

Applying deep reinforcement learning to active flow control in weakly turbulent conditions

Cite as: Phys. Fluids **33**, 037121 (2021); <https://doi.org/10.1063/5.0037371>

Submitted: 13 November 2020 • Accepted: 23 February 2021 • Published Online: 19 March 2021

 Feng Ren,  Jean Rabault and  Hui Tang



View Online



Export Citation



CrossMark

ARTICLES YOU MAY BE INTERESTED IN

[Robust active flow control over a range of Reynolds numbers using an artificial neural network trained through deep reinforcement learning](#)

Physics of Fluids **32**, 053605 (2020); <https://doi.org/10.1063/5.0006492>

[Accelerating deep reinforcement learning strategies of flow control through a multi-environment approach](#)

Physics of Fluids **31**, 094105 (2019); <https://doi.org/10.1063/1.5116415>

[From active learning to deep reinforcement learning: Intelligent active flow control in suppressing vortex-induced vibration](#)

Physics of Fluids **33**, 063607 (2021); <https://doi.org/10.1063/5.0052524>

APL Machine Learning

Open, quality research for the networking communities

Now Open for Submissions

LEARN MORE



Applying deep reinforcement learning to active flow control in weakly turbulent conditions

Cite as: Phys. Fluids **33**, 037121 (2021); doi: [10.1063/5.0037371](https://doi.org/10.1063/5.0037371)

Submitted: 13 November 2020 · Accepted: 23 February 2021 ·

Published Online: 19 March 2021





View Online



Export Citation



CrossMark

Feng Ren,¹  Jean Rabault,^{2,3}  and Hui Tang^{1,a)} 

AFFILIATIONS

¹Research Center for Fluid Structure Interactions, Department of Mechanical Engineering, The Hong Kong Polytechnic University, Hong Kong, China

²Department of Information Technology, Norwegian Meteorological Institute, 0371 Oslo, Norway

³Department of Mathematics, University of Oslo, 0851 Oslo, Norway

^{a)} Author to whom correspondence should be addressed: h.tang@polyu.edu.hk

ABSTRACT

Machine learning has recently become a promising technique in fluid mechanics, especially for active flow control (AFC) applications. A recent work [Rabault *et al.*, *J. Fluid Mech.* **865**, 281–302 (2019)] has demonstrated the feasibility and effectiveness of deep reinforcement learning (DRL) in performing AFC over a circular cylinder at $Re = 100$, i.e., in the laminar flow regime. As a follow-up study, we investigate the same AFC problem at an intermediate Reynolds number, i.e., $Re = 1000$, where the weak turbulence in the flow poses great challenges to the control. The results show that the DRL agent can still find effective control strategies, but requires much more episodes in the learning. A remarkable drag reduction of around 30% is achieved, which is accompanied by elongation of the recirculation bubble and reduction of turbulent fluctuations in the cylinder wake. Furthermore, we also perform a sensitivity analysis on the learnt control strategies to explore the optimal layout of sensor network. To our best knowledge, this study is the first successful application of DRL to AFC in weakly turbulent conditions. It therefore sets a new milestone in progressing toward AFC in strong turbulent flows.

Published under license by AIP Publishing. <https://doi.org/10.1063/5.0037371>

I. INTRODUCTION

Active flow control (AFC) is a long-standing topic in fluid mechanics. Using actuators, it alters flow behavior to improve aerodynamic/hydrodynamic performance, such as lift augmentation, drag reduction, flow-induced-vibration suppression, and mixing or thermal convection enhancement. AFC can be either open-loop or closed-loop, depending on whether measured flow information is used to adjust the control. Compared with open-loop controls, well-designed closed-loop controls can be adaptive and effective in a wider range of flow conditions. However, when flow is turbulent involving strong nonlinearity and multiple spatial and temporal scales, it is quite challenging to design suitable closed-loop control laws in an explicit form.

In the past few years, AFC started to benefit from advances in the field of machine learning (ML). Genetic Programming (GP) was probably the first ML technique applied in AFC. For example, Gautier *et al.*¹ applied GP to search explicit control laws for reducing the recirculation zone behind a backwards-facing step. Zhou *et al.*² applied the linear GP to control the dynamics of a turbulent jet and discovered novel wake patterns. Ren *et al.*³ adopted GP-identified control laws to

successfully suppress vortex-induced vibrations in a numerical simulation environment.

Recently, a novel ML technique, deep reinforcement learning (DRL), has been attracting increasing attention in the fluid mechanics community,^{4–6} following its many successes in robotics control⁷ and sophisticated game playing such as Go.⁸ Applications were mainly focused on agile maneuvering and biomimeticism. For example, Reddy *et al.*⁹ used DRL to train a glider to fly autonomously by exploiting thermal currents in sunny weathers. Verma *et al.*¹⁰ studied the locomotion of fish schoolings, and using DRL trained rear fishes to harness energy from the wake of leading fishes. In these studies, owing to the limitations of early DRL algorithms, discretized control was used, where the control space was limited to a few discrete values rather than spanning a continuous range. With continuous efforts from ML community, however, novel DRL algorithms have been developed to overcome such limitations and, in particular, the so-called “policy-based methods” are now well suited to continuous-control problems.

By applying a policy-based method, i.e., the proximal policy optimization (PPO) method that is now regarded as one of the state-of-the-art methods for continuous control,^{11,12} Rabault *et al.*¹³ achieved

a drag reduction of approximately 8% for a circular cylinder immersed in a laminar channel flow at a diameter-based Reynolds number $Re = 100$, using a pair of anti-phase jets that are issued transversely from the top and bottom of the cylinder. To speed up the simulation-based training, Rabault and Kuhnle¹⁴ also proposed a multi-environment approach, which opens the way to performing DRL-trained controls at higher Reynolds numbers. Following this work, Tang *et al.*¹⁵ were able to design a robust DRL controller, which can effectively control the flow around the cylinder in the range of Reynolds numbers from 60 to 400. Compared to other control strategies, such as the phasor control of Pastoor *et al.*,¹⁶ DRL is more flexible in that it can adapt both the phase and amplitude of the control signal dynamically as the flow configuration evolves, as well as not being limited to some specific forms of forcing, such as sinusoidal forcing.

Given these recent results, two main milestones remain to be demonstrated. First, one needs to fully qualify DRL for control applications in turbulent conditions. Some works already suggest that DRL can satisfactorily control chaotic systems (see, for example, the DRL control of the Kuramoto–Sivashinsky equation presented by Bucci *et al.*,¹⁷ and the control of the Lorentz attractor discussed in their Appendix by Beintema *et al.*¹⁸), but no clear evidence has been presented yet in the context of the control of the full Navier–Stokes equations. Second, the present 2D results with a few actuation locations need to be extended to both larger 2D systems with many actuators, and fully 3D systems. While this may seem challenging, some promising results in this direction have been presented recently, such as the work by Belus *et al.*¹⁹ In this work, the authors demonstrated that an adequate network architecture can take advantage of physical invariants of the system in order to perform effective control at a constant training cost for an arbitrary large number of actuators. Another recent work that suggests that DRL may be effective in more challenging configurations is the experimental work of Fan *et al.*²⁰ In this work, the authors built a setup with two small rotating cylinders located behind a larger cylinder. The aim is to reduce drag by adjusting the rotation rate of the small control cylinders. The authors present satisfactory results, though the final control strategy (close to constant opposite rotation rate for both control cylinders) is of relatively low complexity.

Therefore, we aim in this study to further push the boundary of DRL applications to AFC by conducting more challenging control in weakly turbulent conditions. Specifically, we investigate the same AFC problem as in Rabault *et al.*¹³ and Tang *et al.*,¹⁵ but at a higher Reynolds number $Re = 1000$, where the cylinder wake becomes turbulent and difficult to control. Indeed, while the flow up to $Re = 400$ as studied by Tang *et al.*¹⁵ presents some weakly chaotic instabilities in the cylinder wake, something similar to what is described by other works,^{21,22} it is still dominated by a few well-established harmonic peaks as evidenced in the energy spectrum (see Fig. 1). By contrast, the energy spectrum at $Re = 1000$ is dominated by the $-5/3$ energy cascade and has no well-defined harmonic peaks. This corresponds well to the changes in flow characteristics discussed by Panton.²³ Hence, the flow at $Re = 1000$ has more pronounced turbulent properties that make it more challenging to control than what was presented previously in the literature. Our results show that the DRL-trained AFC not only continues to perform well in this weak turbulent condition, but also achieves a remarkable drag reduction of around 30%.




FIG. 1. Comparison of the turbulent kinetic energy spectra obtained 0.75D downstream from the cylinder at the centerline of the channel, at $Re = 100, 400$, and 1000 (the flow configuration is described in Fig. 2 and Sec. II A). A clear difference is visible among these three cases. While the flow at $Re = 100$ is entirely dominated by the pseudo-harmonic vortex shedding and a few harmonics, the flow at $Re = 400$ presents both a richer frequency content and a weak (i.e., much lower than the frequency peaks) $-5/3$ energy cascade. By contrast, the flow at $Re = 1000$ does not show well-defined harmonic peaks, and the $-5/3$ energy cascade becomes the dominant feature.

II. METHODOLOGY

A. Flow configuration

In the present work, we adopt a flow configuration similar to that in Rabault *et al.*,¹³ i.e., a 2D flow past a circular cylinder, except for the Reynolds number being increased from 100 to 1000 to consider possible turbulence. More specifically, as sketched in Fig. 2(a), a circular cylinder of diameter D is located in the centerline of a narrow channel, $2D$ downstream of the inlet boundary and about $20D$ upstream of the outlet boundary. The incoming flow from the inlet boundary has a parabolic velocity profile. The Reynolds number is then defined based on the mean incoming velocity U and the cylinder diameter D as $Re = UD/\nu$, where ν is kinematic viscosity of the fluid. In the following, all results will be presented in nondimensional form, nondimensionalized with combinations of U, D , and a reference time $T = D/U$.

An array of velocity sensors is employed to perceive the flow environment. As illustrated by orange dots around and behind the cylinder in Fig. 2(a), in total 151 sensors are used, each providing two




FIG. 2. Schematics of the computational domain, boundary conditions and layout of velocity sensor array (a) and the cylinder with a pair of anti-phase jets (b). A sinusoidal velocity profile is applied to the jets.

time-dependent signals, i.e., the streamwise and transverse velocity components. As having been shown in Rabault *et al.*,¹³ this choice of sensor number provides sufficient flow information to perform good AFC at low Reynolds number, which is thus fixed in the present study.

As actuators for the flow control, a pair of jets being issued transversely are implemented at the top and bottom of the cylinder, each covering an arc of 10°, as depicted in Fig. 2(b). This jet pair operates in anti-phase, realizing a zero net mass flux at anytime. A sinusoidal profile is applied to the jet velocities, so that the no-slip boundary condition (BC) is satisfied at the slot edges. The jet centerline velocity is confined in the range of [−1.62, 1.62], consistent with the range set in Rabault *et al.*¹³

The goal of the DRL agent is to reduce drag and meanwhile to mitigate lift fluctuation. To achieve it, we adopt a similar reward function defined in Rabault and Kuhnle:¹⁴

$$r = -\langle C_D \rangle_S - w \langle |C_L| \rangle_S, \tag{1}$$

where C_D and C_L are the drag and lift coefficients, respectively. $\langle \cdot \rangle_S$ indicates an average over a typical actuation period. w is a weighting factor that weights the contributions of drag and lift fluctuation in the reward. In this study it is set as 1, different from that used in Rabault *et al.*¹³ due to a significant increase in $|C_L|/C_D$ from the $Re = 100$ case to the $Re = 1000$ case. The choice of an absolute value term for lift means that both lift bias and fluctuations are penalized. As discussed in Rabault *et al.*,¹³ the choice of the value of w allows to avoid degenerated control strategies. An additional effect of using a reward function that combines the drag and lift coefficients is to provide different time-scales to the reward value. Indeed, the lift part of the reward reacts fast to actuation, as the blowing is in the spanwise direction, while the drag component takes time to evolve following actuation.

B. Flow solver

In prior studies by Rabault *et al.*¹³ and Tang *et al.*,¹⁵ the flow solver used failed at Reynolds numbers roughly larger than 500. To overcome this, here we adopt a well-established lattice-Boltzmann-method (LBM) code for flow simulation.^{3,24,25} In this code, we use the multi-relaxation-time collision scheme,²⁶ the He–Luo incompressible model,²⁷ as well as a uniform Cartesian mesh in the entire computational domain. The BCs are similar to those in Rabault *et al.*:¹³ a constant parabolic velocity profile is applied at the inlet and a zero pressure condition is applied at the outlet. Both BCs are implemented

using the nonequilibrium extrapolation scheme.²⁸ The halfway bounce-back scheme²⁹ is used to satisfy no-penetration and no-slip BC at the top and bottom walls. As for the cylinder with jets, we apply the double linear interpolation method for curved boundary treatment,³⁰ and the corrected momentum exchange method for hydrodynamic force calculation.³¹

In ML-based AFC, it is vital to reduce the time taken by the flow solver to perform each training simulation, as many such simulations are required to find an effective control strategy. Thus, instead of conducting accurate but time-consuming direct numerical simulations (DNS) for the training, we resort to large eddy simulations (LES) when simulating flows at $Re = 1000$. In LES, fluid viscosity involves both molecular viscosity and eddy viscosity, the latter modeling subgrid-scale dissipation and being derived from a local velocity derivative tensor. In this study, the Vreman model³² is implemented to realize LES, where velocity derivatives are calculated using a second-order finite difference scheme. This model has been successfully implemented in our previous works.^{24,25}

To validate the current flow solver, we conduct simulations using different configurations and make comparisons with prior benchmark results, as summarized in Table I. For the three configurations at $Re = 100$, the intermediate mesh (configuration II) generates C_D and $|C_L|$ values fairly close to those obtained with the finest mesh (configuration I). The maximum C_D and C_L values as well as the vortex shedding Strouhal number St match the benchmark results in Schäfer *et al.*³³ and approximate those in Tang *et al.*¹⁵ For the three configurations at $Re = 1000$, although LES with the coarsest mesh (configuration VI) gives results with errors of 5% in \bar{C}_D and 7% in $|C_L|$ compared with the highly resolved DNS (configuration IV), it has good numerical stability and, more importantly, takes only 2% of the latter’s computational time, showing a very good balance between accuracy and efficiency. In addition, with a similar mesh resolution (overall 445, 938 mesh nodes, and 397 mesh nodes located along the cylinder surface), the case performed using the finite volume method (FVM) based DNS by the commercial software ANSYS-Fluent obtains consistent results as configuration V. Based on these results, the settings in configurations II and VI are adopted for the DRL training at $Re = 100$ and 1000, respectively. At $Re = 1000$, each trial simulation running for a duration of $32T$ only takes about 5 min, using our in-house flow solver accelerated with a NVIDIA K40c graphics processing unit (GPU). These efficient simulations to a great extent make the present DRL-based AFC feasible. On the other hand, once the training

TABLE I. Validation and convergence study. Note: $\delta x = 1$ and $\delta t = 1$ are the lattice unit length and time step, respectively, used in our LBM code.

Re	Configuration	Method	$D/\delta x$	$T/\delta t$	\bar{C}_D	$C_{D,max}$	$ C_L $	$C_{L,max}$	St
100	I	DNS	94	2000	3.204	3.244	0.646	1.021	0.3021
	II	DNS	47	1000	3.200	3.240	0.639	0.999	0.3021
	III	DNS	23.5	500	3.196	3.236	0.608	0.948	0.3030
	Schäfer <i>et al.</i>	DNS				3.22–3.24		0.99–1.01	0.295–0.305
	Tang <i>et al.</i>	DNS				3.230		1.032	0.3020
1000	IV	DNS	282	6000	3.476		2.515		
	V	DNS	141	3000	3.438		2.463		
	VI	LES	70.5	1500	3.293		2.339		
	Fluent	FVM-DNS	126	300	3.421		2.530		




FIG. 3. Snapshots of the flow fields obtained at $Re = 100$ (a) and $Re = 1000$ (b). The vorticity contours are normalized by U/D and scaled in a $[-2, 2]$ range. Vortices are identified using the λ_{cf} criterion³⁴ and enclosed with gray lines. The chaotic features in the flow at $Re = 1000$ is clearly visible.

is done and effective control strategies are identified, the settings in the highly resolved DNS configuration, i.e., configuration V, are used to evaluate the well-trained PPO agents, for which the computational cost is not a big concern.

Figure 3 presents snapshots of typical velocity fields obtained at $Re = 100$ and 1000 . One can observe that, compared with the laminar case at $Re = 100$, the simulation at $Re = 1000$ has revealed some chaotic characteristics: vortices shed from the cylinder lose the spatial-temporal symmetry and strongly interact with the channel walls. This can also be clearly seen from the video associated with Fig. 7(a) (Multimedia view). Some vorticity development is visible along the walls near the cylinder, which is due to the blocking effect of the cylinder on the flow. Indeed, we minimized the domain width in order to reduce the computational costs. While this slightly affects the exact value of the drag coefficient compared with the free flow case, it has no implications on either the ability of the DRL agent to control the flow, or the main findings regarding vortex shedding control. Finally, we observe a clear turbulent cascade behind the cylinder, as revealed in Fig. 1 in Sec. I and discussed further in Sec. III when analyzing the effect of the actuation on the flow (see Fig. 10 there).

C. Deep reinforcement learning

The PPO DRL setup for performing AFC is similar to that in Rabault *et al.*¹³ The closed-loop interaction between the cylinder, i.e., the DRL agent, and the fluid environment is depicted in Fig. 4. The sensor array collects the velocity information at selected locations, i.e., the state, from the simulation. The cylinder uses the jet-pair actuation, i.e., the action, to alter the fluid environment. The performance of the AFC is then evaluated using the reward r [defined in Eq. (1)]. In the present study, the DRL algorithm is implemented with an in-house code using Python. More implementation details can be found in Appendix A.




FIG. 4. Schematic of the DRL loop.

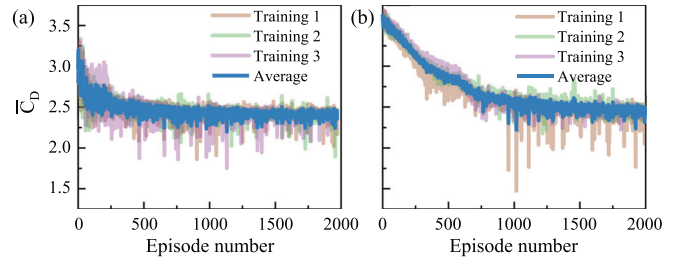


FIG. 5. Learning curves of DRL-trained AFC at $Re = 1000$: starting from scratch (a) and starting from the strategy learnt at $Re = 100$ (b).

III. RESULTS AND DISCUSSIONS

To build confidence on our in-house flow solver and DRL algorithm, we start with benchmarking our DRL-trained AFC at $Re = 100$ against those reported in Rabault *et al.*¹³ Both the identified strategies and the control performance are in good agreement, as detailed in Appendix B. This constitutes an additional validation for the present methods. For the sake of brevity, in the following we only focus on discussing the DRL training at $Re = 1000$. To deal with high-frequency turbulence at this Re , in the learning process we run each episode for $32T$ and adjust the jet actuation five times per T , both significantly longer than those adopted at $Re = 100$. We adopt two different learning strategies for more challenging controls at $Re = 1000$: the first one is to start the learning from a randomly initialized policy, whereas the second one is the so-called “transfer learning,” i.e., to start the learning from the well-trained policy at $Re = 100$. Learning curves are presented in Fig. 5, where three independent trainings are performed for each learning strategy. Very similar learning trends are observed for the trainings under the same strategy, demonstrating the robustness of these two learning strategies. By adopting different strategies, however,

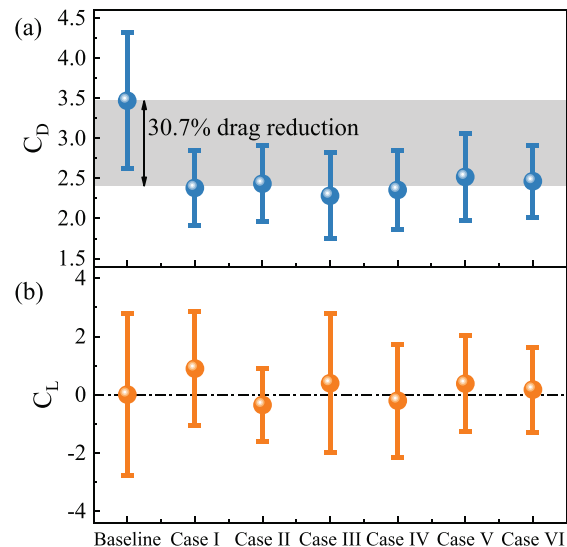


FIG. 6. Comparison of mean drag coefficient (a) and mean lift coefficient (b) among the baseline (i.e., uncontrolled) and six controlled cases that are evaluated using well trained policies and in deterministic mode. The error bars denote the standard deviations of the data. Multimedia view: <https://doi.org/10.1063/5.0037371.1>




FIG. 7. Videos showing the evolution of the cylinder wakes without control (a), with control learnt in case II (b), and with control learnt in case VI (c). Multimedia views: <https://doi.org/10.1063/5.0037371.2>; <https://doi.org/10.1063/5.0037371.3>

an obvious difference in the learning curves is observed: the learnings using transfer learning start with much higher initial \bar{C}_D values than those started from randomly initialized policies. This indicates that the policies learnt at $Re = 100$ do not work well at $Re = 1000$, due to the big difference in the flow dynamics. Nevertheless, all the learning curves eventually approach to similar low \bar{C}_D values, mutually

verifying the effectiveness of learnings using both strategies. Note that, compared to the learnings at $Re = 100$, all these learnings take much more episodes to converge, revealing the difficulty in controlling chaotic flow systems.

Once the trainings are done, the converged policies from each training are evaluated in the deterministic mode using the highly resolved DNS configuration (i.e., configuration V in Table I). As shown in Fig. 6(a), the evident reductions in \bar{C}_D (ranging from 27.4% to 34.2%, with a mean value 30.7%) predicted by DNS simulations confirm that the policies trained using less accurate LES simulations are valid. The slight variation is within expectation, which arises from the eminently random exploration mechanism present in the PPO algorithm and the strong nonlinearity and chaoticity of the turbulent flow considered here. Meanwhile, the fluctuation in C_D is also greatly mitigated by the control. As revealed in Fig. 6(b), lift fluctuations are also reduced by the control. The maximum reduction of 55.2% occurs in case II. However, the control generally leads to asymmetric lift fluctuations, resulting in nonzero \bar{C}_L . This is similar to the findings of several previous works,^{35–37} in which it was found that control leading to an asymmetric flow configuration is what provides the largest drag reduction behind the symmetric fluidic pinball, though at the price of the creation of a biased lift.

Two cases with the smallest lift fluctuations, i.e., cases II and VI, are chosen to reveal more details about control effects of these learnt policies. The evolution of the cylinder wakes in these two controlled cases as well as the uncontrolled case in compared in the videos associated with Fig. 7. The temporal variations of C_D , C_L , and jet velocity u_{jet} presented in Fig. 8 clearly show the irregular feature of the control in




FIG. 8. Temporal variations of drag coefficient (a) and (b), lift coefficient (c) and (d), and jet velocity (e) and (f) for two representative cases evaluated using well trained policies (right) and in deterministic mode: case II (left column) and case VI (right column).




FIG. 9. Contours of mean and turbulence quantities of the flow without (left column) and with (middle column, case II; right column, case VI) control: mean streamwise velocity \bar{u} (a)–(c), mean transverse velocity \bar{v} (d)–(f), streamwise Reynolds stress $\overline{u'u'}$ (g)–(i), transverse Reynolds stress $\overline{v'v'}$ (j)–(l), Reynolds shear stress $\overline{u'v'}$ (m)–(o). The white lines in (a)–(c) are iso-lines of zero streamwise velocity, enclosing the recirculation bubble. The data are obtained through processing 10 000 snapshots of the flow field collected from $t = 100$ to 200, where the control starts at $t = 0$ from a statistically stationary state.

turbulence conditions, contrarily to those at $Re = 100$ as shown in Fig. 15 in Appendix B.

Contours of the first- and second-order turbulence quantities are presented in Fig. 9. Comparison among the mean-streamwise-velocity

contours in Figs. 9(a)–9(c) reveals that the control significantly elongates the recirculation bubble in the wake, by 211% in case II and 195% in case VI, which is accompanied by significant reduction of hydrodynamic drag as revealed earlier. The control also significantly reduces the Reynolds stresses, i.e., $\overline{u'u'}$, $\overline{v'v'}$ and $\overline{u'v'}$, in the wake region, indicating that the wake flow becomes less fluctuating. The turbulent-kinetic-energy (TKE) spectra at two selected locations in the wake, one in the recirculation bubble at $0.75D$ downstream of the cylinder center and the other outside the bubble at $3D$, are presented in




FIG. 10. TKE spectra at two locations in the centerline of the channel, i.e., $0.75D$ (a) and $3D$ (b) downstream of the cylinder center. The frequency is represented by the Strouhal number $St = fD/U$.




FIG. 11. Time-averaged flow field of the half flow domain. The arrowed lines are streamlines, and the background is colored by mean streamwise velocity.




FIG. 12. Distribution of total sensitivity indices (S_T) at all the 151 sensor locations. The sensors are colored with eight equal levels ranging from gray ($S_T=0$) to black ($S_T=0.01$) to reflect their influence on the control output. A larger S_T corresponds to a larger influence.

Fig. 10. It is seen that the peak at $St \approx 0.7$ in the baseline (i.e., uncontrolled) case disappears in the controlled cases, reflecting the fact that the natural vortex formation and shedding process is significantly altered by the jet actuation. In addition, with the control, obvious decrease in the spectra at higher frequencies is observed. This is consistent with the significant reduction of the Reynolds stresses revealed in **Fig. 9**.

From the results shown in **Figs. 8–10**, it is seen that the flow characteristics in cases II and VI are very similar, indicating that different learning strategies eventually give similar control policies. This similarity can also be clearly seen from the video in **Figs. 6(a)–6(c)** (Multimedia view).

To further evaluate the effectiveness of the learnt control policies, we conduct a comparative study similar to that in Bergmann *et al.*,³⁸ which was also adopted in, among others, Rabault *et al.*¹³ and Tang *et al.*¹⁵ Bergmann *et al.*³⁸ suggests that the drag experienced by the cylinder consists of two main components, one arising from a “symmetric base flow” around the cylinder, and the other arising due to the vortex shedding from the cylinder. We estimate the symmetric-base-flow drag by performing a simulation of flow around a half cylinder at $Re=1000$ with a symmetrical boundary condition employed at the centerline of the channel. The resulting flow is shown in **Fig. 11**. In this case, the mean drag coefficient on the half cylinder is read as $\bar{C}_D = 0.927$. Therefore, according to Bergmann *et al.*,³⁸ it is deduced that for the whole cylinder $\bar{C}_D = 1.854$ if the vortex shedding is fully suppressed. This is a 47% reduction from the uncontrolled \bar{C}_D listed in **Table I**. Compare to this idealized value, the maximum reduction of

34% (case III in **Fig. 6**) obtained using the present DRL-trained AFC is still quite remarkable. We anticipate that further improvements can be made by using finer grained actuation, for example, through multiple jet pairs deployed on the cylinder.

Finally, we perform a sensitivity analysis on cases II and VI to explore the optimal layout of sensor network, using the Python library SALib that is based on the method introduced by Saltelli³⁹ and Sobol.⁴⁰ For each case, we generate 10 000 groups of random samples for both the streamwise and transverse velocity components (u and v) collected by the 151 sensors, each ranging between -2 and 2 . Two observations are made from the results shown in **Fig. 12**: first, the sensors placed on the windward side of the cylinder are generally less important; second, the control is more sensitive to v component than u component, except for sensors deployed in the near leeward side of the cylinder. By comparing the distributions of the total sensitivity indices (S_T) in the two cases, one can clearly see the difference in the pattern of most influential sensors located in the mid-wake region. This indicates that the pattern of influential sensors in this region differs with control strategies. Hence, in order to achieve effective control with an essential number of sensors, this sensitivity analysis should be conducted on more cases to reach a statistically meaningful answer, which deserves a further investigation.

IV. SUMMARY AND CONCLUSION

In the present work, we performed the first PPO DRL trained AFC in weakly turbulent conditions, with the aim to reduce the drag

and mitigate lift fluctuations experienced by a circular cylinder at $Re = 1000$. The findings are summarized as follows:

1. At intermediate Reynolds numbers where the flow shows weakly turbulent features, DRL can still find effective control strategies. Due to the much stronger nonlinear flow features involved; however, the learning requires much more episodes to converge than the learning in the laminar flow regime.
2. Both randomly initialized and transfer-learning strategies perform well to reach a similar drag reduction level, i.e., around 30%. The latter strategy does not show obvious advantages over the former strategy as expected, mainly due to the large difference in flow dynamics caused by turbulence.
3. Through analyzing the AFC results, two flow features associated with the drag reduction are identified: first, the recirculation bubble is greatly elongated, similar to what has been observed in the laminar regime. Second, turbulence levels in the wake, especially in the near wake, are significantly reduced by the control.
4. The sensitivity analysis identifies the most influential sensors as well as the influential velocity component, which can greatly help reduce the number of sensors and hence the complexity of ANN architecture for the control.

This work further qualifies DRL as a useful ML tool for solving AFC problems, and sets a new milestone by illustrating the effectiveness of DRL in a case much more complex than previous studies. We anticipate that more relevant studies will be conducted in much stronger turbulent conditions to further progress toward real-world applications.

AUTHORS' CONTRIBUTIONS

F.R. and J.R. contributed equally to this work.

ACKNOWLEDGMENTS

Feng Ren and Hui Tang gratefully acknowledge financial support from the Research Grants Council of Hong Kong under the General Research Fund (Project Nos. 15249316 and 15214418). Jean Rabault acknowledges funding obtained through the Petromaks II project (Grant No. 280625).

APPENDIX A: DEEP REINFORCEMENT LEARNING

Here we present a brief description about the proximal policy optimization (PPO) algorithm. For more details, readers are referred to any of the many discussions on this topic, such as Heess *et al.*,¹² Schulman *et al.*,¹¹ and Rabault *et al.*^{5,13}

In each episode, the PPO agent applies the control policy N times and collects a sequence of state-action-reward combinations, i.e.,

$$\tau = (s_1, a_1, r_1), (s_2, a_2, r_2), (s_t, a_t, r_t), \dots, (s_N, a_N, r_N). \tag{A1}$$

To optimize against a long-term objective, the learning process is driven by a discounted reward

$$R_t = \sum_{t' > t} \gamma^{t'-t} r_{t'}, \tag{A2}$$

where $0 < \gamma < 1$ is a discount factor usually close to 1 (herein set as 0.97), such that later rewards contribute more to the discounted reward.

The policy, π_{Θ} , is modeled by an artificial neural network (ANN) having a set of weights Θ . As shown in Fig. 13, the PPO algorithm uses two sets of ANNs: an “actor” network whose input is the state and output the action, and a “critic” network whose input is the state and output a prediction of the discounted reward. In both ANNs, we use two fully connected hidden layers with 512 neurons in each layer.

In order to perform training, an appropriate loss function must be defined for each ANN. When training the critic network, an intermediate variable, i.e., the “advantage,” is used to evaluate the difference between the predicted and actual discounted rewards

$$\hat{A}_t = R_t - V_{\Theta}(s_t). \tag{A3}$$

Then, the objective of the critic network is to minimize a loss function measuring the discrepancy between the predicted and actual discounted reward, i.e.,

$$J_{critic} = \hat{E}_t(-\hat{A}_t^2), \tag{A4}$$

where \hat{E}_t denotes the empirical expectation over time.

As learning progresses, the PPO agent always attempts to increase its cumulative reward. To achieve this, the actor network is used to generate actions so that the agent can interact with the environment. In return, this network is also trained using the reward information. In the present PPO implementation, we follow the work in Schulman *et al.*,¹¹ where a clipped surrogate objective function is used, i.e.,

$$J_{actor} = \hat{E}_t[\min(q_t(\Theta)\hat{A}_t, \text{clip}(q_t(\Theta), 1 - \epsilon, 1 + \epsilon)\hat{A}_t)], \tag{A5}$$

where $q_t(\Theta) = \pi_{\Theta}(a_t|s_t)/\pi_{old}(a_t|s_t)$ is the ratio of the probability of current policy π_{Θ} in adopting action a_t according to state s_t to the probability of previous policy π_{old} . The clipped term inside the above equation means that $q_t(\Theta)$ is constrained to an interval $[1 - \epsilon, 1 + \epsilon]$, where ϵ is a hyper-parameter set as 0.2 as recommended by Schulman *et al.*¹¹ Therefore, excessively large policy updates, which would make the training process unstable, are avoided.




FIG. 13. General actor-critic setup used in the PPO algorithm.




FIG. 14. Learning process at $Re = 100$: (a) using weighting factor $w = 0.2$, (b) using $w = 1.0$.

When updating the policy, we use the Adam (short for “adaptive moment estimation”) optimizer, which performs better in fast convergence than conventional stochastic gradient descent optimizers.⁴¹ For both ANNs, the learning rate is fixed as 0.001. To deal with continuous control, the actor network does not directly generate actions. Instead, it generates a combination of parameters for a certain probability distribution for actions. In this study, we choose the beta distribution, from which the actions are sampled in a predefined range.⁴²

Once the learning converges and the performance reaches a satisfactory level, deterministic runs can be performed to reveal more details. In these runs, the agent no longer conducts learning from the sampled data. Instead, it directly generates deterministic actions, which come with the highest probability in the distribution, i.e., no random process is involved.

APPENDIX B: DRL CONTROL OF LAMINAR FLOW

Since in the present work different flow solver and DRL implementation are employed, we cross-validate our methods by performing the same AFC problem at $Re = 100$ as in Rabault *et al.*¹³ We adopt the same settings for the learning. The learning curves, performed using two different weighting factors for lift, i.e., $w = 0.2$ and 1.0 , are presented in Fig. 14, all showing good learning trends and converge within about 200 episodes. During the learning progress, each episode, i.e., a complete run of the simulation starting from a uncontrolled, fully developed flow, runs for $24T$, corresponding to about 7.3 vortex shedding periods in the uncontrolled case. During each episode the action is adjusted 60 times according to the latest policy. To mitigate possible instability in the simulation, each adjustment of action is gradually realized using a smooth function. The control policy learnt by the PPO agent is updated every 20 episodes.

Once the training is done, deterministic controls are performed using the learnt policies. Results from three representative cases (i.e., DRL I, II, and III) are presented in Fig. 15. The achieved drag reduction rates are 7.6%, 7.5%, and 7.1%, respectively, close to 8% as reported in Rabault *et al.*¹³ Note that, different weighting factors for lift give different results. As $w = 0.2$ is relatively small (i.e., DRL I and II), the agent pursues larger drag reduction rates. However, as revealed in Figs. 15(g) and 15(h), the jet forcing is significantly biased, i.e., one side of the jet pair always blows and the other side always sucks. As a result, the cylinder will experience




FIG. 15. Three typical control strategies observed in the deterministic run at $Re = 100$: temporal variations of drag coefficient (a)–(c), lift coefficient (d)–(f), and jet velocity (g)–(i). Results shown in the left and middle columns use lift weighting factor $w = 0.2$, and results shown in the right column correspond to $w = 1.0$. The black lines represent the baseline (uncontrolled) case. In each subfigure, the gray background indicates the transient process, whereas the white background indicates the steady-state process.

nonzero mean lift force, as evidenced in Figs. 15(d) and 15(e), which is unexpected and could cause serious problems to the structure. On the other hand, as $w = 1.0$ is relatively large (i.e., DRL III), the agent weights the lift fluctuation more. Although a smaller drag reduction rate is obtained, the undesirable asymmetric lift fluctuation is significantly mitigated, as evidenced in Fig. 15(f). Moreover, the significant reduction in the fluctuation amplitude, compared to that in the baseline case, also suggests that the vortex formation and shedding is well suppressed by the control.

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

- ¹N. Gautier, J.-L. Aider, T. Duriez, B. Noack, M. Segond, and M. Abel, "Closed-loop separation control using machine learning," *J. Fluid Mech.* **770**, 442–457 (2015).
- ²Y. Zhou, D. Fan, B. Zhang, R. Li, and B. R. Noack, "Artificial intelligence control of a turbulent jet," *J. Fluid Mech.* **897**, A27 (2020).
- ³F. Ren, C. Wang, and H. Tang, "Active control of vortex-induced vibration of a circular cylinder using machine learning," *Phys. Fluids* **31**, 093601 (2019).
- ⁴S. L. Brunton, B. R. Noack, and P. Koumoutsakos, "Machine learning for fluid mechanics," *Annu. Rev. Fluid Mech.* **52**, 477–508 (2020).
- ⁵J. Rabault, 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," *J. Hydrodyn.* **32**, 234–246 (2020).
- ⁶F. Ren, H.-B. Hu, and H. Tang, "Active flow control using machine learning: A brief review," *J. Hydrodyn.* **32**, 247–253 (2020).
- ⁷V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, "Human-level control through deep reinforcement learning," *Nature* **518**, 529–533 (2015).
- ⁸D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot *et al.*, "Mastering the game of Go with deep neural networks and tree search," *Nature* **529**, 484 (2016).
- ⁹G. Reddy, A. Celani, T. J. Sejnowski, and M. Vergassola, "Learning to soar in turbulent environments," *Proc. Natl. Acad. Sci.* **113**, E4877–E4884 (2016).
- ¹⁰S. Verma, G. Novati, and P. Koumoutsakos, "Efficient collective swimming by harnessing vortices through deep reinforcement learning," *Proc. Natl. Acad. Sci.* **115**, 5849–5854 (2018).
- ¹¹J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," [arXiv:1707.06347](https://arxiv.org/abs/1707.06347) (2017).
- ¹²N. Heess, T. B. Dhruva, S. Sriram, J. Lemmon, J. Merel, G. Wayne, Y. Tassa, T. Erez, Z. Wang, S. Eslami *et al.*, "Emergence of locomotion behaviours in rich environments," [arXiv:1707.02286](https://arxiv.org/abs/1707.02286) (2017).
- ¹³J. Rabault, M. Kuchta, A. Jensen, U. Réglade, and N. Cerardi, "Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control," *J. Fluid Mech.* **865**, 281–302 (2019).
- ¹⁴J. Rabault and A. Kuhnle, "Accelerating deep reinforcement learning strategies of flow control through a multi-environment approach," *Phys. Fluids* **31**, 094105 (2019).
- ¹⁵H. 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," *Phys. Fluids* **32**, 053605 (2020).
- ¹⁶M. Pastoor, L. Henning, B. R. Noack, R. King, and G. Tadmor, "Feedback shear layer control for bluff body drag reduction," *J. Fluid Mech.* **608**, 161–196 (2008).
- ¹⁷M. A. Bucci, O. Semeraro, A. Allauzen, G. Wisniewski, L. Cordier, and L. Mathelin, "Control of chaotic systems by deep reinforcement learning," *Proc. R. Soc. A* **475**, 20190351 (2019).
- ¹⁸G. Beintema, A. Corbetta, L. Biferale, and F. Toschi, "Controlling Rayleigh-Bénard convection via reinforcement learning," *J. Turbul.* **21**, 585–521 (2020).
- ¹⁹V. Belus, J. Rabault, J. Viquerat, Z. Che, E. Hachem, and U. Réglade, "Exploiting locality and translational invariance to design effective deep reinforcement learning control of the 1-dimensional unstable falling liquid film," *AIP Adv.* **9**, 125014 (2019).
- ²⁰D. 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.* **117**, 26091–26098 (2020).
- ²¹H. Jiang and L. Cheng, "Flow separation around a square cylinder at low to moderate Reynolds numbers," *Phys. Fluids* **32**, 044103 (2020).
- ²²H. Jiang, "Separation angle for flow past a circular cylinder in the subcritical regime," *Phys. Fluids* **32**, 014106 (2020).
- ²³R. L. Panton, *Incompressible Flows*, 4th ed. (Wiley Publishing, 2013).
- ²⁴F. Ren, B. Song, and H. Hu, "Lattice Boltzmann simulations of turbulent channel flow and heat transport by incorporating the Vreman model," *Appl. Therm. Eng.* **129**, 463–471 (2018).
- ²⁵F. Ren, B. Song, Y. Zhang, and H. Hu, "A GPU-accelerated solver for turbulent flow and scalar transport based on the lattice Boltzmann method," *Comput. Fluids* **173**, 29–36 (2018).
- ²⁶D. D'Humières, I. Ginzburg, M. Krafczyk, P. Lallemand, and L. S. Luo, "Multiple-relaxation-time lattice Boltzmann models in three dimensions," *Philos. Trans. R. Soc. A* **360**, 437–451 (2002).
- ²⁷X. He and L. S. Luo, "Lattice Boltzmann model for the incompressible Navier-Stokes equation," *J. Stat. Phys.* **88**, 927–944 (1997).
- ²⁸Z.-L. Guo, C.-G. Zheng, and B.-C. Shi, "Non-equilibrium extrapolation method for velocity and pressure boundary conditions in the lattice Boltzmann method," *Chin. Phys.* **11**, 366 (2002).
- ²⁹X. He, Q. Zou, L.-S. Luo, and M. Dembo, "Analytic solutions of simple flows and analysis of nonslip boundary conditions for the lattice Boltzmann BGK model," *J. Stat. Phys.* **87**, 115–136 (1997).
- ³⁰D. Yu, R. Mei, L.-S. Luo, and W. Shyy, "Viscous flow computations with the method of lattice Boltzmann equation," *Prog. Aerosp. Sci.* **39**, 329–367 (2003).
- ³¹Y. Chen, Q. Cai, Z. Xia, M. Wang, and S. Chen, "Momentum-exchange method in lattice Boltzmann simulations of particle-fluid interactions," *Phys. Rev. E* **88**, 013303 (2013).
- ³²A. Vreman, "An eddy-viscosity subgrid-scale model for turbulent shear flow: Algebraic theory and applications," *Phys. fluids* **16**, 3670–3681 (2004).
- ³³M. Schäfer, S. Turek, F. Durst, E. Krause, and R. Rannacher, "Benchmark computations of laminar flow around a cylinder," *Flow Simulation with High-Performance Computers II* (Springer, 1996), pp. 547–566.
- ³⁴J. Zhou, R. J. Adrian, S. Balachandar, and T. Kendall, "Mechanisms for generating coherent packets of hairpin vortices in channel flow," *J. Fluid Mech.* **387**, 353–396 (1999).
- ³⁵N. Deng, B. R. Noack, M. Morzyński, and L. R. Pastur, "Low-order model for successive bifurcations of the fluidic pinball," *J. Fluid Mech.* **884**, A37 (2020).
- ³⁶C. Raibaudo, P. Zhong, B. R. Noack, and R. J. Martinuzzi, "Machine learning strategies applied to the control of a fluidic pinball," *Phys. Fluids* **32**, 015108 (2020).
- ³⁷G. Y. Cornejo Maceda, B. R. Noack, F. Lusseyran, N. Deng, L. Pastur, and M. Morzynski, "Artificial intelligence control applied to drag reduction of the fluidic pinball," *PAMM* **19**, e201900268 (2019).
- ³⁸M. Bergmann, L. Cordier, and J.-P. Brancher, "Optimal rotary control of the cylinder wake using proper orthogonal decomposition reduced-order model," *Phys. Fluids* **17**, 097101 (2005).
- ³⁹A. Saltelli, "Making best use of model evaluations to compute sensitivity indices," *Comput. Phys. Commun.* **145**, 280–297 (2002).
- ⁴⁰I. M. Sobol, "Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates," *Math. Comput. Simul.* **55**, 271–280 (2001).
- ⁴¹D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014).
- ⁴²P.-W. Chou, D. Maturana, and S. Scherer, "Improving stochastic policy gradients in continuous control with deep reinforcement learning using the beta distribution," in *Proceedings of the 34th International Conference on Machine Learning (IJMLR.org, 2017)*, Vol. 70, pp. 834–843.