



MOX-Report No. 107/2024

**Active Flow Control for Bluff Body under High Reynolds Number
Turbulent Flow Conditions Using Deep Reinforcement Learning**

Chen, J.; Ballini, E.; Micheletti, S.

MOX, Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

mox-dmat@polimi.it

<https://mox.polimi.it>

Active Flow Control for Bluff Body under High Reynolds Number Turbulent Flow Conditions Using Deep Reinforcement Learning

J. Chen[†], E. Ballini[#], S. Micheletti[#]

December 19, 2024

[†] Politecnico di Milano, Piazza L. da Vinci 32, I-20133 Milano, Italy
jingbo.chen@mail.polimi.it

[#] MOX – Modellistica e Calcolo Scientifico, Dipartimento di Matematica
Politecnico di Milano, Piazza L. da Vinci 32, I-20133 Milano, Italy
{enrico.ballini, stefano.micheletti}@polimi.it

Abstract

This study employs Deep Reinforcement Learning (DRL) for active flow control in a turbulent flow field of high Reynolds numbers at $Re = 274000$. That is, an agent is trained to obtain a control strategy that can reduce the drag of a cylinder while also minimizing the oscillations of the lift. Probes are placed only around the surface of the cylinder, and a Proximal Policy Optimization (PPO) agent controls nine zero-net mass flux jets on the downstream side of the cylinder. The trained PPO agent effectively reduces drag by 29% and decreases lift oscillations by 18% of amplitude, with the control effect demonstrating good repeatability. Control tests of this agent within the Reynolds number range of $Re = 260000$ to 288000 show the agent's control strategy possesses a certain degree of robustness, with very similar drag reduction effects under different Reynolds numbers. Analysis using power spectral energy reveals that the agent learns specific flow frequencies in the flow field and effectively suppresses low-frequency, large-scale structures. Graphically visualizing the policy, combined with pressure, vorticity, and turbulent kinetic energy contours, reveals the mechanism by which jets achieve drag reduction by influencing reattachment vortices. This study successfully implements robust active flow control in realistically significant high Reynolds number turbulent flows, minimizing time costs (using two-dimensional geometrical models and turbulence models) and maximally considering the feasibility of future experimental implementation.

1 Introduction

In aircraft, navigation, and related industries, we often encounter the technical task of minimizing drag. To address these challenges, we usually depend on passive

control or active control [18]. An approach to reducing drag in specific cases is by modifying the aerodynamic shape of an object[49], for example, the dimpled design of golf balls [10].

For items such as large trucks, trains, wind turbines and aircrafts that have geometric shapes set and where reducing drag leads to a decrease in energy consumption [3, 30, 51, 8], active control is required after aerodynamic optimization has already been completed.

Active control can be categorized as open-loop active control [43] and closed-loop active control. Open-loop active control is easier to set up, but is very unstable when dealing with external disruptions or internal parameter variations. Closed-loop active control is frequently utilized in intricate situations for flow control [9, 11, 1, 19, 24], particularly in the context of Active Flow Control (AFC) [40].

In AFC, the input consists of physical measurements collected by probes within the flow field, while the output involves altering the flow field using actuators. Establishing a reliable transfer function between input and output is difficult, particularly for turbulent flow fields, and the stability of the transfer function is weak due to several disruptions in real-world settings [5]. Therefore, it is crucial to integrate machine learning (ML) for active closed-loop control of the flow field [28, 25, 35, 37, 31, 17, 26, 46, 53, 5, 16]. The integration of machine learning with fluid dynamics is closely linked and has been applied in research areas such as flow field reconstruction and prediction [42], turbulence modeling [15, 47], and flow control [36, 34, 48, 54]. Yunfei Li [27] and Steven L. Brunton [6] offer insightful overviews of recent advances in the field. Reinforcement Learning (RL) [45, 12] or Deep Reinforcement Learning (DRL) [2] approaches are frequently selected to attain control. DRL utilizes Neural Networks (NN) to address complicated, high-dimensional problems within the subject of ML, whereas RL utilizes tabular approaches and Q-Learning to develop strategies. Reinforcement learning is not feasible for complex control issues that involve continuous high-dimensional action spaces. Hence, we opt to utilize DRL with Deep Neural Networks (DNN) to address intricate control challenges [38, 14, 8, 7, 29, 33, 32].

Deep Neural Networks in DRL acts as transfer functions in control theory, serving as a black-box model to accurately represent the relationships between inputs and outputs. DNN is well-suited for modeling complicated nonlinear systems, like those needing turbulence control, which are difficult to model using traditional methods. Therefore, substantial computational effort is required for proper fitting of the nonlinear function between inputs and outputs, along with a suitable update approach to update the weights of DNN. Recent research focuses on machine learning-based active flow control at low Reynolds numbers [35, 31, 46, 53]. Researchers such as Feng Ren [35] and Hongwei Tang [46] have successfully reduced the drag on bluff bodies under weakly turbulent laminar and turbulent conditions, achieving effective control across various Reynolds numbers. Active flow control at high Reynolds numbers is important in various engineering applications, including reducing drag for cars [41, 20, 22], high-speed trains and wings [21], as well as attitude control for aircraft without rudders. Studies such as those of Zhi-

jie Zhao have utilized synthetic jets for attitude control of such aircraft [52] and Enrico Ballini successfully implemented flow control over bluff bodies at higher Reynolds numbers using the PPO algorithm [4].

Optimizing control for high Reynolds number flow configurations using DRL encounters numerous hurdles. Fluid mechanics simulations are time-intensive, and training the network requires detailed flow field data like velocity and pressure, resulting in a huge number of network inputs. The significant time cost for training the agent is related to the high non-linearity between inputs and outputs. When computational resources are limited, it is practical to begin by utilizing simplified models to investigate efficient flow field control. This can involve employing two-dimensional geometric models or a range of turbulence models during the early stages of study. Efficiently exploiting samples is key in DNN's update process. This is because a large amount of samples requires significant data processing, and the majority of deep-reinforcement learning agents use the Gradient Descent Method (GDM) for updates. Increasing the number of samples improves the accuracy of gradient computation, but also leads to a substantial increase in computational workload, requiring a trade-off. It is crucial to reduce the total time cost of agent training in turbulent flow circumstances by examining the agent's ability to learn flow field properties and optimizing control based on features. The stability of the training results can be assessed by examining the reproducibility of the training and the robustness and reaction speed of the selected policy control mechanism after training, focusing on control aspects.

Previous research on flow control with DRL has mainly used the PPO algorithm [35, 31, 46, 39]. This label-free learning technique streamlines data processing and is commonly employed in flow control for many reasons. Flow control is a challenge due to the continuous action space, and the PPO algorithm excels at addressing such complexities, playing a vital role in improving decision-making. Furthermore, PPO includes a clipping mechanism that inhibits sudden changes in actions. PPO has an exploratory tendency that helps prevent getting stuck in local optima, although this can lead to slower convergence rates. Defining the reward function enables the optimization of many objectives at the same time.

In view of these PPO algorithmic features, this article also employs PPO as the basis for its investigation of flow control-based drag reduction. Virtual probes are positioned only on the cylinder surface to prevent interference from upstream probes in the wake from impacting downstream probes, considering practicality in real experimental settings and engineering uses, as well as the significance of experimental-simulation comparisons. Given these limitations, the maximum amount of flow field data is collected, and the flow field is managed using jetting. Xinhui Dong et al. placed sensors around the cylinder to monitor pressure and trained an agent for suction control, experimentally achieving a reduction in the adverse effects of vortex shedding at the rear of the cylinder and suppressing lift fluctuations [13].

The policy developed in this study significantly reduced the drag on the cylinder by 29% in a flow field at $Re = 274000$ and exhibited a very high response speed.

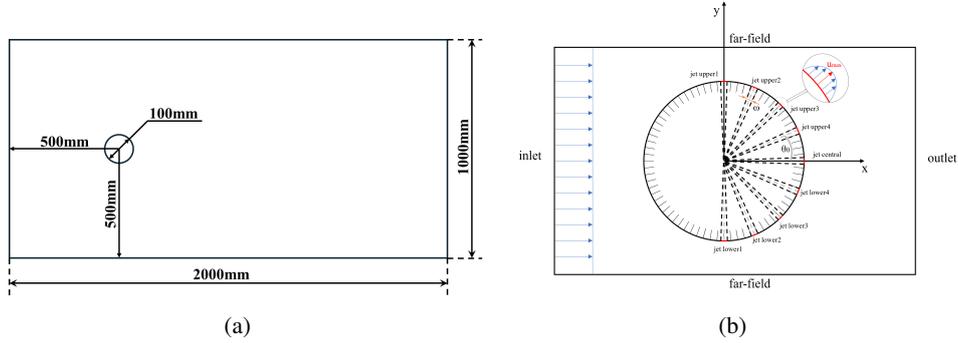


Figure 1: Geometrical configuration of the circular cylinder and the computational domain: (a) The relative position of the cylinder in the computational domain. (b) Schematic of the jets on the cylinder with geometrical dimensions and boundary conditions with the profile of velocity inlet and the profile of jets.

This control strategy is also robust and capable of handling variations in Reynolds numbers within a certain range without excessively increasing the oscillation of vortex shedding. The paper is framed as follows. Section 2 introduces the numerical methods and the PPO algorithm, as well as the methodology for coupling the agent with the simulation environment. Section 3 discusses the trained agent based on the repeatability of the training results, analyzing the impact of control on the flow field through velocity, pressure, and vorticity contour plots. The effectiveness of the control is further examined by comparing and analyzing data and graphs. Section 4 summarizes the work presented in this paper and describes directions for future research.

2 Methodology

2.1 Flow Configuration

The configuration in this investigation is a turbulence flow past a circular cylinder in a two-dimensional setting with a Reynolds number (Re) of 274000 where $Re = \bar{U}D/\nu$. \bar{U} is the mean velocity of inflow from infinity, D is the diameter of the cylinder, and ν is the kinematic viscosity of the fluid. The diameter of the cylinder is 100mm and the domain of the numerical simulation is 2000mm(*length*) \times 1000mm(*height*), the center of the cylinder is placed 500mm away from the inlet boundary and 500mm away from the far field boundary (see Fig. 1(a)).

For flow control, we use jets to influence the flow field around a cylinder to achieve the control objective. Figure 1(b) shows the configuration of nine control jets which are installed symmetrically on the surface of the cylinder with different angular positions relative to the x -axis. We can define the position of the jets through θ , $\theta_i(\text{jet upper}) = 90^\circ, 67.5^\circ, 45^\circ, 22.5^\circ$ with $i = 1, 2, 3, 4$ and

$\theta_i(\text{jetlower}) = -90^\circ, -67.5^\circ, -45^\circ, -22.5^\circ$ with $i = 1, 2, 3, 4$ and finally the central jet at 0° . The dimension of the jet outlet is limited by the angle ω and the angle ω is always equal to 5° . To avoid the discontinuity of velocity when we define the boundary condition of the jets and no-slip wall on the surfaces of the cylinder, we choose the velocity profile at each jet outlet following the expression defined below [46]:

$$u = u_{\max} \cos\left(\frac{\pi}{\omega}(\theta - \theta_0)\right),$$

where u_{\max} is limited between $-2m/s$ and $2m/s$ and because it is an incompressible flow, the density of the fluid is constant, which means that the Artificial Neural Network (ANN) delivers the velocities, so the agent controls the velocity of nine jets directly: when the jet velocity is positive, it means blowing; otherwise it means suction. The nine jets are zero net mass flux jets, implying that at any given moment the total flux through the boundaries of these nine jets is zero, eliminating the need for additional mass injection into the flow field.

2.2 Problem description

In the present study, the flow is considered viscous and incompressible. The stationary no-slip wall is applied to the cylinder. The governing equations are the two-dimensional Reynolds-averaged Navier-Stokes equations (RANS).

The continuity equation for incompressible flow reads:

$$\frac{\partial \bar{u}_i}{\partial x_i} = 0.$$

The momentum equation for incompressible flow reads:

$$\rho \left(\frac{\partial \bar{u}_i}{\partial t} + \bar{u}_j \frac{\partial \bar{u}_i}{\partial x_j} \right) = -\frac{\partial \bar{p}}{\partial x_i} + \mu \frac{\partial^2 \bar{u}_i}{\partial x_j^2} - \rho \frac{\partial \overline{u'_i u'_j}}{\partial x_j}.$$

To close the system, we need to model the Reynolds stresses $\overline{u'_i u'_j}$. In this regard, we adopt the $k - \omega$ SST turbulence model, which can calculate flow separation relatively accurately, according to the research of Wen and Qiu [50, 23]. The use of the Large Eddy Simulation (LES) method, which directly resolves large-scale eddies, imposes severe requirements on the mesh and demands more computational resources. The 2D and 3D RANS models still exhibit shortcomings in capturing the drag crisis phenomenon [44], but the calculated drag coefficients remain within a reasonable range [50]. Given the already substantial expense of carrying out numerical simulations in three dimensions, combined with the extensive time needed for training the agent, the task becomes exceedingly challenging. Therefore, we chose to simplify this model to a two-dimensional problem and apply the RANS model. In the present work, we use the SIMPLE scheme to solve the incompressible RANS model problem. The numerical simulation is based on the commercial software ANSYS-Fluent. Drag, F_D , and lift, F_L , can be calculated by integrating

the viscous stress and pressure acting on the surface of the cylinder, and in this case, the reaction force contributed by jet propulsion should be considered. The drag force and lift force, the drag coefficient, C_D , and the lift coefficient, C_L , are defined below.

$$F_D = \int (\boldsymbol{\sigma} \cdot \boldsymbol{n}_c) \cdot \boldsymbol{e}_x dS,$$

$$F_L = \int (\boldsymbol{\sigma} \cdot \boldsymbol{n}_c) \cdot \boldsymbol{e}_y dS,$$

$$C_D = \frac{2F_D}{\rho \bar{U}^2 D},$$

$$C_L = \frac{2F_L}{\rho \bar{U}^2 D}.$$

The stress tensor is denoted by $\boldsymbol{\sigma}$, \boldsymbol{n}_c represents the unit outward normal vector to the outer surface of the cylinder, and $\boldsymbol{e}_x = (1, 0)$ and $\boldsymbol{e}_y = (0, 1)$. The Strouhal number, St , is a dimensionless number used to represent the frequency of the dominant oscillation of the flow, and is defined as follows:

$$St = \frac{f_s D}{\bar{U}}.$$

The quantity f_s is a natural frequency, representing the vortex shedding frequency that can be calculated from the periodic oscillation of the lift coefficient, C_L .

2.3 Mesh independence

The unstructured mesh used in the simulation is made up of 182612 triangular grids (Fig. 2). To further reduce the computational load without compromising the accuracy of the numerical simulations, we divide the computational domain into three nested regions with increasing grid density. In addition, the grid around the cylindrical wall is much more refined because of the boundary layer. Table 1 illustrates that convergence is reasonably assumed with respect to the mesh size and time step. The mean values of C_D for meshes of 182612 and 283061 elements at a time step of $\delta_t = 0.00025s$ vary by just 0.5%. Therefore, considering computational efficiency, we opted for the mesh with 182612 elements for this study.

2.4 Deep Reinforcement Learning

The basic framework of PPO is illustrated in Fig. 3, and the details of the PPO algorithm are presented in Appendix A, while the employed hyper parameters are listed in Appendix B. It is essentially a closed-loop control system where observations and rewards are input variables, and actions are output variables during neural network training. For a pre-trained neural network, inputting an observation allows the agent to quickly select the corresponding policy, which is the action. The training process can be viewed as constructing a transfer function or, more precisely, approximating a transfer function. In general, the agent interacts with

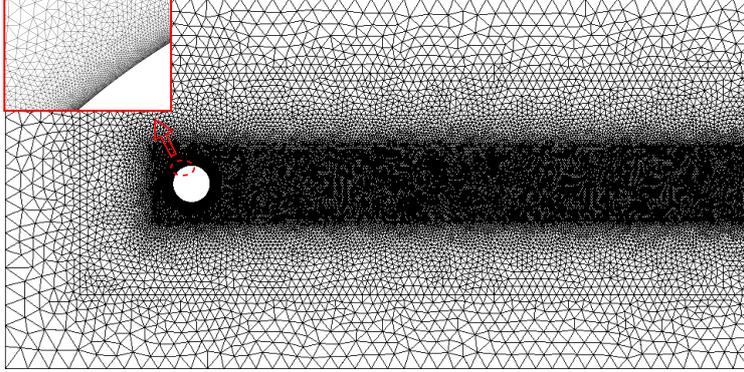


Figure 2: Computational mesh: the innermost mesh is refined and the mesh around the cylinder is highly refined considering the accuracy of the simulation in the boundary layer.

Table 1: Mesh convergence with mesh resolution and time step size referring drag coefficient for the flow around the circular cylinder at $Re = 274000$.

N. elements	28874	85849	182612	283061
$\delta_t = 0.001$				
C_{D_mean}	0.52522	0.78816	0.79094	
$\delta_t = 0.0005$				
C_{D_mean}		0.87778	0.84921	
$\delta_t = 0.00025$				
C_{D_mean}		0.84996	0.89337	0.89754

the environment, which in our study is a virtual environment corresponding to the simulation of a flow field. The data generated from the flow field simulation is transmitted across a channel to the agent as an observation input. This corresponds to the velocity and pressure information obtained from the probes. Additionally, the lift coefficient and drag coefficient are calculated by simulation to compute the reward input for the PPO agent:

$$Reward = -|C_D| - 0.2|C_L|.$$

A specific weighting for drag and lift is used to minimize drag without significantly boosting lift, emphasizing the importance of this optimization target in relation to the research objectives. The agent uses observations and rewards to train the DNN and computes actions. This study examines the maximum values of nine jet vectors. The values are used as input for the velocity profile function of the jets, and the equations act as boundary conditions for the numerical simulations. Executing the action in the environment completes a full interaction cycle between



Figure 3: Schematic of the PPO agent closed-loop control.

the PPO agent and the environment.

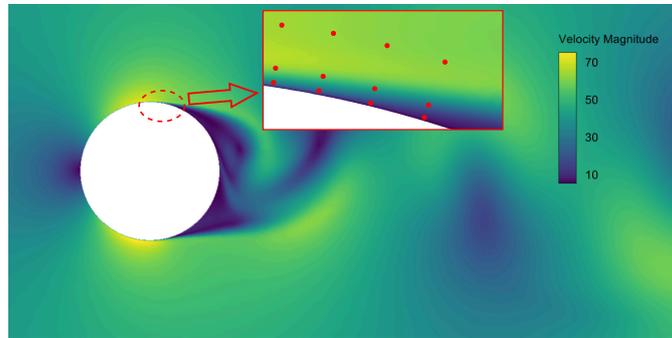


Figure 4: Schematic diagram of three layers of sensors placed around the surface of the cylinder, ensuring that the first layer of sensors is placed within the boundary layer.

Zooming in on the figure reveals that we have placed three layers of probes around the cylinder, marked with red dots, as shown in Fig. 4. The first layer is placed at a height of 0.1mm from the cylinder surface and can be used to measure airflow data over the surface of the cylinder; the second and third layers are placed at 1mm and 2.5mm , respectively. It is crucial to ensure that the first layer of probes is within the boundary layer, enabling us to gather information about the flow separation points. The second and third layers of the probes aim to obtain information on pressure and velocity gradients, as well as information on the effects of flow separation on the wall surface, allowing the agent to learn and determine whether flow separation is exacerbated under the influence of the jets. In simpler terms, we wish the fluid to flow closely along the surface of the cylinder. Each layer consists of 128 probes, each probe capable of simultaneously capturing two types of information: velocity and pressure.

To facilitate the interaction between the agent and the environment, we establish a framework as shown in Fig. 5. Using a Text User Interface (TUI), a dialogue is established between Matlab and Fluent, with Matlab constructing the PPO agent. Matlab outputs the control variable, which is the maximum value of jet velocity, and this is fed into Fluent via the TUI, followed by the complete input of veloc-

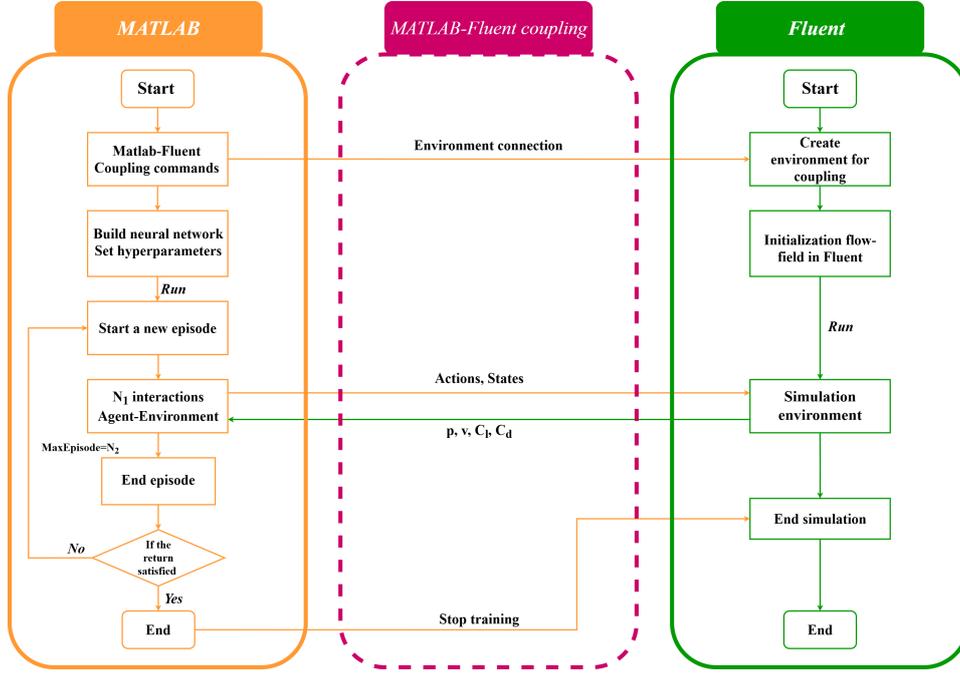


Figure 5: Flow chart of the Matlab-Fluent coupling with training process and agent-environment interactions.

ity inlet through the velocity profile equation. Once all boundary conditions have been entered, Fluent runs a simulation for one time step, which includes n iterative steps. A complete interaction between Matlab and Fluent of N_1 times constitutes one episode, with N_2 being the total number of episodes in a training session. In our study, rather than resetting the flow field to its initial state at the end of an episode, the flow field is simulated for several time steps in an environment with only an inlet and no jets, allowing the flow field to return to an uncontrolled state, effectively conducting a random initialization.

3 Results and Discussions

As discussed previously, training an agent for control in high Reynolds-number flows is a time-consuming and uncertain endeavor. Each complete training session encompasses 100 episodes, with each episode containing 100 interactions, and each training session takes approximately 120 hours (hardware:i7-12700H @2,3GHz RAM 64GB). The return for each episode is the result of the weighted summation of the rewards over an entire trajectory. To demonstrate the reproducibility of the training, three independent training sessions were conducted, as shown in Fig. 6, and the returns of these sessions were averaged to obtain the mean return. It can be seen from the graph that the fluctuation in rewards during each

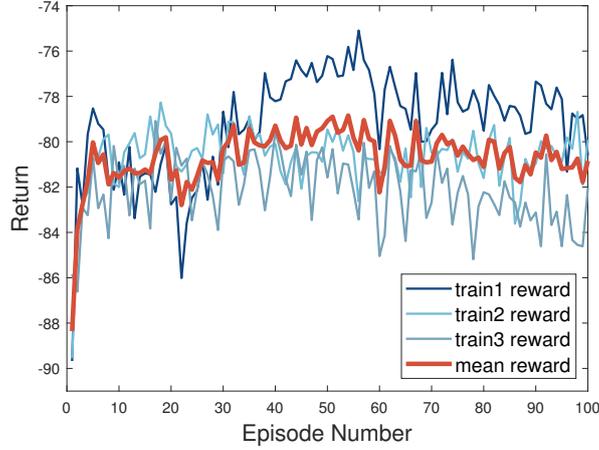


Figure 6: Training curve of PPO agent with three independent trainings.

training session is significant. Given the high time cost and the complexity of the environment, it is undesirable for the agent to remain in a local optimum during training. The agent should be more exploratory, since staying in a local optimum would result in a more stable and convergent return curve, but achieving a similarly effective control outcome would require exponentially more time. Thus, balancing training costs while obtaining a reliable control strategy becomes a challenging dilemma. We found that the mean return stabilizes around -81 after some fluctuations, indicating that individual training sessions are repeatable. Although we cannot conclude that this is the optimal control strategy (as there is no definitive criterion for determining an optimal control strategy), it suggests that finding a better control strategy than the current training results is extremely difficult. In this case, we need to verify the repeatability of the control strategy of the trained agent, in other words, whether agents with similar return values obtained from different trainings have similar control effects. We also need to assess whether the agent's control strategy meets the requirements at the control level, such as quick response speed and robustness. These aspects will be discussed further.

In the coefficient of drag graph, we disregard the data from the flow field transition phase and focus solely on the fully developed flow field starting from $0.05s$ onward, as shown in Fig. 7(a). This graph compares the drag coefficient of the cylinder without control to that under agent control after training completion. The control transition area, highlighted in a darker shade of gray, represents the process through which the flow field transitions from an uncontrolled state to a stable controlled state. The control strategy is observed to respond in approximately 0.03 seconds, as shown in Fig. 7(a), quickly transitioning the flow field to a controlled state with a significant reduction of 29% in the drag coefficient, demonstrating the practical importance of this control strategy. In the coefficient of lift graph, it is evident that the amplitude of the oscillation is reduced. However, an asymmet-

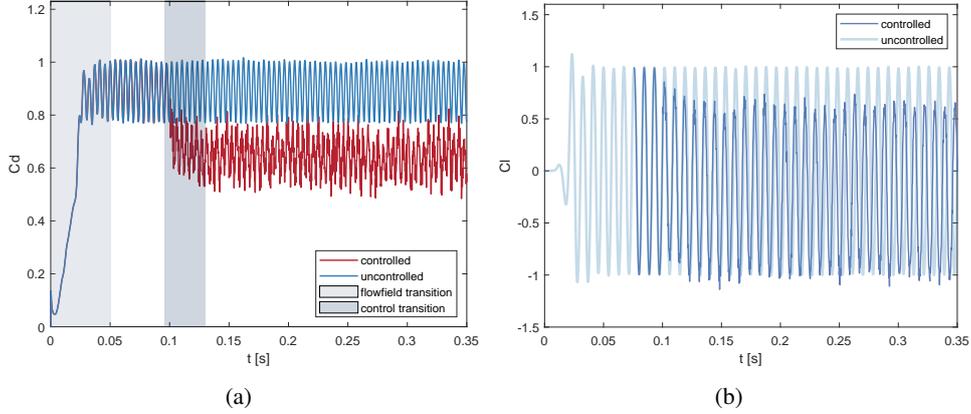


Figure 7: Comparison of the drag coefficient and lift coefficient of the cylinder under agent control versus uncontrolled conditions.

ric phenomenon in the vertical resultant force is also noticeable compared to the uncontrolled condition, as shown in Fig. 7(b). This implies that the symmetry of the agent’s policy under this flow characteristic can be further improved based on specific requirements. When comparing the drag and lift curve graphs, it is not difficult to notice that the oscillation amplitude of the lift predominates.

From a statistical perspective, the flow field data are averaged over time to obtain the result of the difference between the flow field without control and the flow field under the control of a trained agent. In particular, the average is performed after simulating 100 periods of oscillation. By calculating the difference, we can more intuitively observe the changes in the flow field caused by the jet, as the maximum speed of the jet is only $2m/s$, while the maximum speed around the cylinder can reach $70m/s$. The wake of the flow field is observed to be relatively symmetric, indicating that the flow field is statistically stable under control. However, slight asymmetry can be noticed around the cylinder through the mean streamwise velocity, the mean transverse velocity, and the mean static pressure, with a velocity difference of approximately $2m/s$. It is important to note that the maximum speed of the jet is also controlled at $2m/s$, with the inlet flow speed at $40m/s$. This slight asymmetry is understandable compared to the results achieved by the control. Possibly, this asymmetry is due to the uncertainties in the flow and the training, which prevent the agent from quickly mastering the macroscopic periodic patterns of the flow. For future work, it could be considered to integrate algorithms that can recognize signal features to explore whether agents can quickly learn the macroscopic periodic patterns of the fluid, and to investigate whether this method could accelerate agent training or stabilize the training at a better control effect. Additionally, the symmetry in the mean vorticity diagram is relatively good, indicating that the jet has a minimal impact on vorticity.

To verify the repeatability of the control strategy, we selected an agent trained from one of the three independent training sessions conducted after the mean re-

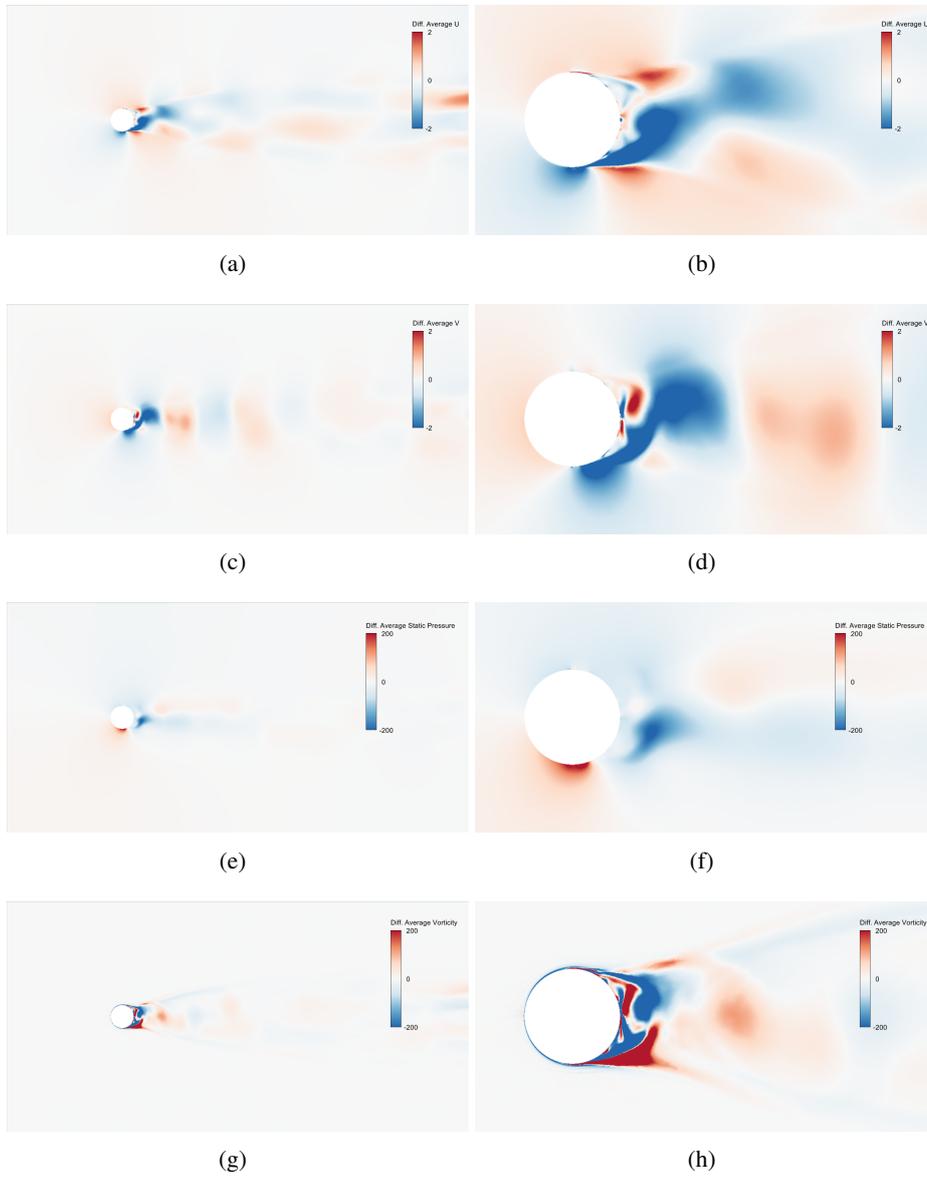


Figure 8: Contour of averaged difference of streamwise velocity (a), averaged difference of transverse velocity (b), averaged difference of static pressure (e) and averaged difference of vorticity (g) over 100 oscillation periods; the contours on the right column (b), (d), (f), (h) are enlarged views.

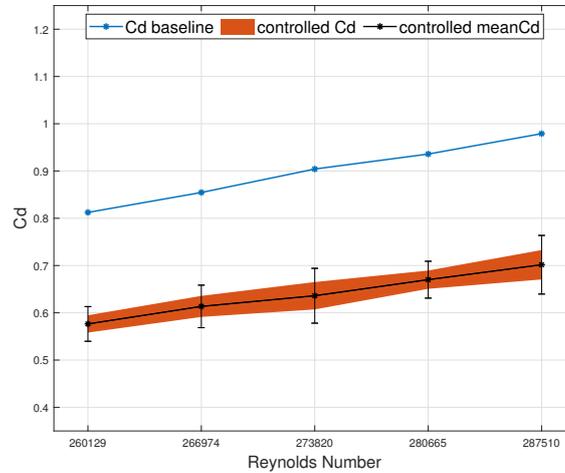


Figure 9: Comparison of the coefficient of drag before and after control under a range of Reynolds number. The blue line is the baseline without control and the black line is the coefficient of drag averaged by five independent repeated control tests. The error bars denote the difference between the maximum and minimum coefficient of drag among these five tests. The orange area represents the standard deviation range.

ward curve stabilized. This agent was subjected to five repeated control tests, each lasting 100 periods of oscillation, during which the drag coefficient was calculated. This process was also repeated with variations in Reynolds numbers to assess the robustness of the control strategy, as shown in Fig. 9. The blue curve represents the variation in the drag coefficient under different Reynolds numbers without control. The orange area indicates the range of variance in the drag coefficient after five tests, while the black curve represents the average drag coefficient under different Reynolds numbers after testing. The error bars show the difference between the maximum and minimum drag coefficients obtained in five tests at each Reynolds number. The results of five tests with different Reynolds numbers show that the drag under agent control varies within a small range, proving that the control strategy is repeatable and stable. Comparison of drag coefficients before and after control under different Reynolds numbers shows that the control strategy possesses a certain degree of robustness.

In tests carried out in various Reynolds numbers using an agent trained in $Re = 244000$, the robustness of the procedure was established. The study delved into the impact of the same agent on flow field structures under different flow conditions from a frequency domain perspective. The blue and orange curves in the graph represent the lift coefficient curves without control and under agent control, respectively, post Fast Fourier Transform (FFT). These curves illustrate the distribution of energy at different frequencies due to vertical oscillations acting on the cylinder,

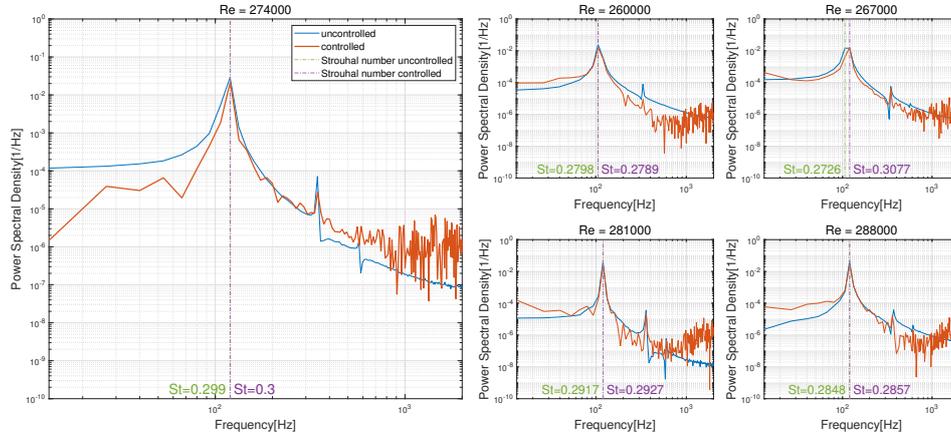
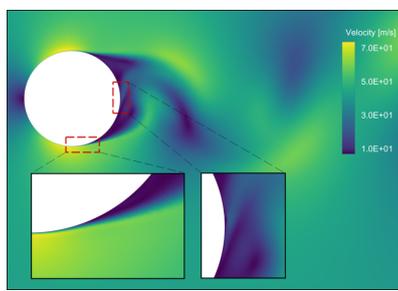


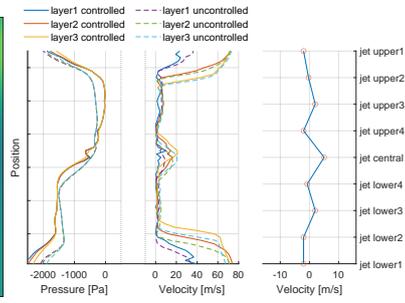
Figure 10: Diagram of the comparison of Power Spectral Density (PSD) before and after control at different Reynolds number, marked with Strouhal number.

as shown in Fig. 10. The low-frequency portion of the graph indicates large-scale structures within the flow, while the high-frequency portion indicates smaller-scale oscillations. The vertical dashed lines in green and purple, respectively, denote the Strouhal numbers without control and under agent control, corresponding to the peaks of the two curves at the vortex-shedding frequency on the horizontal axis. This frequency is characteristic of the von Kármán vortex streak. Post-control, the vortex-shedding frequency exhibits varying degrees of change across different Reynolds numbers. In particular, only at $Re = 274000$ is there a significant suppression of low-frequency energy, indicating an effective mitigation of lift oscillations caused by large-scale vortices formed as the fluid passes the cylinder. However, this phenomenon was not as effectively observed under the control of this agent at other Reynolds numbers, suggesting that the agent's learning was specifically tuned to the frequency characteristics of the flow conditions at training. Due to the use of the RANS model, the impact of turbulent fluctuations across different scales in the frequency domain was not very pronounced, leading to relatively smoother curves. Nevertheless, a comparison of the orange curves across different Reynolds numbers reveals significant fluctuations in the energy in the high-frequency region post-control, primarily showing an increase. This suggests that a significant portion of the flow control impact targets small-scale flow structures. Within the studied range of Reynolds numbers, it is evident that high-frequency, small-scale structures in the flow field carry less energy. The blue curve further illustrates that as the scale of flow decreases, energy is transferred from large-scale to small-scale structures. The slope of the curve to the right of the vortex shedding frequency indicates that the control strategy does not significantly alter the distribution of energy across different scales or the inherent energy at the vortex shedding frequency.

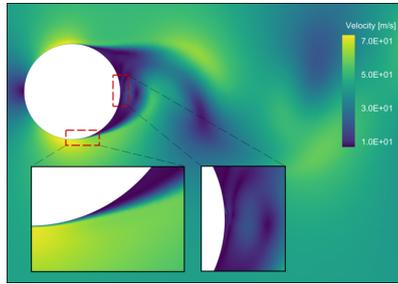
The series of velocity contours at time levels $t_0 + 0.001n$ ($n = 0, 1, 2, 3, 4, 5$),



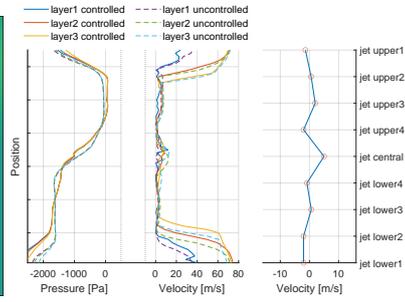
(a)



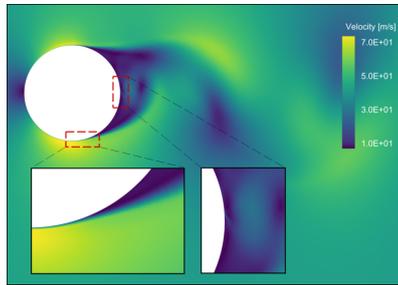
(b)



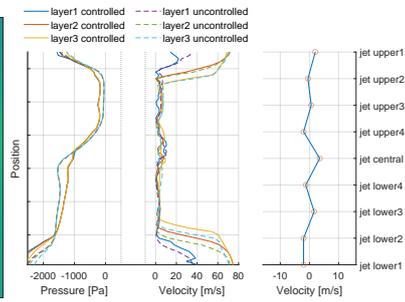
(c)



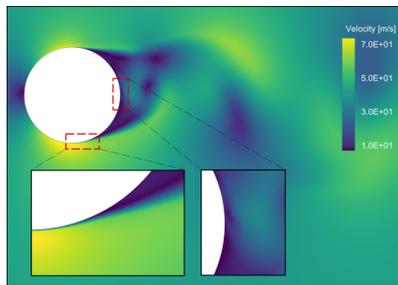
(d)



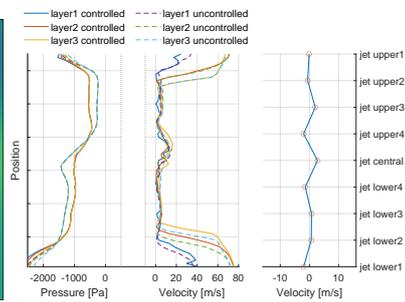
(e)



(f)



(g)



(h)

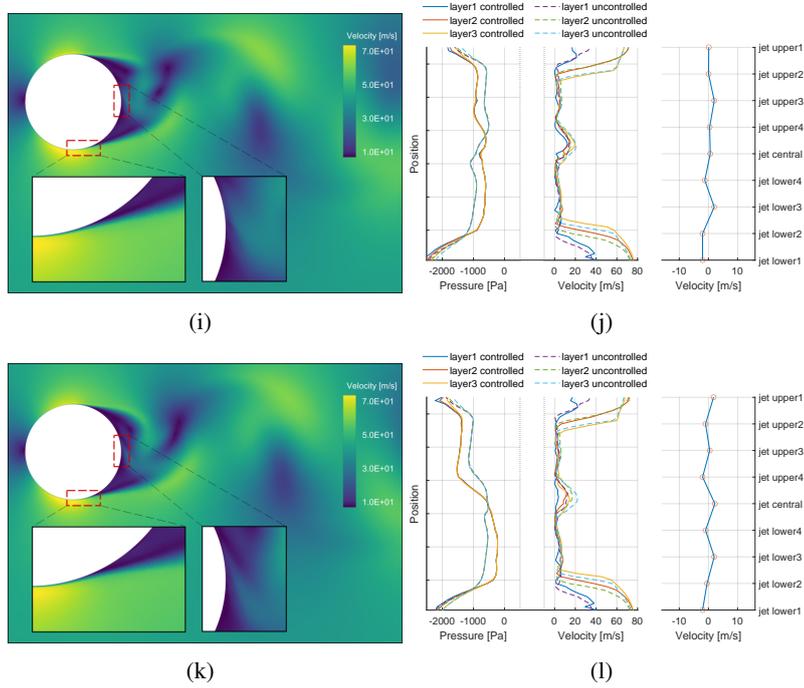
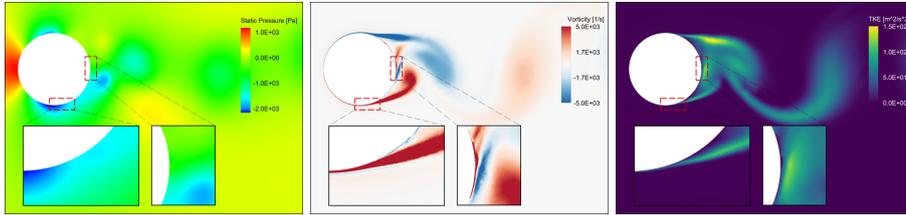


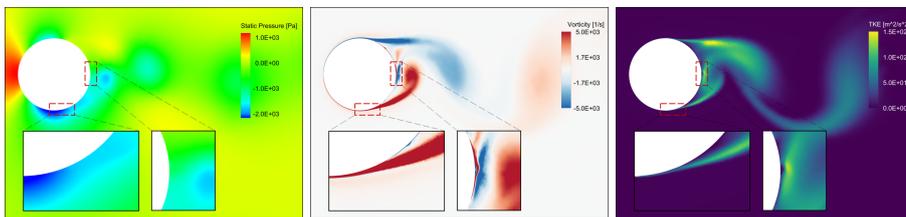
Figure 11: Visualization of the agent policy: Contour of velocity magnitude with zoomed in views on the left column: (a), (c), (e), (g), (i), (k). The comparison of the velocity and pressure detected by three layers of sensors on the right half of the cylinder corresponding the velocity of nine jets on the right column.



(a)

(b)

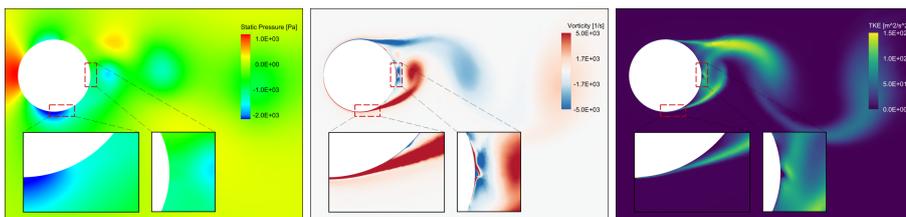
(c)



(d)

(e)

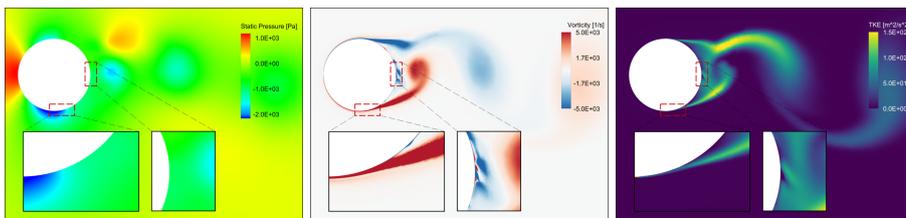
(f)



(g)

(h)

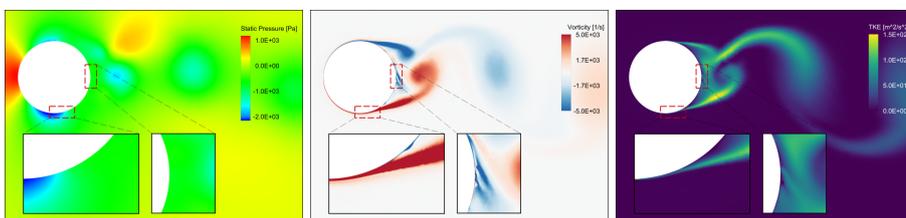
(i)



(j)

(k)

(l)



(m)

(n)

(o)

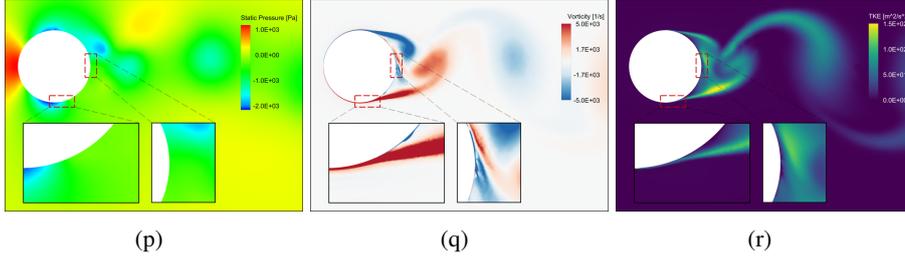


Figure 12: Comparison of the static pressure, vorticity and turbulent kinetic energy (TKE) at different times with time interval 0.001s.

as shown in Fig. 11, along with the velocity and pressure data extracted from three layers of sensors on the right-hand surface of the cylinder and the jet velocities of nine jets, visualize the control strategy or policy employed by the agent. The changes in velocity contour across different times and the fluctuations of pressure on the cylinder's surface elucidate the oscillatory process of vortex shedding. The enlarged images reveal the variations in the flow separation zone on the cylinder surface and how the shedding vortices reattach to the cylinder's back, impacting the pressure there. In Fig. 11(b), it is evident that under the influence of the central jet, the pressure at the back of the cylinder increases, resulting in a peak in the pressure curve.

The velocity diagram of the nine jets indicates that the central jet predominantly employs a jetting strategy, which results in the jet's reactive force directly propelling the cylinder and elevating the static pressure at its back, thereby reducing pressure drag. The velocities of jet upper1 and jet lower1 are minimal and exhibit little change, as increasing the velocity of these two jets would lead to premature flow separation and a larger separation zone, further decreasing the pressure at the cylinder's back and consequently increasing pressure drag. Therefore, the agent avoids employing a high-speed jetting strategy for these two jets. The other jets adjust their strategy periodically in response to the pulsations of the flow separation zone and the changes in the reattachment vortices, adopting a periodic strategy to optimize the control effect.

This series of contours reflect the dynamic changes in the flow field at time levels $t_0 + 0.001n$ ($n = 0, 1, 2, 3, 4, 5$), as shown in Fig. 12. By examining the zoomed-in views, the characteristics of the local flow can be observed. However, it is crucial to reiterate that, due to the use of the RANS model to simplify turbulence by time averaging for computational efficiency, not all details within the flow field are captured with high precision. However, by comparing the interaction between jets and vortex systems at different times across these three figures, we can reveal how jets achieve control objectives by influencing the flow field. The development of shedding vortices behind the cylinder can be seen through the changes in six sets of vorticity and turbulent kinetic energy contours, with alternating vortex streets forming above and below the cylinder. Zoomed-in views reveal a reattachment zone of

shedding vortices on the cylinder’s right-hand side (downstream), characterized by a lower static pressure and higher turbulent kinetic energy, indicating more disorderly flow. The vorticity contours also show that the flow in the reattachment zone exhibits periodicity, with these small structural vortex systems corresponding to the high-frequency portion mentioned previously in the power spectral density graph. Notably, a region near the central jet’s position exhibits significantly low turbulent kinetic energy (identified in the vorticity contour between the red and blue vortices in the white area on the right zoomed-in view), resulting from small vortex systems generated around the jet. This is indicative of the control strategy employed by the agent, which reduces turbulent kinetic energy near the surface at the rear of the cylinder, thus increasing static pressure at the back, reducing pressure drag, and achieving drag reduction. However, given the limited output power of the jets, a further improvement is needed to modify the flow in the reattachment zone. From the perspective of the dynamic development of shedding vortices, it is clear that when vortices of opposite rotation directions meet, the turbulent kinetic energy in the region decreases, the energy is dissipated, and the vorticity is reduced. For future research, it is essential to delve deeper into the flow mechanisms, particularly under high Reynolds number conditions, to elucidate the interactions between jets and vortices of various scales. To this end, transitioning to a three-dimensional model for the geometry is necessary.

4 Conclusions

The study used the Proximal Policy Optimization (PPO) method from Deep Reinforcement Learning (DRL) as a flow control model and simulations were performed using the $k - \omega$ SST turbulence model. The trained agent successfully reduced drag for a bluff body in high Reynolds number turbulence by periodically controlling the flow field. The results consist of the following.

- This study introduces Deep Reinforcement Learning (DRL) and establishes a zero net flow active flow control (AFC) model to reduce drag for bluff bodies under complex flow conditions at high Reynolds numbers.
- Training with random initialization of the flow field for each episode resulted in three distinct training sessions producing stabilized curves within a certain range, demonstrating that the training is reproducible.
- Under high Reynolds number conditions, the flow field is very intricate and displays significant nonlinearity. The agent gathered restricted data from the flow field due to potential interference between upstream and downstream probes; therefore, only data regarding velocity and pressure near the cylinder walls was obtained. The trained agent can achieve a fast response and control with a certain degree of robustness. The agent control method successfully decreased the drag on the cylinder by around 30% and showed reliable robustness within a specific range of Reynolds numbers.

- The agent, trained using Active Flow Control (AFC), mainly influenced the high-frequency small-scale structures of the flow field under the current conditions, with little effect on the frequency and energy of vortex shedding. It successfully dampened the low-frequency longitudinal oscillations affecting the cylinder, leading to a decrease in the fluctuating lift forces.

Based on this research, the primary goal for the future is to further explore the use of temporal algorithms such as LSTM (Long-Short Term Memory) to enable the agent to better and more quickly learn the periodic characteristics of the flow, thereby reducing training time costs. On this basis, our aim is to train the agent using simulation data within a $\pm 10\% - 15\%$ range of Reynolds numbers to further enhance the control robustness at high Reynolds numbers. Considering the practical conditions in the experiments, further studies are needed on the limitations with regard to the number of probes and their placement positions.

Acknowledgments

The present research is part of the activities of the "Dipartimento di Eccellenza 2023-2027". The second and third authors are members of the "Gruppi Nazionali di Ricerca INdAM". The authors express sincere thanks to Shumin Li for her helpful insights and contributions.

Data Availability Statement

Data supporting the findings of this study are available from the corresponding author on a reasonable request.

A Deep Reinforcement Learning - Proximal Policy Optimization

A Markov Decision Process (MDP) is built based on a set of interactive objects, as shown in Fig. 13, called an agent and environment[45]. Its elements include states, actions, and rewards. The agent will perceive the current state of the system, implement actions on the environment according to the policy, thus changing the state of the environment, and receiving a reward. In general, the time steps of the MDP can be finite or infinite; for the finite MDP there exists a terminal state which can be defined by episode; in this circumstance, the space of state and action are finite, so is the return. If there is no terminal state, in correspondence with the state, the action and return are infinite. Based on this process, the trajectory is given:

$$S_0, A_0, R_0, S_1, A_1, R_1, S_2, A_2, R_2 \dots S_t, A_t, R_t.$$

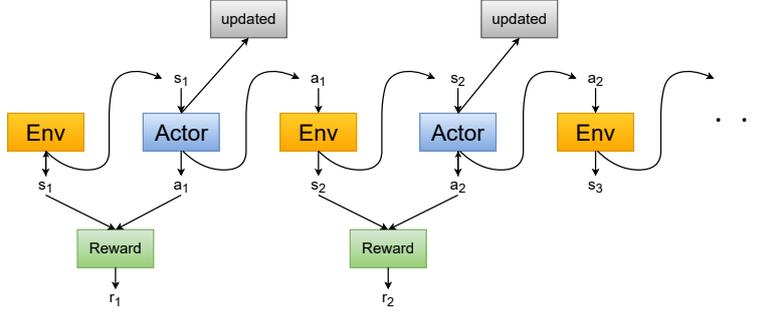


Figure 13: Markov decision process. The agent starts at S_t , the actor takes action A_t according to the state S_t , then, the environment generates a new state S_{t+1} .

This trajectory describes the problem in a discrete time sequence, so the subscript t means the discrete time or time step.

We introduce the return G_t , which is the sum of the rewards R of each discrete time t for an episode;

$$G_t = R_{t+1} + R_{t+2} + \dots + R_T.$$

When modeling real problems, according to the definition above, unless G_t tends to converge, it is inappropriate to consider infinite reward. To change this situation, we need to introduce the discounted return for episodic tasks,

$$G_t = R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \dots = \sum_{k=0}^{\infty} \gamma^k R_{t+k+1} = R_{t+1} + \gamma G_{t+1}.$$

The discount factor γ is a hyperparameter, in the range $[0, 1]$ and indicates importance of future rewards at the current time. If $\gamma = 0$, the agent is focused only on the current reward without looking at the future. If $\gamma = 1$, it recovers to the undiscounted return. So, using the discounted return, it seems to appear an infinite return in an infinite MDP task and makes the solution of MDP feasible.

The policy is specifically the action in each state, defined by the function π ,

$$\pi(a|s) = P(A_t = a | S_t = s).$$

Mathematically, the function π is the density of probability. We also define $\pi(a|s, \theta)$ as the policy network, which aims to approximate the policy $\pi(a|s)$ by using the policy network $\pi(a|s, \theta)$. The property of the policy network satisfies:

$$\sum_{a \in A} \pi(a|s, \theta) = 1.$$

The value function, also called state-value function, mathematically is an expectation of the reward for subsequent actions determined by the policies, which is used

to evaluate the quality of the state:

$$V_{\pi}(s) = E_{\pi}[G_t | S_t = s] = E_{\pi}[\sum_{k=0}^{\infty} \gamma^k R_{t+k+1} | S_t = s].$$

When the agent starts to train, we want the value function to be maximized by the critic network, and we hope to get as much reward as possible in the shortest possible time, and the value function only depends on the current state. We have another kind of value function, the Q -function, also called action-value function, and differently it depends on the current state and the current action. It is the expectation of the reward that can be obtained in each subsequent state in a given state s , after taking an action a ,

$$Q_{\pi}(s, a) = E_{\pi}[G_t | S_t = s, A_t = a] = E_{\pi}[\sum_{k=0}^{\infty} \gamma^k R_{t+k+1} | S_t = s, A_t = a].$$

We can relate the state-value function and the action-value function by following:

$$V_{\pi}(s) = \sum_a (E_{\pi}[\sum_{k=0}^{\infty} \gamma^k R_{t+k+1} | S_t = s, A_t = a]) \pi(a|s) = \sum_a \pi(a|s) Q_{\pi}(s, a).$$

We can also use $v(s)$ and $q(s, a)$ to represent the approximated state-value function and the approximated action-value function.

If we consider a situation where all the rewards are positive, in this case, whatever the action the agent takes, to increase the expectation of the value function, we need to increase the probability of this token action. Thus, it is extremely necessary to introduce the advantage function:

$$A_{\pi}(s, a) = Q_{\pi}(s, a) - V_{\pi}(s).$$

The advantage function provides a standard for evaluating the policy. If the advantage function is positive, it means that in state s , taking a new action a is better than the action given by the current policy, then we need to increase the probability of this new action in the RL algorithm by maximizing the expectation of the advantage function in the critic network. The critic network in Fig. 14 is used to approximate the advantage function of appendix A. In this scenario, the network's input is made up of 768 neurons (each of the 3 layers utilizes 128 probes to measure two parameters-pressure and velocity-for gathering the necessary information). The actor and critic networks both feature four hidden layers, each matching the size of the network's input. The critic network's output includes a single neuron, whereas the actor network's two paths produce outputs corresponding to the number of jets, which is nine for the mean and nine for the standard deviation in appendix A. For updating the policy network (actor network), generally we choose a type of gradient descent method to find the parameter θ for the network, and it is necessary to calculate the gradient of the objective function,

$$\begin{aligned} \nabla_{\theta} J(\theta) &= \nabla_{\theta} \{E_{S \sim \eta, A \sim \pi} [\log \pi(A|S, \theta_t) q_{\pi}(S, A)]\} \\ &= E_{S \sim \eta, A \sim \pi} [\nabla_{\theta} \log \pi(A|S, \theta_t) q_{\pi}(S, A)]. \end{aligned}$$

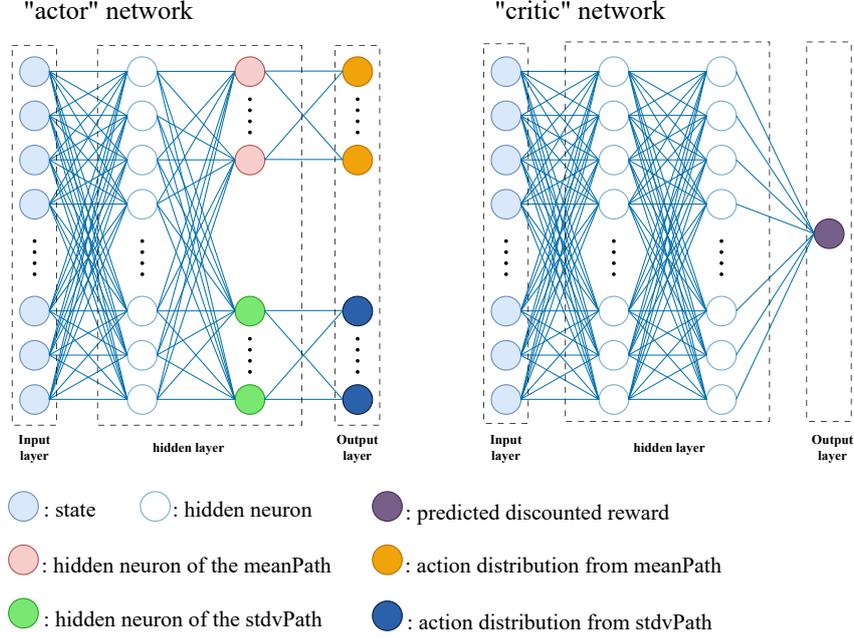


Figure 14: Schematic diagram of actor-critic network used in the PPO algorithm.

Here we introduce a baseline $b(S)$, which is a scalar function of the state space S .

$$\nabla_{\theta} J(\theta) = E_{S \sim \eta, A \sim \pi} \{ \nabla_{\theta} \log \pi(A|S, \theta_t) [q_{\pi}(S, A) - b(S)] \}.$$

Then we need to demonstrate that the introduction of $b(S)$ does not change the gradient value.

$$\begin{aligned}
\nabla_{\theta} J(\theta) &= E_{S \sim \eta, A \sim \pi} [\nabla_{\theta} \log \pi(A|S, \theta_t) b(S)] \\
&= \sum_{s \in S} \eta(s) \sum_{a \in A} \pi(a|s, \theta_t) \nabla_{\theta} \log \pi(a|s, \theta_t) b(s) \\
&= \sum_{s \in S} \eta(s) \sum_{a \in A} \nabla_{\theta} \pi(a|s, \theta_t) b(s) \\
&= \sum_{s \in S} \eta(s) b(s) \sum_{a \in A} \nabla_{\theta} \pi(a|s, \theta_t) \\
&= \sum_{s \in S} \eta(s) b(s) \nabla_{\theta} \sum_{a \in A} \pi(a|s, \theta_t) \\
&= \sum_{s \in S} \eta(s) b(s) \nabla_{\theta} 1 = 0.
\end{aligned}$$

In this case, we choose the value function:

$$b(S) = E_{A \sim \pi} [q(s, A)] = v(s).$$

It is difficult or maybe impossible to calculate the expectation as we said before, so we need to resort to a Monte Carlo estimation, but particularly, in this case, as an

off-line method, the data set we sample from can be changed.

$$E_{x \sim p_0}[x] \approx \bar{f} = \frac{1}{n} \sum_{i=1}^n f(x_i) = \frac{1}{n} \sum_{i=1}^n \frac{p_0(x_i)}{p_1(x_i)} x_i.$$

Usually, for the PPO algorithm, we sample from the old policy, combined with the advantage function, and we have the policy gradient:

$$\nabla_{\theta} J(\theta) = E \left[\frac{\pi_{new}(A|S, \theta)}{\pi_{old}(A|S, \theta)} \nabla_{\theta} \log \pi_{new}(A|S, \theta) (q_{\pi}(S, A) - v_{\pi}(S)) \right].$$

The update of the parameter θ is:

$$\theta_{t+1} = \theta_t + \alpha_{\theta} \frac{\pi_{new}(a_t|s_t, \theta_t)}{\pi_{old}(a_t|s_t, \theta_t)} \nabla_{\theta} \log \pi_{new}(a_t|s_t, \theta_t) (q_t(s_t, a_t) - v_t(s_t)).$$

After updating the parameter, we can get the action distribution at the current observation (states) from the actor network. The distribution of the action can be written as:

$$p(a) = \frac{1}{\xi \sqrt{2\pi}} e^{-\frac{(\alpha - \mu)^2}{2\xi^2}},$$

where ξ is the standard deviation, and μ is the mean value; these two parameters are the outputs of the actor network. Having this distribution, the agent chooses an action as an input into environment.

For avoiding a destructive large weight update, we introduce the PPO clip function:

$$L^{CLIP}(\theta) = E[\min\{r(\theta)\hat{a}_{\pi}, \text{clip}(r(\theta), 1 - \epsilon, 1 + \epsilon)\hat{a}_{\pi}\}],$$

$$\begin{aligned} \hat{a}_{\pi} &= q_{\pi}(S, A) - v_{\pi}(S), \\ r(\theta) &= \frac{\pi_{new}(a_t|s_t, \theta_t)}{\pi_{old}(a_t|s_t, \theta_t)}, \end{aligned}$$

where \hat{a}_{π} is the approximate advantage function. There are two possible cases depending on the positivity of the advantage function. The first is when the advantage function is positive, when L^{CLIP} turns to:

$$L^{CLIP}(\theta) = E[\min\{r(\theta), (1 + \epsilon)\}\hat{a}_{\pi}].$$

In this case, the object is to maximize $L^{CLIP}(\theta)$, giving the approximate advantage function, thus we need to maximize the ratio $r(\theta)$ between the new policy and the old policy, but if the difference is large, it makes no sense to the physical environment. So, with this limiting function, the ratio will reach a ceiling $(1 + \epsilon)\hat{a}_{\pi}$. The second case is when the advantage function is negative, where the PPO clip function becomes:

$$L^{CLIP}(\theta) = E[\max\{r(\theta), (1 - \epsilon)\}\hat{a}_{\pi}].$$

B Hyperparameters PPO

The following are the hyperparameters used for training the PPO agent in this study:

Minibatch size	64
Experience horizon	64
Discount factor	0.97
Clip factor	0.5
Entropy loss weight	0.03
Number of epochs	3
Learning rate	0.001

Table 2: PPO hyperparameters.

References

- [1] O. M. AAMO AND M. KRSTIĆ, *Feedback control of particle dispersion in bluff body wakes*, International Journal of Control, 77 (2004), pp. 1001–1018.
- [2] K. ARULKUMARAN, M. P. DEISENROTH, M. BRUNDAGE, AND A. A. BHARATH, *Deep reinforcement learning: A brief survey*, IEEE Signal Processing Magazine, 34 (2017), pp. 26–38.
- [3] S. AUBRUN, A. LEROY, AND P. DEVINANT, *A review of wind turbine-oriented active flow control strategies*, Experiments in Fluids, 58 (2017), pp. 1–21.
- [4] E. BALLINI, A. S. CHIAPPA, AND S. MICHELETTI, *Reducing the drag of a bluff body by deep reinforcement learning*, arXiv, 2305.03647 (2023).
- [5] S. L. BRUNTON AND B. R. NOACK, *Closed-loop turbulence control: Progress and challenges*, Applied Mechanics Reviews, 67 (2015), p. 050801.
- [6] S. L. BRUNTON, B. R. NOACK, AND P. KOUMOUTSAKOS, *Machine learning for fluid mechanics*, Annual Review of Fluid Mechanics, 52 (2020), pp. 477–508.
- [7] M. A. BUCCI, O. SEMERARO, A. ALLAUZEN, G. WISNIEWSKI, L. CORDIER, AND L. MATHELIN, *Control of chaotic systems by deep reinforcement learning*, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences, 475 (2019), p. 20190351.
- [8] E. BØHN, E. M. COATES, S. MOE, AND T. A. JOHANSEN, *Deep reinforcement learning attitude control of fixed-wing UAVs using proximal policy op-*

- timization*, in 2019 International Conference on Unmanned Aircraft Systems (ICUAS), IEEE, 2019, pp. 523–533.
- [9] H. CHOI, W.-P. JEON, AND J. KIM, *Control of flow over a bluff body*, Annual Review of Fluid Mechanics, 40 (2008), pp. 113–139.
- [10] H. CHOWDHURY, B. LOGANATHAN, Y. WANG, I. MUSTARY, AND F. ALAM, *A study of dimple characteristics on golf ball drag*, Procedia Engineering, 147 (2016), pp. 87–91. The Engineering of SPORT 11.
- [11] L. CHUAN EUN, A. SHAKRINE MOHD RAFIE, S. WIRIADIDJAJA, AND O. FARUQI MARZUKI, *An overview of passive and active drag reduction methods for bluff body of road vehicles*, International Journal of Engineering & Technology, 7 (2018), pp. 53–56.
- [12] S. COLABRESE, K. GUSTAVSSON, A. CELANI, AND L. BIFERALE, *Flow navigation by smart microswimmers via reinforcement learning*, Physical Review Letters, 118 (2017), p. 158004.
- [13] X. DONG, H. HONG, X. DENG, W. ZHONG, AND G. HU, *Surrogate model-based deep reinforcement learning for experimental study of active flow control of circular cylinder*, Physics of Fluids, 35 (2023), p. 105147.
- [14] Y. DUAN, X. CHEN, R. HOUTHOOFT, J. SCHULMAN, AND P. ABBEEL, *Benchmarking deep reinforcement learning for continuous control*, in Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML'16, JMLR.org, 2016, pp. 1329–1338.
- [15] K. DURAISAMY, G. IACCARINO, AND H. XIAO, *Turbulence modeling in the age of data*, Annual Review of Fluid Mechanics, 51 (2019), pp. 357–377.
- [16] T. DURIEZ, S. L. BRUNTON, AND B. R. NOACK, *Machine Learning Control – Taming Nonlinear Dynamics and Turbulence*, Springer International Publishing, 2017.
- [17] 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*, Proceedings of the National Academy of Sciences, 117 (2020), pp. 26091–26098.
- [18] M. GAD-EL HAK, *Flow Control: Passive, Active, and Reactive Flow Management*, Cambridge University Press, Aug. 2000.
- [19] N. GAUTIER, J.-L. AIDER, T. DURIEZ, B. R. NOACK, M. SEGOND, AND M. ABEL, *Closed-loop separation control using machine learning*, Journal of Fluid Mechanics, 770 (2015), pp. 442–457.

- [20] T. HEINEMANN, M. SPRINGER, H. LIENHART, S. KNIESBURGES, C. OTHMER, AND S. BECKER, *Active flow control on a 1:4 car model*, *Experiments in Fluids*, 55 (2014), pp. 1–11.
- [21] S. KASMAIEE, M. TADJFAR, AND S. KASMAIEE, *Machine learning-based optimization of a pitching airfoil performance in dynamic stall conditions using a suction controller*, *Physics of Fluids*, 35 (2023), p. 095121.
- [22] S. KRAJNOVIĆ AND J. FERNANDES, *Numerical simulation of the flow around a simplified vehicle model with active flow control*, *International Journal of Heat and Fluid Flow*, 32 (2011), pp. 192–200.
- [23] R. LANGTRY AND F. MENTER, *Correlation-based transition modeling for unstructured parallelized computational fluid dynamics codes*, *AIAA Journal*, 47 (2009), pp. 2894–2906.
- [24] C. LECLERCQ, F. DEMOURANT, C. POUSSOT-VASSAL, AND D. SIPP, *Linear iterative method for closed-loop control of quasiperiodic flows*, *Journal of Fluid Mechanics*, 868 (2019), pp. 26–65.
- [25] C. LEE, J. KIM, D. BABCOCK, AND R. GOODMAN, *Application of neural networks to turbulence control for drag reduction*, *Physics of Fluids*, 9 (1997), pp. 1740–1747.
- [26] J. LI AND M. ZHANG, *Reinforcement-learning-based control of confined cylinder wakes with stability analyses*, *Journal of Fluid Mechanics*, 932 (2022), p. A44.
- [27] Y. LI, J. CHANG, C. KONG, AND W. BAO, *Recent progress of machine learning in flow modeling and active flow control*, *Chinese Journal of Aeronautics*, 35 (2022), pp. 14–44.
- [28] Y. MAO, S. ZHONG, AND H. YIN, *Active flow control using deep reinforcement learning with time delays in Markov decision process and autoregressive policy*, *Physics of Fluids*, 34 (2022), p. 053602.
- [29] V. MNIH, K. KAVUKCUOGLU, D. SILVER, A. A. RUSU, J. VENESS, M. G. BELLEMARE, A. GRAVES, M. RIEDMILLER, A. K. FIDJELAND, G. OSTROVSKI, S. PETERSEN, C. BEATTIE, A. SADIK, I. ANTONOGLU, H. KING, D. KUMARAN, D. WIERSTRA, S. LEGG, AND D. HASSABIS, *Human-level control through deep reinforcement learning*, *Nature*, 518 (2015), pp. 529–533.
- [30] R. PEREIRA, W. TIMMER, G. DE OLIVEIRA, AND G. VAN BUSSEL, *Design of HAWT airfoils tailored for active flow control*, *Wind Energy*, 20 (2017), pp. 1569–1583.

- [31] J. RABAUULT, M. KUCHTA, A. JENSEN, U. RÉGLADE, AND N. CERARDI, *Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control*, *Journal of Fluid Mechanics*, 865 (2019), pp. 281–302.
- [32] J. RABAUULT AND A. KUHNLE, *Accelerating deep reinforcement learning strategies of flow control through a multi-environment approach*, *Physics of Fluids*, 31 (2019), p. 094105.
- [33] J. RABAUULT, F. REN, W. ZHANG, H. TANG, AND H. XU, *Deep reinforcement learning in fluid mechanics: A promising method for both active flow control and shape optimization*, *Journal of Hydrodynamics*, 32 (2020), pp. 234–246.
- [34] C. RAIBAUDDO, P. ZHONG, B. R. NOACK, AND R. J. MARTINUZZI, *Machine learning strategies applied to the control of a fluidic pinball*, *Physics of Fluids*, 32 (2020), p. 015108.
- [35] F. REN, J. RABAUULT, AND H. TANG, *Applying deep reinforcement learning to active flow control in weakly turbulent conditions*, *Physics of Fluids*, 33 (2021), p. 037121.
- [36] F. REN, C. WANG, AND H. TANG, *Active control of vortex-induced vibration of a circular cylinder using machine learning*, *Physics of Fluids*, 31 (2019), p. 093601.
- [37] ———, *Bluff body uses deep-reinforcement-learning trained active flow control to achieve hydrodynamic stealth*, *Physics of Fluids*, 33 (2021), p. 093602.
- [38] J. SCHMIDHUBER, *Deep learning in neural networks: An overview*, *Neural Networks*, 61 (2015), pp. 85–117.
- [39] J. SCHULMAN, F. WOLSKI, P. DHARIWAL, A. RADFORD, AND O. KLIMOV, *Proximal policy optimization algorithms*, arXiv, 1707.06347 (2017).
- [40] S. SCOTT COLLIS, R. D. JOSLIN, A. SEIFERT, AND V. THEOFILIS, *Issues in active flow control: theory, control, simulation, and experiment*, *Progress in Aerospace Sciences*, 40 (2004), pp. 237–289.
- [41] A. SEIFERT, O. STALNOV, D. SPERBER, G. ARWATZ, V. PALEI, S. DAVID, I. DAYAN, AND I. FONON, *Large Trucks Drag Reduction using Active Flow Control*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2009, pp. 115–133.
- [42] V. SEKAR, Q. JIANG, C. SHU, AND B. C. KHOO, *Fast flow field prediction over airfoils using deep learning approach*, *Physics of Fluids*, 31 (2019), p. 057103.

- [43] A. F. SHAHRABI, *The control of flow separation: Study of optimal open loop parameters*, Physics of Fluids, 31 (2019), p. 035104.
- [44] S. P. SINGH AND S. MITTAL, *Flow past a cylinder: shear layer instability and drag crisis*, International Journal for Numerical Methods in Fluids, 47 (2005), pp. 75–98.
- [45] R. S. SUTTON AND A. G. BARTO, *Reinforcement Learning: An Introduction*, A Bradford Book, Cambridge, MA, USA, 2018.
- [46] H. TANG, J. RABAULT, A. KUHNLE, Y. WANG, AND T. WANG, *Robust active flow control over a range of Reynolds numbers using an artificial neural network trained through deep reinforcement learning*, Physics of Fluids, 32 (2020), p. 053605.
- [47] N. THUEREY, K. WEISSENOW, L. PRANTL, AND X. HU, *Deep learning methods for Reynolds-averaged Navier–Stokes simulations of airfoil flows*, AIAA Journal, 58 (2020), pp. 25–36.
- [48] S. VERMA, G. NOVATI, AND P. KOUMOUTSAKOS, *Efficient collective swimming by harnessing vortices through deep reinforcement learning*, Proceedings of the National Academy of Sciences, 115 (2018), pp. 5849–5854.
- [49] J. VIQUERAT, J. RABAULT, A. KUHNLE, H. GHRAIEB, A. LARCHER, AND E. HACHEM, *Direct shape optimization through deep reinforcement learning*, Journal of Computational Physics, 428 (2021), p. 110080.
- [50] P. WEN AND W. QIU, *Investigation of drag crisis phenomenon using CFD methods*, Applied Ocean Research, 67 (2017), pp. 306–321.
- [51] Y. ZHANG, H. ZHANG, X. YI, B. WU, X. GUAN, AND J. XIONG, *Analysis method and experimental study of ice accumulation detection signal based on lamb waves*, Chinese Journal of Aeronautics, 37 (2024), pp. 388–403.
- [52] Z. ZHAO, Z. LUO, J. LIU, X. DENG, W. PENG, AND S. LI, *Flight test of aircraft three-axis attitude control without rudders based on distributed dual synthetic jets*, Chinese Journal of Theoretical and Applied Mechanics, 54 (2022), pp. 1220–1228.
- [53] C. ZHENG, T. JI, F. XIE, X. ZHANG, H. ZHENG, AND Y. ZHENG, *From active learning to deep reinforcement learning: Intelligent active flow control in suppressing vortex-induced vibration*, Physics of Fluids, 33 (2021), p. 063607.
- [54] H. ZHU, T. TANG, H. ZHAO, AND Y. GAO, *Control of vortex-induced vibration of a circular cylinder using a pair of air jets at low Reynolds number*, Physics of Fluids, 31 (2019), p. 043603.

MOX Technical Reports, last issues

Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

- 106/2024** Brunati, S.; Bucelli, M.; Piersanti, R.; Dede', L.; Vergara, C.
Coupled Eikonal problems to model cardiac reentries in Purkinje network and myocardium
- 105/2024** Bartsch, J.; Barakat, A.A.; Buchwald, S.; Ciaramella, G.; Volkwein, S.; Weig, E.M.
Reconstructing the system coefficients for coupled harmonic oscillators
- 104/2024** Cerrone, D.; Riccobelli, D.; Vitullo, P.; Ballarin, F.; Falco, J.; Acerbi, F.; Manzoni, A.; Zunino, P.; Ciarletta, P.
Patient-specific prediction of glioblastoma growth via reduced order modeling and neural networks
- 103/2024** Fois, M.; Gatti, F.; de Falco, C.; Formaggia, L.
A comparative analysis of mesh-based and particle-based numerical methods for landslide run-out simulations
- 102/2024** Bucelli, M.
The lifex library version 2.0
- 101/2024** Bonetti, S.; Corti, M.
Unified discontinuous Galerkin analysis of a thermo/poro-viscoelasticity model
- 100/2024** Farenga, N.; Fresca, S.; Brivio, S.; Manzoni, A.
On latent dynamics learning in nonlinear reduced order modeling
- 99/2024** Ragni, A.; Masci, C.; Paganoni, A. M.
Analysis of Higher Education Dropouts Dynamics through Multilevel Functional Decomposition of Recurrent Events in Counting Processes
- 98/2024** Castiglione, C.; Arnone, E.; Bernardi, M.; Farcomeni, A.; Sangalli, L.M.
PDE-regularised spatial quantile regression
- 97/2024** Ferro, N.; Mezzadri, F.; Carbonaro, D.; Galligani, E.; Gallo, D.; Morbiducci, U.; Chiastra, C.; Perotto, S.
Designing novel vascular stents with enhanced mechanical behavior through topology optimization of existing devices