa.gov/citations/19930081472>.
- [43] I. E. Garrick, On some reciprocal relations in the theory of nonstationary flows, Tech. Rep. NACA-TN-629, NACA, 1938 (unpublished).