

# Reinforcement-learning-based control of confined cylinder wakes with stability analyses

Jichao Li<sup>1</sup> and Mengqi Zhang<sup>1,†</sup>

<sup>1</sup>Department of Mechanical Engineering, National University of Singapore, 117575 Republic of Singapore

(Received 28 March 2021; revised 1 November 2021; accepted 21 November 2021)

This work studies the application of a reinforcement learning (RL)-based flow control strategy to the flow past a cylinder confined between two walls to suppress vortex shedding. The control action is blowing and suction of two synthetic jets on the cylinder. The theme of this study is to investigate how to use and embed physical information of the flow in the RL-based control. First, global linear stability and sensitivity analyses based on the time-mean flow and the steady flow (which is a solution to the Navier–Stokes equations) are conducted in a range of blockage ratios and Reynolds numbers. It is found that the most sensitive region in the wake extends itself when either parameter increases in the parameter range we investigated here. Then, we use these physical results to help design RL-based control policies. We find that the controlled wake converges to the unstable steady base flow, where the vortex shedding can be successfully suppressed. A persistent oscillating control seems necessary to maintain this unstable state. The RL algorithm is able to outperform a gradient-based optimisation method (optimised in a certain period of time) in the long run. Furthermore, when the flow stability information is embedded in the reward function to penalise the instability, the controlled flow may become more stable. Finally, according to the sensitivity analyses, the control is most efficient when the probes are placed in the most sensitive region. The control can be successful even when few probes are properly placed in this manner.

**Key words:** machine learning, wakes, flow control

## 1. Introduction

The flow past a cylinder constitutes one of the classical problems in fluid mechanics. It has intriguing flow phenomena that have attracted generations of researchers hoping to unveil the essential fluid mechanics underneath (Williamson 1996). We consider here a specific type of the wake flow confined between two parallel walls (Schäfer *et al.* 1996). This flow is representative in many industrial and engineering scenarios such as a flow

<sup>†</sup> Email address for correspondence: [mpezmq@nus.edu.sg](mailto:mpezmq@nus.edu.sg)

past dividers in polymer processing and turbulence promoters in the liquid-metal blankets of fusion reactors (Kanaris, Grigoriadis & Kassinos 2011). More importantly, we place the study of the confined wake flow in the context where we want to use a reinforcement learning (RL)-based strategy to control the flow to evaluate the control performance of RL, possibly boosted by the flow physics knowledge we obtain from the stability analyses. Studies on machine-learning-based algorithms in fluid mechanics are burgeoning in recent years (Duraisamy, Iaccarino & Xiao 2019; Brunton, Noack & Koumoutsakos 2020), but works on applying RL to flow control are still relatively few (see Rabault *et al.* 2020; Viquerat, Meliga & Hachem 2021; Garnier *et al.* 2021 for a more complete literature review). There exist some unexplored topics in this field and some of them will be addressed in this work. In the following, we will first summarise the past works on confined and unconfined cylinder wake flows (especially their flow instability) and then review the recent development of machine learning in fluid mechanics (only relevant works will be discussed).

### 1.1. *Cylinder wake flows and their flow instability*

In the confined wake flow, the ratio of the cylinder diameter to the wall distance is termed as the blockage ratio (Chen, Pritchard & Tavener 1995; Sahin & Owens 2004). Coutanceau & Bouard (1977) investigated this flow via experimental visualisation techniques and identified the limits of the Reynolds number ( $Re$ ) range in which the twin vortices exist and adhere stably to the cylinder. Chen *et al.* (1995) showed that the formation of the steady vortex pair at the rear of the cylinder was not owing to the bifurcation of the full dynamic system but instead was probably associated with a bifurcation of a restricted kinematical problem. They identified the Hopf bifurcation point by solving an eigenvalue problem resulting from linearisation and showed that the flow stability was lost through a symmetry-breaking Hopf bifurcation. Anagnostopoulos, Iliadis & Richardson (1996) investigated the flows past a cylinder with three different blockage ratios at  $Re = 106$  (their description implied that they used the cylinder diameter as the reference length). They found that the size of the standing vortices decreased with the blockage ratio before the Hopf bifurcation and the spacing of the vortices decreased with increasing blockage ratio when the wake became unsteady. Sahin & Owens (2004) systematically investigated two-dimensional flows past a confined circular cylinder with different blockage ratios. The neutral stability curve was obtained via the linear global stability analysis. They identified four regions in the parameter space of  $Re$  (based on the cylinder diameter and maximum inlet fluid speed) and the blockage ratio, and each region corresponded to one type of flow motion: steady symmetric flow; symmetric vortex shedding; steady asymmetric flow and asymmetric vortex shedding.

As mentioned above, when the Reynolds number exceeds a critical value, the confined flows experience a Hopf bifurcation from a steady symmetric state towards a time-periodic non-symmetric state. This is similar to a flow past an unconfined cylinder (Provansal, Mathis & Boyer 1987; Sreenivasan, Strykowski & Olinger 1987). Thus, it is necessary and instructive to review the works on the unconfined wake flow past a cylinder. Continuous efforts have been made to understand the mechanism underneath the vortex shedding phenomenon, which is usually unwanted or even harmful. It has been shown that the global instability (Noack & Eckelmann 1994) is responsible for the onset of the vortex shedding process (Jackson 1987). Pier (2002) showed that the two-dimensional time-periodic vortex shedding regime observed in the cylinder wake at moderate Reynolds numbers may be interpreted as a nonlinear global structure. Barkley (2006) studied the stability of a (time-)mean flow and showed that the eigenfrequency of the mean flow tracked almost

exactly the Strouhal number of the (nonlinear) vortex shedding. Through a global weakly nonlinear analysis, Sipp & Lebedev (2007) further confirmed for the cylinder flow that the mean flow was approximately marginally stable and showed that the linear dynamics of the mean flow yielded the frequency of the saturated Stuart–Landau limit cycle. Leontini, Thompson & Hourigan (2010) conducted linear global stability analysis on the mean flows and showed that the mean cylinder wake for  $Re \leq 600$  was marginally stable and the eigenfrequency of the leading global mode was close to the saturated vortex shedding frequency.

In addition to the linear stability analysis, sensitivity analysis has also been applied to unconfined cylinder flows. Based on the insights provided by stability analyses, Strykowski & Sreenivasan (1990) managed to suppress the vortex shedding behind circular cylinders over a limited range of Reynolds numbers by a proper placement of a much smaller cylinder in the near wake of the main cylinder. Their results revealed that this part of the flow is important for flow control. Hill (1992) applied a sensitivity analysis based on the adjoint method (Jameson 1988; Luchini & Bottaro 2014) to the flow past a cylinder and computed the sensitivity of the least stable growth rate to the placement of a second smaller cylinder. The sensitivity analysis reproduced the most sensitive regions that have been experimentally identified by Strykowski & Sreenivasan (1990). Giannetti & Luchini (2007) performed an analysis of the eigenvalue sensitivity to structural perturbations in the linearised governing equations and identified the ‘wavemaker’ region by overlapping the direct and adjoint perturbation modes, which agreed well with the experimental data obtained by Strykowski & Sreenivasan (1990). Using the linear stability theory and the adjoint method, Marquet, Sipp & Jacquin (2008) presented a general theoretical formalism to assess how base-flow modifications alter the flow stability. Boujo & Gallaire (2014) used the adjoint method to identify the regions that were the most sensitive to volume forcing and wall blowing/suction. The control strategies designed based on the sensitivity analyses were proven to be effective via validations using the full nonlinear Navier–Stokes simulations. Furthermore, sensitivity analyses have also been shown to be effective in the control of flow disturbances for the optimal transient growth (Corbett & Bottaro 2001), the noise amplification in a globally stable flat-plate boundary layer (Brandt *et al.* 2011) and other flow control problems. Nevertheless, because the adjoint-based optimal control law is obtained by minimising the cost function such as time-averaged drag or flow fluctuation via the simulations of Navier–Stokes equations, it is still not computationally affordable to use this method in real-time active control of fluid flows. As an alternative, the RL control strategy will be studied in this work.

### 1.2. Reinforcement learning as a flow control strategy

Recently, reinforcement learning (RL), which has been used in some complex systems including automated driving and game playing, has been applied in the field of flow control. Modelling the control as a Markov decision process, an RL-based control agent is trained to take actions (to exert influences on the environment) to maximise the expected cumulative gains (reward) in a period. RL can be treated as a black-box technique from the user side and establishes a control law from scratch. Verma, Novati & Koumoutsakos (2018) used the RL-based control to find an efficient collective swimming strategy of fishes by harnessing vortices. Rabault *et al.* (2019) applied reinforcement learning in active flow control for drag reduction in a confined wake flow at a moderate Reynolds number ( $Re = 100$  based on the averaged velocity and the cylinder diameter), and Rabault & Kuhnle (2019) further presented a multi-environment approach to accelerate the training

of the RL agent. Tang *et al.* (2020) trained the RL agent to achieve a robust control of the drag reduction in the flow past the confined cylinder at multiple Reynolds numbers. Xu *et al.* (2020) used RL-based control to stabilise the wake of the main cylinder by rotating two small cylinders located at two symmetrical positions downstream of the main cylinder. Paris, Beneddine & Dandois (2021) used a stochastic gated input layer in the RL agent to select an optimal subset from some initially placed probes. Ren, Rabault & Tang (2021) performed a follow-up study of Rabault *et al.* (2019) and presented a successful application of the RL control in weakly turbulent conditions ( $Re = 1000$ ) with a drag reduction of 30%. Beintema *et al.* (2020) applied RL in the suppression of Rayleigh–Bénard convection and discussed limitations in controlling unstable and chaotic dynamics. Overall, most of the state-of-the-art works are focused on the validation of RL-based control in two-dimensional applications, which may not be persuasive enough for industrial applications in real-world flows with three-dimensional effects. Fan *et al.* (2020) demonstrated, for the first time, the effectiveness of RL in experimental fluid mechanics by applying it in the drag reduction of circular cylinders in a turbulent flow. The configuration chosen by Fan *et al.* (2020) resulted in a simple control strategy (basically anti-clockwise rotations of control cylinders), so, regarding the use of RL for controlling turbulent flows, there would be much to investigate in more complex cases.

Recent developments of machine learning algorithms applied in fluid mechanics point to a very important line of research where we ought not to entirely rely on a brute-force strategy when designing and applying a machine learning algorithm in a flow problem, but should rather consider embedding some of the most fundamental physical or mathematical constraints or using some prior knowledge in the construction of the algorithm. This idea is drawing increasing attention in the broad fields of physics and engineering (von Rueden *et al.* 2021) as it will significantly reduce the searching space or guide the algorithm to advance in a more physically relevant direction, which will help to converge to the sought solutions more rapidly. For example, in the work of Ling, Kurzawski & Templeton (2016), a deep learning approach to RANS turbulence modelling that embedded Galilean invariance into the network using a higher-order multiplicative layer was presented. This approach ensured that the predicted anisotropy tensor lies on an invariant tensor basis and was shown to have significantly more accurate predictions than a generic neural network that did not have any embedded invariance properties. Raissi, Perdikaris & Karniadakis (2019b) presented a physics-informed deep learning framework that synergistically combines mathematical models and training data, enabling scientific prediction and discovery from incomplete models and incomplete data. Raissi *et al.* (2018) used the fluid mechanics governing equations as regularisation mechanisms in the loss function of the deep learning network and demonstrated that this physics-informed deep learning algorithm is particularly effective for multi-physics problems such as vortex-induced vibrations of cylinders. Similar applications can also be found in discovering turbulence models (Raissi, Babaei & Givi 2019a), estimating hydraulic conductivity in Darcy flows (Tartakovsky *et al.* 2020) and so on. In the case of RL, Belus *et al.* (2019) found it useful to embed translational invariance into the architecture of the RL agent via the control of a one-dimensional depth-integrated falling liquid film. Zeng & Graham (2021) studied the Kuramoto–Sivashinsky equation using RL to minimise the dissipation rate and power cost in the chaotic system. Importantly, they trained the RL in a symmetry-reduced space (Budanur & Cvitanović 2017), showcasing the significance of considering embedding some physical constraint in the RL design. We believe that there still remains a lot to do in the research of RL-based flow control that uses prior knowledge of the flows in the construction of the algorithm. This is the theme of the current work. In particular, we will obtain useful information provided by the flow stability and sensitivity analyses and

use it in the RL-based control policies to suppress vortex shedding. The flow instability mechanism, for example being absolutely or convectively unstable (Huerre & Monkewitz 1990), affects the choice of control strategies. Using the information of flow instability (coupled with linear control theory Kim & Bewley 2007) to modify the mean flow structure can be very efficient even by small-amplitude perturbations. Delaunay & Kaiktsis (1999) found that in an unconfined flow past a cylinder, a slight suction destabilises the wake in the subcritical  $Re$  regime and a slight blowing stabilises the flow in the supercritical regime. As shown by Sahin & Owens (2004), flows at different  $Re$  and blockage ratios have such different characteristics that the optimal control strategy for them may differ. Thus, being aware of some fluid information is helpful to design the control strategy. In fact, in this work, we have experienced that directly applying RL to some challenging flow control problems (e.g. in the range of parameters that are difficult to control) without analysing the flow mechanism may fail. The RL-based flow control by Rabault *et al.* (2019) has shown its limitation in the cases where the drag reduction performance of the control policy becomes unstable (Tang *et al.* 2020). Analysing the instability mechanism may help to shed light on how to improve the control performance of RL and hopefully obtain more effective RL-based control strategies. To the best of our knowledge, there are currently no studies in the literature on reporting how stability or sensitivity analysis can be effectively used in RL-based flow control.

### 1.3. *The position of the current work*

The primary aim of the current work is to explore the application of RL in fluid mechanics by harnessing the stability and sensitivity analyses in the RL-based control of confined wake flows. This parallels the efforts of embedding/using the flow physics in machine learning studies to better leverage the power of the latter to obtain more physically relevant results, exemplified in the important works such as Ling *et al.* (2016), Raissi *et al.* (2019b) among others, as reviewed above.

The flow physics of the confined wake flow will first be investigated. We will reproduce some reported results on the global stability analysis of the confined wake flow as validation of our computations and will also apply the sensitivity analysis (which has not been applied to the confined wake flow) to describe the important flow structures/patterns and discuss their dynamics. These results will serve to guide the design of efficient control strategies in RL, which is the core theme of our study. In this part, a vanilla RL-based control method will first be applied to suppress the vortex shedding at different blockage ratios and Reynolds numbers. The pros and cons of RL-based control in different regimes are analysed. As a comparison, with the flow sensitivity used as *a priori* knowledge (for guiding the probe placement) and the stability information (embedded in the reward design), we will show that the performance of the RL-based control can be improved.

The paper is structured as follows. In § 2, we introduce the confined cylinder wake problem and the control facilities. In § 3, we introduce the methodologies used in this work. The results on the flow stability and sensitivity analyses are reported in § 4. In § 5, we discuss different facets of the RL-based control and present how to use the results of stability and sensitivity analyses to improve the control performance (especially the placement of the probes). Finally, in § 6, we conclude the paper with some discussions.

## 2. Problem formulation

We investigate the wake flow past a two-dimensional cylinder in a confined space (Chen *et al.* 1995; Sahin & Owens 2004), as shown in figure 1. We work with a Cartesian

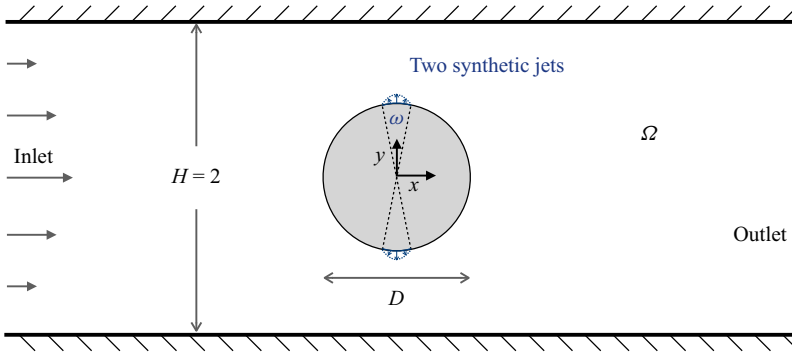


Figure 1. Cylinder symmetrically confined by two parallel no-slip walls. The channel height is  $H = 2.0$  and the cylinder diameter is  $D$ . The synthetic jet width is  $w = \pi/18$ .

coordinate with  $x$  and  $u$  in the horizontal (or streamwise) direction and  $y$  and  $v$  in the vertical (or wall-normal) direction. The length is non-dimensionalised by  $D$  (which is the diameter of the cylinder) and the velocity by  $U_{max}$  (which is the maximum velocity of the parabolic inflow, to be discussed below). The cylinder is placed in the middle of the confined channel and the coordinates of the cylinder centre are  $(0,0,0,0)$ . We use the blockage ratio  $\beta = D/H$  to quantify the degree of blockage. The non-dimensionalised inflow profile is given as

$$u(y) = 1 - y^2, \quad v(y) = 0. \tag{2.1}$$

At the surface of the cylinder and on both sides of the channel, we apply the no-slip boundary condition. On the right-hand side of the computational domain, we impose an outflow condition with  $(p\mathbf{I} - (1/Re)\nabla\mathbf{u}) \cdot \mathbf{n} = 0$ , where  $p$  is the pressure,  $Re$  is the Reynolds number,  $\mathbf{u} = (u, v)^T$  is the velocity and  $\mathbf{n}$  is the outward normal. The Reynolds number is defined as  $Re = U_{max}D/\nu$ , where  $\nu$  is the kinematic viscosity. Note that some of the previous works on the confined cylinder wake flow (Chen *et al.* 1995; Sahin & Owens 2004) have also used this definition of  $Re$  or its variant. A comparison with their results will be made below.

When the Reynolds number exceeds a critical value (see § 4.2), the confined cylinder wake flow starts to fluctuate and exhibit vortex shedding, which may excite a structural vibration. With active suction or blowing of the synthetic jet flows (Glezer 2011; Rabault *et al.* 2019), the stability of the confined cylinder wake can be modified and then the vortex shedding may be damped. As shown in figure 1, two synthetic jets are placed at the top and bottom tips of the cylinder with a width of  $w = \pi/18$ . For both jets, as in Rabault *et al.* (2019), a velocity boundary condition of a cosine-like profile is applied:

$$\left. \begin{aligned} f_{\theta, Q_{upper}}(x, y, t) &= (f_x, f_y)^T = Q_{upper}(t) \frac{\pi}{2\omega R^2} \cos\left(\frac{\pi}{\omega}(\theta - 0.5\pi)\right) (x, y)^T \\ f_{\theta, Q_{lower}}(x, y, t) &= (f_x, f_y)^T = Q_{lower}(t) \frac{\pi}{2\omega R^2} \cos\left(\frac{\pi}{\omega}(\theta - 1.5\pi)\right) (x, y)^T \end{aligned} \right\}, \tag{2.2}$$

where  $\theta$  is the radian angular coordinate of an arbitrary point  $(x, y)$  on the surface of the jets; and  $f_x$  and  $f_y$  are the velocity components along the  $x$  and  $y$  directions, respectively. The flow rates of the upper and lower jets are controlled by changing the scalar values of  $Q_{upper}$  and  $Q_{lower}$ , respectively. The condition  $Q_{upper} + Q_{lower} = 0$  is enforced to ensure that there is no additional mass added to the flow. An effective active control law of the

synthetic jet flow rates is vital to the suppression performance and we will use RL to learn such a control policy.

### 3. Methodologies

#### 3.1. Direct numerical simulation

Flow simulations are performed by solving the two-dimensional (2-D) incompressible Navier–Stokes equations in the computational domain  $\Omega$ :

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \frac{1}{Re} \nabla^2 \mathbf{u}, \quad \nabla \cdot \mathbf{u} = 0. \quad (3.1)$$

The open-source Nek5000 code developed by Fischer, Lottes & Kerkemeier (2017) is used. The spatial discretisation in Nek5000 is based on the spectral element method (SEM). In each spectral element, the velocity space is represented by  $N$ th-order Legendre polynomial interpolants based on tensor-product arrays of Gauss–Lobatto–Legendre (GLL) quadrature points. The SEM has been shown to have little numerical dispersion and dissipation, which is important in the stability analysis. We use the two-step backward differentiation formula for time integration in the unsteady flow simulation with a time step of  $5 \times 10^{-3}$  unit time. Based on the mesh convergence study (§ 4.1), we choose a mesh with 273 elements of order  $N = 7$ , which leads to 17 472 grid points. The flow field is advanced from a certain initial flow field; however, owing to the convective effect in the flow, the initial conditions are not important in the following analysis, and we will analyse and control the period of vortex shedding. The lift force  $F_L$  and drag force  $F_D$  are computed by integrating forces on the cylinder surface, and the lift coefficient and drag coefficient are defined as  $C_l = F_L/0.5DU_{max}^2$  and  $C_d = F_D/0.5DU_{max}^2$ , respectively.

#### 3.2. Linear stability analysis

We will conduct the linear stability analysis to study the flow stability/instability of the confined wake flow. To linearise the incompressible Navier–Stokes equations, the total flow states  $(\mathbf{u}, p)$  are decomposed as a sum of steady base states  $(U_b, P_b)$  and infinitesimal perturbations  $(\tilde{\mathbf{u}}, \tilde{p})$ . Based on the specific problems to be analysed below,  $U_b$  can be chosen as the mean flow or the steady-state solution to the nonlinear Navier–Stokes equations, which are called respectively the mean flow and the base flow in this work. The mean flow can be easily obtained by time-averaging the DNS results. For the base flow, however, when the  $Re$  number is greater than the critical  $Re_c$ , the cylinder wake flow experiences a Hopf bifurcation (Sahin & Owens 2004) and evolves to be a time-periodic non-symmetric state. Thus, the steady-state solution cannot be obtained by a time-marching method. We use the selective frequency damping (SFD) method developed by Akervik *et al.* (2006) to damp the unsteady temporal oscillations via a low-pass filter.

In the linear stability analysis, the perturbations are assumed to be in the form of normal modes  $(\tilde{\mathbf{u}}(x, y, t), \tilde{p}(x, y, t))^T = (\hat{\mathbf{u}}(x, y), \hat{p}(x, y))^T \exp(\sigma t)$  with  $\sigma = \lambda + i\omega$ , where the real part  $\lambda$  and the imaginary part  $\omega$  are the growth rate and frequency of the mode, respectively, and  $(\hat{\mathbf{u}}(x, y), \hat{p}(x, y))^T$  are called shape functions of the variables. After substituting the normal-mode ansatz into the Navier–Stokes equations and linearising them around the base state  $(U_b, P_b)$ , we obtain

$$\sigma \hat{\mathbf{u}} + \nabla \hat{\mathbf{u}} \cdot U_b + \nabla U_b \cdot \hat{\mathbf{u}} = -\nabla \hat{p} + \frac{1}{Re} \nabla^2 \hat{\mathbf{u}}, \quad \nabla \cdot \hat{\mathbf{u}} = 0. \quad (3.2)$$

The boundaries for the linear direct problem are the same as those in the nonlinear Navier–Stokes equations, except that the inlet boundary condition for the velocity is a Dirichlet type with  $\hat{\mathbf{u}} = \mathbf{0}$ .

The equations in (3.2) lead to an eigenvalue problem and the solutions constitute linear global modes of the problem. For clarity, we use  $\hat{\mathbf{q}}$  to represent  $(\hat{\mathbf{u}}(x, y), \hat{p}(x, y))^T$ . Then, the stability analysis can be investigated by solving the following eigenvalue problem:

$$A\hat{\mathbf{q}} = \sigma\hat{\mathbf{q}}, \quad \text{with } A = \begin{pmatrix} -U_b \cdot \nabla - \nabla U_b \cdot + Re^{-1} \nabla^2 & -\nabla \\ \nabla \cdot & \mathbf{0} \end{pmatrix}. \quad (3.3)$$

As one can see, the Jacobian matrix  $A$  depends on the base state  $(U_b, P_b)$ . This global eigenvalue problem can be solved by an iterative approach, and the most popular one is the Arnoldi algorithm (Arnoldi 1951; Saad 1980). It is a time-stepping-based Jacobian-free method (meaning that one does not need to explicitly construct  $A$ ) and has been widely used in the global stability analyses of complex flow problems, such as by Eriksson & Rizzi (1985), Tezuka & Suzuki (2006) and Barkley, Blackburn & Sherwin (2008) among many others, see also the review paper by Theofilis (2011). The major step is the generation of a Krylov subspace  $\mathcal{K}_m$  by marching the linearised Navier–Stokes equations from a certain initial snapshot  $\hat{\mathbf{q}}_0$  at successive equidistant instants of time ( $\Delta t$ ). An orthogonal basis is then generated by the Gram–Schmidt procedure, transforming the large-scale eigenvalue problem to a smaller one of Hessenberg form that can be solved easily.

### 3.3. Sensitivity analysis

The sensitivity analysis based on the adjoint method will also be performed in the current work. It is an important tool which has been extensively applied in flow control and shape optimisation. As in Giannetti & Luchini (2007) and Marquet *et al.* (2008), the adjoint equation of the linearised Navier–Stokes equations reads

$$\sigma^* \hat{\mathbf{u}}^+ - \nabla \hat{\mathbf{u}}^+ \cdot U_b + (\nabla U_b)^T \cdot \hat{\mathbf{u}}^+ = -\nabla \hat{p}^+ + \frac{1}{Re} \nabla^2 \hat{\mathbf{u}}^+, \quad \nabla \cdot \hat{\mathbf{u}}^+ = 0, \quad (3.4)$$

where  $\hat{\mathbf{u}}^+$  and  $\hat{p}^+$  are the adjoint vectors to  $\hat{\mathbf{u}}$  and  $\hat{p}$ , respectively. In principle, the boundary conditions for the adjoint equation are (following Peplinski *et al.* 2014)

$$\hat{\mathbf{u}}^+ = \mathbf{0} \quad \text{at the inlet and walls,} \quad (3.5a)$$

$$\hat{p}^+ \mathbf{n} - Re^{-1} (\nabla \hat{\mathbf{u}}^+) \cdot \mathbf{n} = (U_b \cdot \mathbf{n}) \hat{\mathbf{u}}^+ \quad \text{at the outlet.} \quad (3.5b)$$

The boundary conditions (3.5b) are not supported in the current SEM flow solver (Peplinski *et al.* 2014). Instead, Giannetti & Luchini (2007) explained that because of the particular structure of the base flow, the adjoint mode decays rapidly away from the cylinder; therefore,  $\hat{\mathbf{u}}^+ \rightarrow \mathbf{0}$  can be considered when the outlet is far enough from the cylinder. This method is adopted here because a far enough outlet has been used in the current work (see the geometry in figure 3 in the following) and we have checked that in our simulations, the amplitude of  $\hat{\mathbf{u}}^+$  at the outlet is almost zero.

Similar to the direct problem, the adjoint can be solved by the Arnoldi method. For certain flows, the structural sensitivity analysis can be used to locate the origin of the instability perturbations, called the wavemaker region, which can help to understand the instability mechanism (Pier & Huerre 2001). As shown by Giannetti & Luchini (2007), the






Figure 2. The reinforcement learning framework in flow control: Agent, neural network; Environment, DNS using Nek5000; Action, adjustment of the synthetic jet flow rates; Reward, reduction of the shedding energy; States, spatial velocities.

wavemaker region  $\eta$  can be identified by overlapping the direct eigenvector  $\hat{u}$  and adjoint eigenvector  $\hat{u}^+$ :

$$\eta = \frac{|\hat{u}| |\hat{u}^+|}{\langle \hat{u}, \hat{u}^+ \rangle}. \quad (3.6)$$

It is noted that for the confined cylinder wake flow, we cannot find results on its sensitivity analysis in the literature. Thus, the results reported below on this analysis will be interesting by themselves (especially, the variation of the wavemaker region when  $Re$  or  $\beta$  changes). However, the linear stability analyses of confined cylinder wake flows have been documented (Chen *et al.* 1995; Sahin & Owens 2004) and we will compare our results (based on the SFD base flow) with them for the validation purpose and will further perform the stability analysis based on the mean flow.

### 3.4. Reinforcement learning

Reinforcement learning trains the control agent from scratch by interacting with the environment and maximising the expected cumulative reward. As shown in figure 2, reinforcement learning is composed of three fundamental components, i.e. the agent, the environment and the reward function. The agent usually contains a neural network such as a multilayer perceptron or convolutional neural network that is used to determine the control action  $a_t$  based on the current state of the environment  $s_t$ ; the action is then applied to the environment; and the reward  $r_t$  (evaluating the quality of control actions) is calculated and recorded for network updates.

As a policy-gradient method, the RL agent network (with parameters  $\theta$ , i.e. weights and biases) is trained to find the optimal policy  $\pi_\theta(a_t|s_t)$ , which is the distribution probability of action  $a_t$  (with respect to the states  $s_t$ ) to maximise the expected cumulative reward:  $R_t = \sum_{k>t} \gamma^{(k-t)} r_t$ , where  $\gamma \in (0, 1]$  is a discount factor. The current RL agent uses the proximal policy optimisation (PPO) method developed by Schulman *et al.* (2017) to update the parameters  $\theta$ . PPO is an episode-based actor–critic algorithm. In addition to the network approximating the policy  $\pi$  for action distributions (called ‘actor’), PPO involves a critic network  $V$  to predict the discounted reward with respect to the states  $s_t$ , which is further used to update the actor network. When training the critic network,  $\hat{A}_t = R_t - V(s_t)$  is defined to measure the discrepancy between the predicted and actual discounted rewards, and the loss function to be minimised can be defined as  $L_{critic} = \widehat{\mathbb{E}}_t(\hat{A}_t^2)$ , where  $\widehat{\mathbb{E}}_t$  is the empirical expectation over time. A clipped surrogate objective function is maximised to update the actor network:  $L^{Clip}(\theta) = \widehat{\mathbb{E}}_t[\min(p_t(\theta)\hat{A}_t, \text{clip}(p_t(\theta), 1 - \epsilon, 1 + \epsilon)\hat{A}_t)]$ , where  $p_t(\theta) = \pi_\theta(a_t|s_t)/\pi_{\theta_{old}}(a_t|s_t)$ . The clip term removes the incentive for moving  $p_t$  outside

of the interval  $[1 - \epsilon, 1 + \epsilon]$  ( $\epsilon$  is 0.2 as recommended) and thus prevents an excessively large policy update. More technical details on PPO can be found in Schulman *et al.* (2017). The adam optimiser is used to update the networks and the learning rate is fixed as 0.001.

In this work, for the suppression of the vortex shedding in the confined wake flow past a cylinder, the RL environment is simulated by direct numerical simulation of the wake flow using Nek5000. Referring to the open-source RL-based cylinder flow control repository developed by Rabault *et al.* (2019) based on the Tensorforce library (Kuhnle, Schaarschmidt & Fricke 2017), we present an open-source Python code to interface the Nek5000 simulation environment with the RL agent, which is available as a GitHub repository: [https://github.com/npuljc/RL\\_control\\_Nek5000](https://github.com/npuljc/RL_control_Nek5000). Because the flow rates of two synthetic jets are confined according to  $Q_{lower} + Q_{upper} = 0$ , the RL action is realised by manipulating the flow rate of the upper synthetic jet, and the lower synthetic jet has the same flow rate but opposite direction. Vortex shedding frequencies of the cases studied in this work are approximately  $0.3 \sim 0.4$  (see § 4.3) and we choose a duration of  $\Delta t = 0.2$  between two control actions (corresponding to  $6\% \sim 8\%$  vortex shedding period) to leave a large degree of control freedom. We define a training episode composed of 16 time units, which corresponds to  $4.8 \sim 6.4$  vortex shedding periods, and thus, 80 actions will be taken in each episode. To avoid abrupt changes, we adopt the same strategy used by Rabault *et al.* (2019) to gradually update the jet flow rate after each time step in DNS; that is,  $u_{jet}^{(t+1)} = u_{jet}^{(t)} + 0.1 \times (u_{action} - u_{jet}^{(t)})$ . Following Rabault *et al.* (2019), we define the policy network as a multilayer perceptron with two hidden layers ( $512 \times 512$ ). The RL agent is updated every 20 episodes in the training process. Probe sensors that monitor velocity components in both directions are placed in the flow field to provide environment states for RL.

Rabault *et al.* (2019) discussed that the number of probes has a direct influence on the control performance in RL. We extend this investigation and further determine that the probes are better placed in the regions that are important in the sensitivity analysis. This heuristic approach may also be helpful to be combined with the optimal searching method proposed by Paris *et al.* (2021). Furthermore, the definition of the reward function is important in RL-based control. Rabault *et al.* (Rabault *et al.* 2019; Rabault & Kuhnle 2019) used a drag-based reward function to control (reduce) the drag force on the cylinder. In the current work, to damp the vortex shedding, we use a reward function defined based on the kinetic energy of vortex shedding and the details are discussed in §§ 5.1 and 5.2.

## 4. Results: stability and sensitivity analyses of confined wake flows

### 4.1. Mesh convergence study

As shown in figure 3, the computational domain is defined by three parts for the ease of mesh generation. Here,  $x_1$  and  $x_2$  are the lengths of the two rectangle subdomains in the streamwise direction. The two rectangle subdomains are discretised using  $n_1 \times n_3$  and  $n_2 \times n_3$  element nodes, respectively. The middle square domain with a side length of 2.0 is discretised using an ‘O’-type mesh with  $4 \times n_3 \times n_4$  element nodes (4 here denotes the four compartments delimited by the blue lines in the figure;  $n_3 = 8$  and  $n_4 = 6$ ). To obtain a reasonable computational mesh, we study the influences of both the computational domain size and the mesh resolution on the numerical results. As shown in table 1, five computational sizes are investigated for  $Re = 200$ ,  $\beta = 0.5$  ( $Re = 200$  is in the upper limit of the Reynolds numbers we investigate in this work) and  $E$  is the total number of elements. For each computational size, we compare numerical results by using meshes of two resolution levels (L1 and L2). The L1 mesh is finer, which is generated by doubling




Figure 3. The computational mesh is composed of two rectangle domains and a square domain.

Mesh	$x_1$	$x_2$	$n_1$	$n_2$	$E$	$\overline{C_d}$ by L2	$\overline{C_d}$ by L1
D1	3.0	10.0	6	15	273	2.422105	2.422176
D2	3.0	20.0	6	20	308	2.422132	2.422176
D3	10.0	10.0	15	15	336	2.422096	2.422168
D4	10.0	20.0	15	20	371	2.422139	2.422168
D5	20.0	20.0	20	20	406	2.422140	2.422168

Table 1. Resolution parameters of L2 meshes in five computational domains and the time-averaged  $C_d$  on the confined cylinder with  $\beta = 0.5$  and  $Re = 200$  evaluated by different meshes. Here,  $E$  is the total number of elements in the spectral element method.

the element numbers used in the L2 mesh. The time-averaged  $C_d$  evaluated by different meshes are shown in table 1 and the reference value evaluated by Sahin & Owens (2004) is 2.4245. The choice of five computational domains does not significantly influence the numerical results and all are close to the reference. This means that using the smallest computational domain (D1) is sufficient. Compared with the finer L1 mesh, the L2 mesh merely introduces an error of less than 0.003 %. In the following, we will use D5-L1 as the configuration for all the stability analyses below and in the case of RL control, to reduce the computational burden, we use D1-L2.

#### 4.2. Critical Reynolds numbers

The dynamics of the confined cylinder wake flow is governed by the Reynolds number and the blockage ratio. The vortex shedding in the wake flow occurs through a symmetry-breaking Hopf bifurcation beyond the critical Reynolds number ( $Re_c$ ). In the confined flow past a cylinder,  $Re_c$  varies with the blockage ratio. To determine the vortex shedding region in the  $Re$ - $\beta$  plane, we solve the eigenvalue problem in the linear stability analysis based on the SFD base flow.

As shown in figure 4, two values of  $\beta$  are studied. For the confined flow with  $\beta = 0.5$ , the growth rate of the leading eigenmode increases monotonically with a rise of  $Re$  from 100 to 150. When the Reynolds number is greater than  $Re_c = 123.6$ , the base flow becomes unstable and then the wake starts meandering. However, for the confined flow with the blockage ratio  $\beta = 0.75$ , two critical  $Re_c$  values,  $Re_{c1} = 108.9$  and  $Re_{c2} = 169.6$ , are identified. With the increase of  $Re$  from 100, the flow stability is lost via the first Hopf bifurcation point at  $Re_{c1}$  and the wake vortex starts shedding subsequently. If the  $Re$  further increases, passing the other bifurcation point at  $Re_{c2}$ , the flow becomes stabilised. As shown in figure 5, the confined flows with  $\beta = 0.25$  and  $\beta = 0.5$  merely have one recirculation zone; whereas for the larger blockage ratio  $\beta = 0.75$ , in addition to the recirculation bubble just downstream of the cylinder, two additional recirculation bubbles develop close to the walls further downstream when  $Re$  is large enough.




Figure 4. Eigenvalues of the leading eigenmode at different Reynolds numbers for two blockage ratios  $\beta = 0.5, 0.75$ . The curves are interpolated using data of the scatter points. The grey shade indicates flow instability. Note that that the data points are our computational results and the lines are fitted to guide the eyes.




Figure 5. Vorticity of SFD base flows of different confined cylinder wake flows. The purple dashed lines and black dash-dotted lines are the recirculation zones boundaries of the SFD base flows and time-mean flows, respectively: (a)  $\beta = 0.25$ ; (b)  $\beta = 0.5$  and (c)  $\beta = 0.75$ .

Similar to observations by Sahin & Owens (2004), the recirculation bubbles on the confinement walls become larger with the increase of  $Re$  and lead to the other Hopf bifurcation point at  $Re_{c2}$ .

By examining the critical Reynolds numbers with different blockage ratios, the vortex shedding region is shown in figure 6 as a grey background in the  $Re$ - $\beta$  plane. The region is determined by two branches of critical points, and we have used 13 and 8 data points in the lower branch and upper branch, respectively. The favourable agreement with the results by Chen *et al.* (1995) and Sahin & Owens (2004) demonstrates the good accuracy of the numerical methods and meshes used in this work. The confined flow becomes more stable as  $\beta$  increases up to 0.5. Then, the flow becomes destabilised as the block ratio increases. When  $\beta$  is larger than  $\sim 0.85$ , the flow is again stabilised. For a relatively large  $\beta$  ( $\lesssim 0.85$ ), another stability region can be identified when  $Re$  is sufficiently large (when  $Re \leq 200$ ). It is noted that the flow phenomena are richer and more complex in the upper-right region of the  $Re$ - $\beta$  plane (see figure 5 of Sahin & Owens 2004) and they are not studied here. We will focus on the grey region where vortex shedding occurs and will use RL control to abate it.

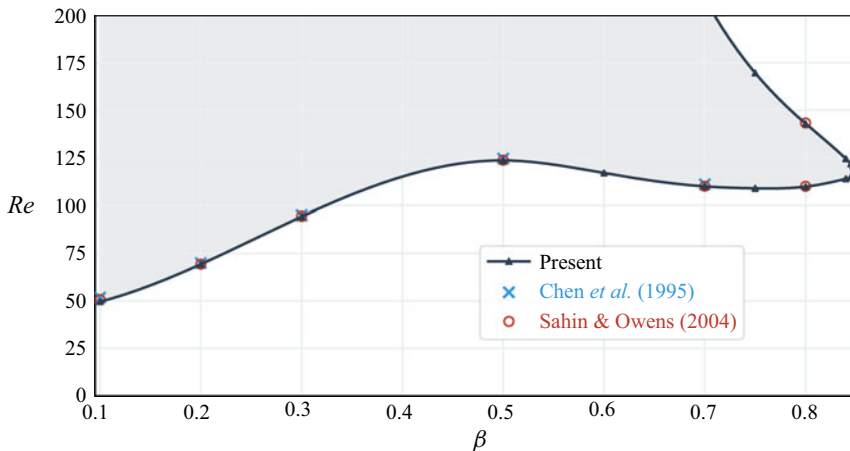


Figure 6. Vortex shedding region (grey) delimited by the neutral stability curve. Flows in the white region are stable. A comparison is made to two previous publications, as shown in the legend.

#### 4.3. Vortex shedding phenomenon

In general, the oscillating Kármán vortex street downstream the cylinder wake occurs owing to the loss of instantaneous reflection symmetry through a Hopf bifurcation. For the unconfined flow past a cylinder, it is suggested by Maurel, Pagneux & Wesfreid (1995) and Noack *et al.* (2003) that the amplitude of the oscillating wake saturates when the time-averaged flow (mean flow) is marginally stable and the mechanism for nonlinear saturation of the oscillating wake flow is the mean flow correction/modification through the formation of Reynolds stresses. The mean flow provides a good profile to predict the shedding frequency of the unconfined cylinder wake (Yang & Zebib 1989; Pier 2002; Barkley 2006). To understand in detail the stability property of the confined cylinder wake, hereafter, we perform the global linear stability analysis of its mean flow.

Figure 7 shows the leading eigenvalues of the confined cylinder wake with  $\beta = 0.25$ ,  $\beta = 0.5$  and  $\beta = 0.75$ . Compared with the stability analysis based on the SFD base flow (black solid lines), the frequencies solved by the linear analysis based on the mean flow (blue dashed lines) agree better with the results of nonlinear DNS (red filled triangles), and the relative discrepancies are almost within 1%. The growth rates of the mean flows are close to zero, which implies that the confined mean flows are marginally stable. The conclusions hold for both the weak and strong confinement cases. This implies that the confined cylinder wake flows approximately have the real-zero imaginary-frequency (RZIF) property (Turton, Tuckerman & Barkley 2015), which is similar to the unconfined cylinder wake. The RZIF property implies that the eigenfrequency of a nonlinearly saturated oscillating flow can be well approximated by a linear analysis based on the time-mean flow (Pier 2002; Barkley 2006; Sipp & Lebedev 2007). Furthermore, the linear stability analysis using the SFD base flow (black lines) generates apparently different results than those using the mean flow and of DNS (except the frequency in the case of  $\beta = 0.75$ ).

To reveal the relationship between the SFD base flow and the mean flow, we perturb and evolve the SFD base flow at  $\beta = 0.25$  at  $Re = 150$  to see how it develops to the saturated state (with vortex shedding). This analysis follows that for the unconfined wake flows in Barkley (2006). As shown in figure 8, because the SFD base flow is unstable, the amplitude of the  $C_l$  oscillation increases and eventually saturates at a periodic vortex shedding state.

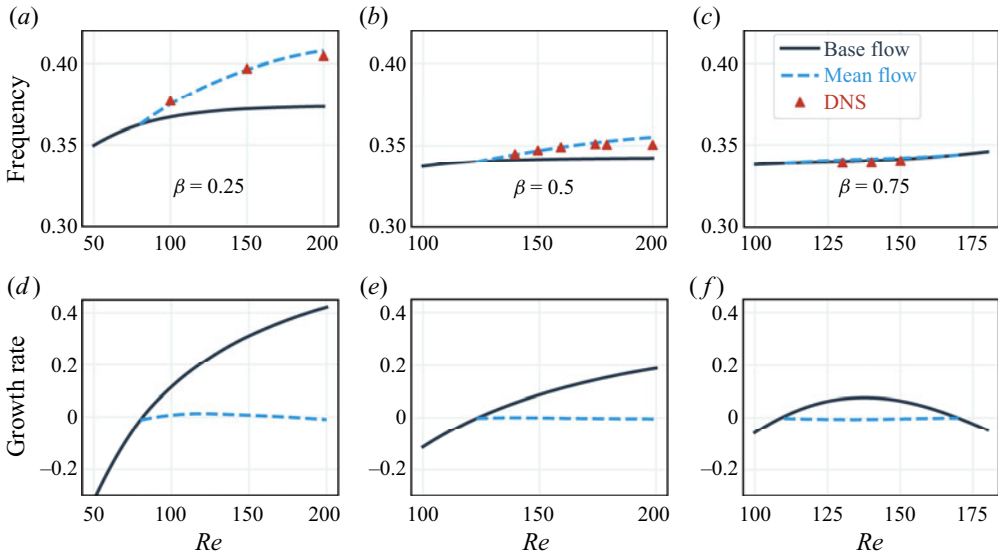


Figure 7. Frequencies and growth rates as a function of  $Re$  in the global linear stability analysis of confined cylinder flows for three values of confinement ratio  $\beta$ .




Figure 8. Relationship between the SFD base flow and the mean flow at  $Re = 150$  with  $\beta = 0.25$ : (a)  $C_l$  as function of  $t$  from the SFD base flow to the saturated flow with vortex shedding; (b) phase diagram of  $C_l$  and  $C_d$  from  $t = 0$  to  $t = 150$ ; (c) growth rates and frequencies of the SFD base flow and the mean flow solved by the linear stability analysis.

The cone-like shape of drag evolution has been explained by Loiseau, Noack & Brunton (2018) based on the sparse identification of nonlinear dynamics, SINDy (Brunton, Proctor & Kutz 2016). Similarly to the unconfined flow past a cylinder, the evolution of the unstable SFD base flow in the confined wake is a nonlinear saturation of oscillations (as shown in panels *a* and *b*), and the selection of vortex shedding amplitude and frequency is based on the marginal stability of the mean flow. In figure 8(c), we show the growth rates and frequencies of the SFD base flow and the saturated mean flow.

#### 4.4. Structural sensitivity

The above section details the frequency and growth rate of the dominant linear mode in the confined cylinder wake flow. Regarding the flow control and manipulation, much more useful information can be obtained by conducting a sensitivity analysis.




Figure 9. Growth rates and frequencies of SFD base flows of confinement cylinder wakes with different blockage ratios and Reynolds numbers.

Structural sensitivity reveals the most sensitive spatial part of the flow to perturbation; this region is traditionally dubbed as a wavemaker region. We investigate the influence of blockage ratios and Reynolds numbers on the wavemaker region in this section. Three Reynolds numbers,  $Re = 115, 150, 185$  are selected, covering subcritical and supercritical cases. The growth rates and frequencies of the leading eigenvalues obtained in the global linear stability analysis are shown in figure 9. When the blockage ratio  $\beta$  is smaller than 0.5, increasing  $\beta$  stabilises the confined cylinder wake, and the leading eigenfrequency decreases simultaneously. For the confined wake flow at  $Re = 115$ , the vortex shedding is even fully suppressed (i.e. the flow becomes stable) when  $0.4 < \beta < 0.6$ . Then, further increasing  $\beta$  destabilises the wake, which becomes unstable again when  $0.6 < \beta < 0.8$ . These results on the growth rate are consistent with those in figure 6. In confined cylinder wakes at  $Re = 150$ , a similar stabilising–destabilising trend is observed when  $0.4 < \beta < 0.6$  but the flow is always unstable with a positive growth rate. A stabilising effect is observed when  $\beta > 0.7$  and the flow becomes stable when  $\beta = 0.8$ . This is because the second critical Reynolds number ( $Re_{c2}$ ) exists in the confined cylinder wakes with  $\beta > 0.7$  and it decreases significantly with the increase of  $\beta$ , see figure 6. For wake flows at  $Re = 185$ , the stabilising effect for  $\beta > 0.7$  is more significant and leads to a more stable flow when  $\beta = 0.8$ . Despite the stability, the wake flow is in an asymmetric status (see figure 10(c) with  $\beta = 0.8$  below), and this phenomenon has been reported by Sahin & Owens (2004). However, the effects of  $Re$  on the leading eigenfrequency are not significant when  $0.4 < \beta < 0.7$ , as shown in figure 9(b). When  $\beta > 0.7$ , the leading eigenfrequency of  $Re = 185$  decreases more significantly.

Next, we show the direct, adjoint modes and wavemaker region in the confined cylinder wake flow with different blockage ratios and Reynolds numbers in figure 10. The amplitudes of adjoint and direct modes for velocity are represented by the red and blue colour maps, respectively. The wavemaker region is shown by nine contour lines using the grey colour map, which illustrates the sensitive domain with  $\eta > 5\% \times \eta^{max}$ , where  $\eta$  is the overlap function in structural sensitive analysis evaluated by (3.6) and  $\eta^{max}$  is the largest  $\eta$  in the computational domain. Similar to unconfined cylinder wakes (Giannetti & Luchini 2007; Marquet *et al.* 2008), the wavemaker region of the confined wake flow is located downstream of the cylinder, and thus to some extent, the upstream domain is less important in terms of flow sensitivity and needs not to be monitored in the suppression of wake vortex shedding. Furthermore, when  $\beta < 0.7$ , the cylinder wake flow with a larger blockage ratio (i.e. a stronger confinement effect) possesses a longer and

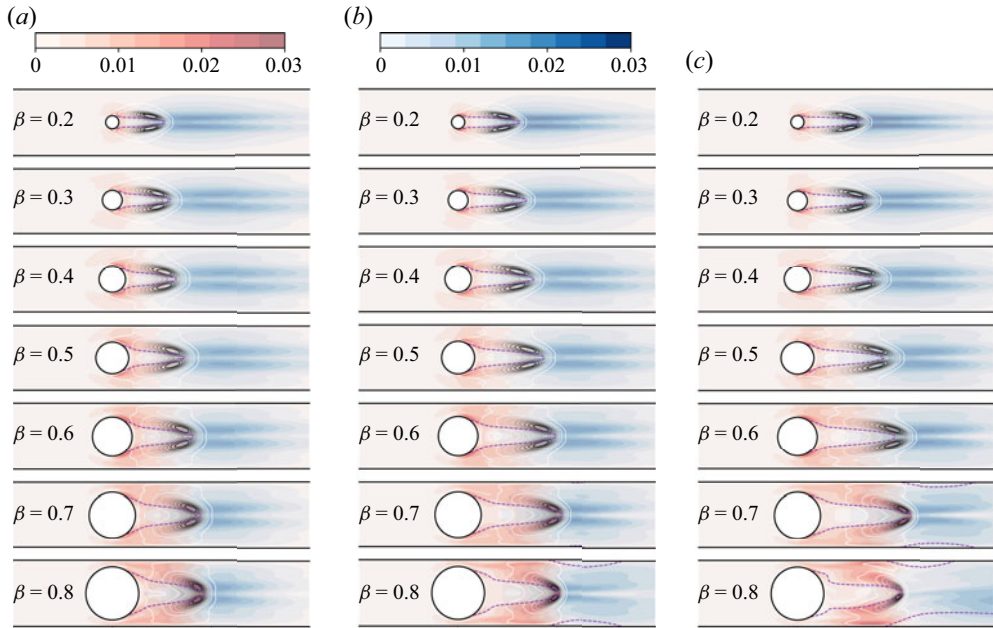


Figure 10. Structural sensitivity analysis of confined cylinder wakes with different blockage ratios and Reynolds numbers. The adjoint mode, direct mode and wavemaker are shown by red, blue and grey colour maps. The recirculation zone is illustrated by the purple dashed line. Note that only part of the domain is shown: (a)  $Re = 115$ ; (b)  $Re = 150$  and (c)  $Re = 185$ .

wider recirculation zone, and correspondingly the wavemaker region also expands and moves downstream. Regarding the length of the recirculation region, Chen *et al.* (1995) also reported longer recirculation zones as a function of  $Re$ , whereas the result of the elongated wavemaker region with  $\beta$  seems not to have been reported in the literature for the confined cylinder wake flow. However, increasing the Reynolds number also leads to longer recirculation zones, and the wavemaker region is pushed downstream (at least for the three  $Re$  values investigated here). Regarding the flow control of the confined cylinder wake flows, the significance of these results is that to more efficiently suppress the vortex shedding, one should monitor the perturbations further downstream from the cylinder when  $\beta$  or  $Re$  increases because the most sensitive region (wavemaker) is located further downstream. Finally, for some wake flows with larger  $\beta$ , i.e.  $Re = 150, \beta = 0.8$  or  $Re = 185, \beta = 0.7, 0.8$ , recirculation bubbles are developed on the confinement walls downstream of the cylinder wake (similar to the results in figure 5), which stabilises the cylinder wake and leads to a second critical  $Re$  as analysed in the previous section.

## 5. Results: RL-based control of confined wake flows

In the above section, we have shown in detail the results of the stability and sensitivity analyses of the confined cylinder wake flow to obtain important flow information such as the wavemaker region. In this section, we analyse the influences of key parameters in the RL algorithm on the control performance.



### 5.1. Vortex shedding suppression via reinforcement learning

After characterising the linear instability and flow sensitivity of the confined cylinder wake flows, we adopt the deep RL to control the vortex shedding in this flow using the two synthetic jets on the cylinder. In the literature, Rabault *et al.* (2019) has studied an RL-based control method to reduce the drag on the cylinder.

To begin, we would like to first obtain some preliminary results by defining the reward as

$$r_0 = - \sum_i^{n_{node}} ((u^i)^2 + (v^i)^2), \quad (5.1)$$

to have a peek of what the controlled wake may become. Here,  $r_0$  is the negative value of a plain sum of the kinetic energy in the computational domain and  $n_{node}$  is the number of grid nodes in DNS. Two points that deserve to be mentioned are that the kinetic energy will not decrease to zero as long as we feed the flow domain with an inflow and that the kinetic energy will decrease with the suppression of the vortex shedding. Thus, the RL algorithm will control the confined wake flow towards the lowest kinetic energy where the vortex shedding is weakest. All the control investigations from §§ 5.1–5.4 are based on 86 velocity probes covering the wavemaker region (see probe distribution in figure 19b). The reason for this selection will be explained in § 5.6.

Following Rabault *et al.* (2019), we use a random reset in the training process so that new episodes have a 20% possibility to start from the given uncontrolled initial condition and otherwise they start from the last state of the previous episode. As shown in figure 11(a), the averaged energy decreases with the training of the RL agent, especially after 100 episodes. The energy spikes of the learning curve result from the starting points of new episodes from the given initial condition. After 400 ~ 500 episodes, the RL training approximately converges. As shown in figure 11(b), we test the RL-based control from  $t = 16$  to  $t = 200$  (the flow starts from the initial condition at  $t = 0$ ). Although the testing control time is shorter than that in the fluidic pinball stabilisation (Maceda *et al.* 2021), the controlled flow shows a convergence to a periodic state in this case. The obtained control policy reduces the kinetic energy and the wake vortex shedding is suppressed (figure 12). At the beginning of the control (briefly after  $t = 16$ ), the energy first increases, and a slightly larger lift is observed consequently (see the inset in figure 11b). Then, the energy decreases exponentially and the oscillations of the lift coefficient are also reduced before reaching a steady oscillation. This result seems to be a common characteristics of the RL control applied to the wake flow in our case and will be seen later in our results. A larger actuation has also been observed in the RL control of Tang *et al.* (2020) whose aim is to reduce the drag on the cylinder. Eventually, the value of the lift coefficient converges to an oscillation with a small amplitude. Figure 12 shows the vorticity of the baseline and controlled flows. RL control increases the recirculation zone when it suppresses the oscillation. Nevertheless, slight vortex shedding downstream from the recirculation zone can still be observed, though the extent and amplitude are largely decreased.

### 5.2. SFD base flow

From figure 12, we see that when the confined wake flow is controlled, one salient feature is that the recirculation zone elongates. This is in fact a strong clue indicating that the controlled wake flow may converge to a flow that is close to the SFD base flow, because the recirculation zone in the latter case is relatively long (see figure 5).




Figure 11. RL-based vortex shedding suppression control using the plain sum of the kinetic energy ( $\beta = 0.25$  and  $Re = 150$ ). In panel (b), the grey shade means that no control is applied until  $t = 16$ . The reward function is (5.1). (a) Training process and (b) RL control performance.




Figure 12. Vorticity of the baseline confined cylinder flow with  $\beta = 0.25$  at  $Re = 150$  and the controlled flow with the RL agent trained by the total energy. RL control using the kinetic energy damps the oscillation by increasing the recirculation zone. Nevertheless, the vortex shedding is not fully suppressed. The reward function is (5.1).

To confirm this hypothesis, we use the fluctuation of the kinetic energy to monitor the strength of the vortex shedding, which is computed by

$$s_e = \sum_i^{n_{node}} ((u^i - u_0^i)^2 + (v^i - v_0^i)^2). \quad (5.2)$$

The shedding wake flow ( $\mathbf{u}$ ) can be decomposed as a sum of a reference part  $\mathbf{u}_0 = (u_0, v_0)$  and a shedding part  $\mathbf{u}_s$ , as in  $\mathbf{u} = \mathbf{u}_s + \mathbf{u}_0$ . To suppress the vortex shedding, we can use  $r = -s_e$  as the reward function in RL, which will minimise the fluctuation shedding energy. The idea is then to use the time-mean flow or the SFD base flow to evaluate the shedding energy in training RL control agents; that is, the mean flow  $\mathbf{u}_b$  (obtained by time-averaging the periodic vortex shedding) and the SFD base flow  $\bar{\mathbf{u}}$  (obtained by the SFD method) are used in (5.2) for  $\mathbf{u}_0$ . If the RL control with the SFD base flow being used in the reward can yield  $r \rightarrow 0$ , then we can understand that the controlled wake flow indeed converges to the SFD base flow. The time-mean flow is also tested for a comparison.

Figure 13 shows the RL control performances of the corresponding agents. Using the mean flow in the shedding energy evaluation cannot effectively suppress the vortex shedding. The result indicates that the time-averaged mean flow is modified to another status where the vortex shedding is slightly suppressed. Thus, the mean flow is not the flow to which the RL algorithm will lead the controlled wake flow. However, using the SFD base flow in the shedding energy evaluation (5.2) leads to a complete suppression. The shedding energy is reduced to approximately zero, which means that the controlled wake is almost the same as the SFD base flow. The fluctuations of the lift coefficient are also almost fully suppressed and no vortex shedding is observed in the vorticity contours. It can be concluded that vortex shedding suppression (by RL) in the confined cylinder

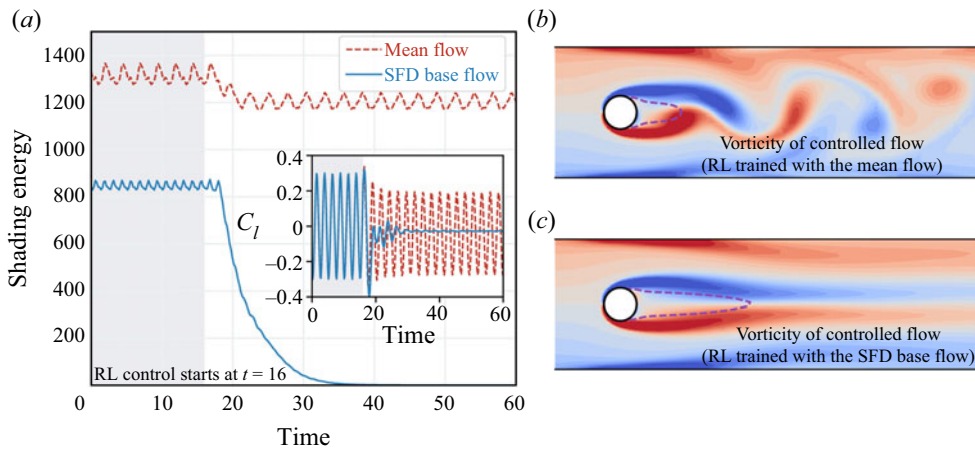


Figure 13. RL-based vortex shedding suppression of the confined cylinder wake ( $\beta = 0.25$  and  $Re = 150$ ) with RL agents trained using different reference flows (mean flow and SFD base flow). It can be seen that using the SFD base flow as the reference in the RL reward evaluation leads to an effective control with the vortex shedding fully suppressed. The reward function is (5.2).

wake flow is realised by modifying the wake flow to be a status which is very close to the SFD base flow. Interestingly, Flinois & Colonius (2015) found similar results in the adjoint-based optimal control of an unconfined cylinder wake that the controlled flow is the same as the SFD base flow. Because of its good performance, we will use the SFD base flow to evaluate the shedding energy for RL control in the following.

### 5.3. Control starting time

The above RL control starts at  $t = 16$ , which is arbitrarily chosen. We next test the robustness of the RL control policy in terms of its application at other phases in the vortex-shedding period. Figure 14 shows the control processes of the confined cylinder wake with the RL-based control starting from three different time stamps. The interval between adjacent time points is 0.8 unit time ( $t_2 - t_1 = t_3 - t_2 = 0.8$ ), which is approximately 1/3 period of the vortex shedding. Among the three cases explored here, a general trend is that the trained RL always uses large positive jet flow rates ahead of the next minimum point of the lift coefficient, when the upper vortex rolls up. Afterwards, relatively small jet rates are needed, and the shedding energy starts to decay. Figure 14 also shows the vorticity fields at the moments with largest jet flow rates, which are exactly the moment of the upper vortex rolling-up. The large blowing of the upper jet produces positive vorticity to offset the rolling-up vorticity, which can be considered to be destructive to the vortex development. This suggests that using large flow rates when the vortex rolls up is an important step in suppressing the cylinder wake fluctuations. The trained RL agent is adaptive to the shedding phase and captures the ‘right’ timing. Owing to the symmetry of the problem at hand, it can be understood that there is a similar RL control policy that uses significant blowing of the lower jet to suppress the rolling-up of the lower vortex.

### 5.4. Necessity of a persistent oscillating control

According to the linear stability analysis, the SFD base flow is unstable; the RL agent maintains this unstable status by actively modifying the jet flow rate in real time.

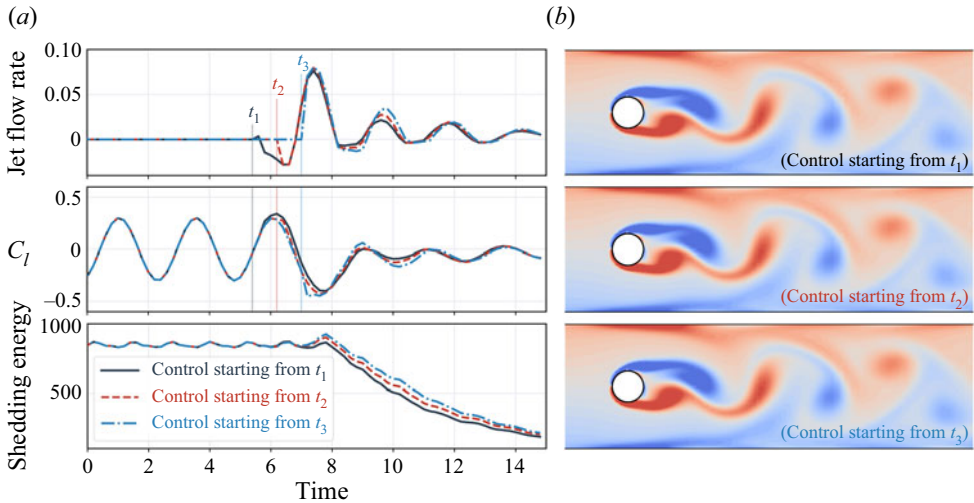


Figure 14. (a) Confined cylinder wakes ( $\beta = 0.25$  at  $Re = 150$ ) with the RL control starting at different time points ( $t_1$ ,  $t_2$  and  $t_3$ ). (b) Wake vorticities at  $t = 7.5$ . It can be seen that no matter when the control starts, RL always chooses to use large flow rates to offset the vorticity at the moment of the upper vortex rolling-up.

To probe how the stability property of the flow changes with time, we will use dynamic mode decomposition (DMD;  $N_{snapshot} = 20$  and  $\Delta t = 0.2$ ) to monitor the status of the flow stability/instability in the control process. In additional tests (not shown), we used more DMD snapshots and found that the results agreed well with those to be presented in the following. In figure 15, the control starts at  $t = 16$ , and in the beginning, again, relatively large jet flow rates are imposed by the RL agent. Then, as the lift coefficient is reduced, the flow rate of the synthetic jets decreases to a small value with a very small oscillating amplitude (see the inset in panel a). Even though the amplitude of the fluctuation is insignificant, we find that this fluctuation in the flow rate is vital to the successful continuous suppression of vortex shedding (we note in passing that the jet flow rate does not oscillate exactly around 0 because the overall effect of blowing/suction will introduce a small degree of asymmetry with respect to the channel centreline in the flow). We perform a numerical experiment in figure 15 to impose a sudden modification to the jet flow rate after  $t = 80$ . If nothing is changed in the RL control, when  $t > 80$ , the controlled wake flow is stable with the slightly fluctuated flow rates (the black lines). However, using a similar but constant jet flow rate ( $u = 0.00352$ , see the red dashed line in panel a) at  $t = 80$  soon triggers flow instability (panel c) and gradually leads to vortex shedding in the cylinder wake (panel b). Figure 16 shows the leading DMD modes of the two flows from  $t = 85$  to  $t = 100$ . With the active jet flow rate fluctuation, significant vorticity is observed around the jets in the leading mode of the RL controlled flow. However, the leading DMD mode remains almost unchanged in the (successful) controlled process from  $t = 85$  to  $t = 100$ . Yet if a constant jet flow rate is forced, the vorticity around the jets of the leading mode vanishes. At the same time, the leading mode becomes unstable and gradually evolves to a state with much stronger vorticities downstream the cylinder. Finally, the flow develops to a saturated periodic vortex shedding state, which has a similar  $C_l$  magnitude to the uncontrolled flow (see figure 15b). To sum up, in the vortex shedding suppression of cylinder wake, the RL-based policy tends to spend large energy at the start of the control in modifying the mean flow structure. Then, much less effort is required to




Figure 15. DMD analysis of the RL-based control process. The grey shade is the confined cylinder wake without control and the RL-based control starts at  $t = 16$ . We test a constant jet flow rate in the controlled cylinder wake from  $t = 80$ , which is shown as the red shade. The black solid lines show results of RL-based control and the red lines are results with the control switched to the constant flow rate.




Figure 16. Vorticities of the leading DMD modes in the controlled flows ( $\beta = 0.25$  at  $Re = 150$ ). Only part of the computational domain is shown for clarity. After the vortex shedding is damped, using a similar but constant jet flow rate soon triggers flow instability and gradually leads to vortex shedding in the cylinder wake.

maintain the stabilised flow. A similar tendency is found in the RL-based drag reduction control of a confined cylinder (Rabault *et al.* 2019). Nevertheless, a persistent active control using the RL agent, even though its oscillating amplitude is small, is required to suppress the vortex shedding in the cylinder wake. This is because the controlled flow is almost identical to the SFD base flow, which is unstable. Thus, when the persistent control is suppressed, the stabilised flow will become unstable and the vortex starts shedding again, as demonstrated in our numerical experiment above.

In real-world applications, the control system always faces uncertain noise. Tang *et al.* (2020) and Paris *et al.* (2021) have shown that the RL-based control can be robust to the Reynolds number variation and systematic noise. There is another uncertainty which is




Figure 17. DMD analysis of the RL-based control process with noisy forces in the inflow domain, which is imposed at  $t \in [60, 80]$ . We test two force magnitudes ( $A = 2.0$  and  $A = 10.0$ ), which are shown as the red and blue lines, respectively. The controlled wake flows are not developed to significant vortex shedding and become stabilised again after the noisy force is removed.

the incoming disturbances penetrating into the inlet domain. To study this influence, we place an external forcing term  $F(x, y)\gamma(t)$  following Hervé *et al.* (2012). Here,  $\gamma(t)$  is a random scaling factor of standard deviation  $\sigma_\gamma = 0.1$  and  $F(x, y)$  is the spatial structure with a Gaussian shape

$$F(x, y) = A \exp\left(\frac{-(x - x_0)^2}{2\sigma_F^2}\right) \exp\left(\frac{-(y - y_0)^2}{2\sigma_F^2}\right) (1, 1)^T, \quad (5.3)$$

where  $\sigma_F = 0.1$  and the spatial centre is  $x_0 = -1.5$  and  $y_0 = 0$ . The magnitude of the external force can be modified by changing  $A$ , and two values ( $A = 2.0$  and  $A = 10.0$ ) are studied. Figure 17 shows the RL-based control of the cylinder wake flow with noisy forces imposed at  $t \in [60, 80]$ . To stabilise the flow, jet flow rates controlled by the RL agent are increased slightly, and the increase magnitude is positively correlated with  $A$ . Although the noisy spatial forces trigger instabilities in the controlled wake flow, the wake does not develop significant vortex shedding. After the noisy force is removed ( $t > 80$ ), the wake gets fully stabilised again after 20 unit time. This numerical experiment shows that RL-based control is robust to spatial noise from the inflow and it is vital to sustain a persistent RL-based control.

### 5.5. Stability-enhanced reward

Next, we discuss how the reward in RL can be changed to incorporate some information on flow stability/instability to improve the control performance. Increasing the Reynolds

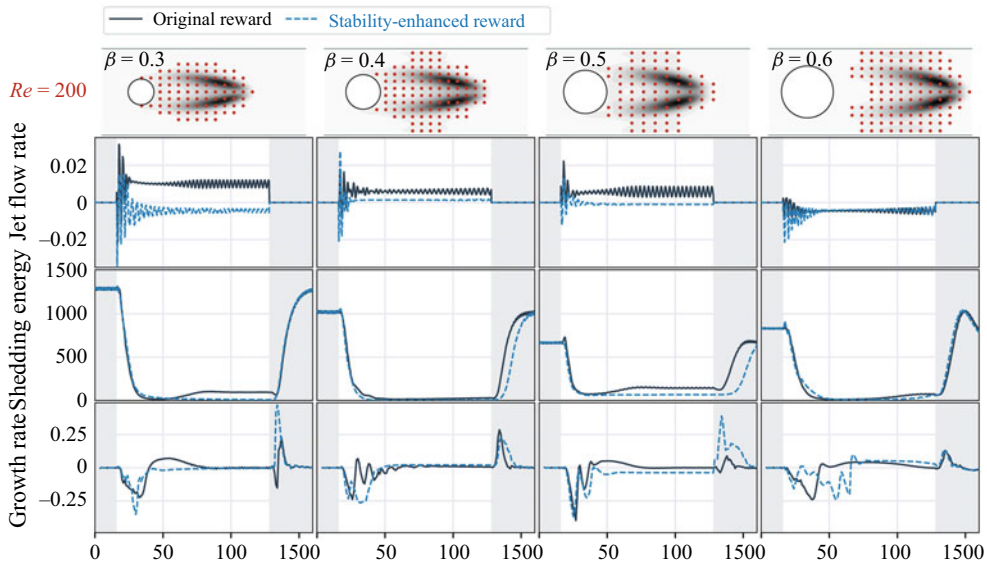


Figure 18. RL control performance of cylinder wakes with different  $\beta$  at  $Re = 200$ . The black region indicates the wavemaker region and the probes are placed to cover it (see the discussion in § 5.6). The stability-enhanced reward reduces unstable instants in the controlled wake flows. From  $t = 128$ , the RL-based control is shut down (shown as the red shade), and the controlled wake flows re-develop to significant vortex shedding, which shows the necessity of a persistent control.

number destabilises the confined wake flows, especially for cylinders with  $\beta \leq 0.5$ . In the RL-based drag reduction control of a confined cylinder wake with  $\beta \sim 0.25$ , Tang *et al.* (2020) found that small drag oscillations still existed in the controlled flow at  $Re = 200$  although four synthetic jets were used; and with a further increase of  $Re$ , although a significant reduction of the averaged drag was achieved, the RL agent could not find a fully stabilised control strategy to completely suppress oscillations in the drag coefficient. This means that the strong instability with the increase of  $Re$  in confined cylinder wakes brings significant challenges to RL-based control, and thus full suppression of the vortex shedding is difficult. We increase the Reynolds number to  $Re = 200$  (note the difference in the definitions of  $Re$  in our work ( $Re = U_{max}D/\nu$ ) and their works ( $Re = \bar{U}D/\nu$ );  $Re = 200$  in our work corresponds to  $Re = 400/3$  in their work) and investigate its influence on vortex shedding suppression with different  $\beta$ . As shown in figure 18, the increase of  $Re$  further pushes the wavemaker domains downstream of the cylinder. We place the probes for RL control to cover the wavemaker region (to be discussed in § 5.6) and train RL agents for vortex suppression at  $Re = 200$  using the original reward in (5.2). It can be seen that not all the controlled wake flows are fully stabilised owing to the increase of  $Re$ .

We used DMD ( $N_{snapshot} = 20$  and  $\Delta t = 0.2$ ) to analyse the controlled flows and found that the growth rate may flip to be positive even after the shedding energy is significantly reduced, which leads to flow instabilities and a subsequent increase of shedding energies. This unstable control is not preferable in practice. To obtain an effective control policy, we use a stability-enhanced reward function for RL, which embeds the information of flow stability/instability in the original reward:

$$\text{Reward} = -s_e e^g, \tag{5.4}$$

where  $s_e$  is the shedding energy defined in (5.2) and  $g$  is the largest growth rate of the flow, which is evaluated by DMD. For a stable flow, the growth rate is smaller than zero

and  $e^g < 1.0$ , which means a larger reward is given to motivate such control policies. For a neutral flow, the growth rate is zero, and (5.4) is equivalent to using the original reward defined by the shedding energy. For unstable flows, the growth rate is larger than zero. In this circumstance, the stability-enhanced term  $e^g > 1.0$ , and the reward decreases correspondingly. With such a reward function, the trained RL agent will come up with a control policy that compensates for the decrease of the reward arising from flow instability. This mitigates the adverse effect of flow instability in the control. We train RL agents using the stability-enhanced reward for vortex suppression with different  $\beta$  at  $Re = 200$ . As shown in figure 18, using the stability-enhanced reward, the instants when the growth rate is larger than zero become fewer in the controlled wake flows. Thus, instabilities in the controlled flows are significantly damped. Such an instability-abated effect is of great value in practical applications such as prevention of aerodynamic buffeting (Gao *et al.* 2017) and aeroelastic flutter (Jonsson *et al.* 2019). In most cases (figure 18), the shedding energy with the stability-enhanced reward is reduced to a lower value that can be maintained for a long run, and, in these circumstances, the amplitude of the control actions is decreased as well. This implies that it is favourable to maintain the controlled cylinder wake flow close to the base flow because it costs lower control energy, which is a beneficial side effect of the proposed stability-enhanced reward. Nevertheless, there is still no guarantee to get a fully stabilised wake flow by using the proposed stability-enhanced reward (see  $t \in [100, 128]$  in the case of  $\beta = 0.6$ ). We also investigate the destabilising process of the controlled flow with the shut-down of control actions (shown as the red shades in figure 18). The results show that a continuous control is necessary; otherwise, the flow will become unstable.

We provide some discussions on these results in terms of flow physics. In the literature of 2-D cylindrical wake flow, the shift mode has been found to be indispensable in a successful reduced-order model (ROM) of the transient dynamics in this flow (Noack *et al.* 2003). Inclusion of this mode in a ROM realises a mean field correction, pointing from the unstable steady flow to the time-averaged mean flow. Our results on the RL of the confined wake flow can be understood in a similar manner. The controlled flow is the SFD base flow, which is unstable and tends to evolve to the periodic flow once the control is turned off (see figure 15). Thus, when we improve the RL agent by telling it how to abate the instability, the control can be more effective and efficient, as shown in our tests (see figure 18). Admittedly, our stability-enhanced reward is still crude as we simply incorporate the information of the first DMD mode in the new reward. In answering a comment by one of the reviewers, we realise that the analysis and RL control can be conducted more delicately by abating the flow instability in the direction of the unstable base flow drifting to the time-mean flow (currently we do not know how much the abatement effect of using the first DMD mode can contribute in this direction). This may lead to a more efficient RL control. More research is required to explore the new design of reward function in RL in the future.

### 5.6. A heuristic strategy of probe placement

The reinforcement learning agent observes environment changes via the probes. In preparing this work, we found that Paris *et al.* (2021) studied the probe placement by modifying the architecture of RL-based control by adding a stochastic gated input layer between the states and agent. The modified RL can select an optimal subset from the initial probe sensors. However, a good control performance in this case may be largely dependent on the initial placement of the probes. It may be the case that the initial placement of the probes is not able to encompass the global optimal solution of the control. To achieve an effective RL-based control, probes should be placed in domains that can convey all the



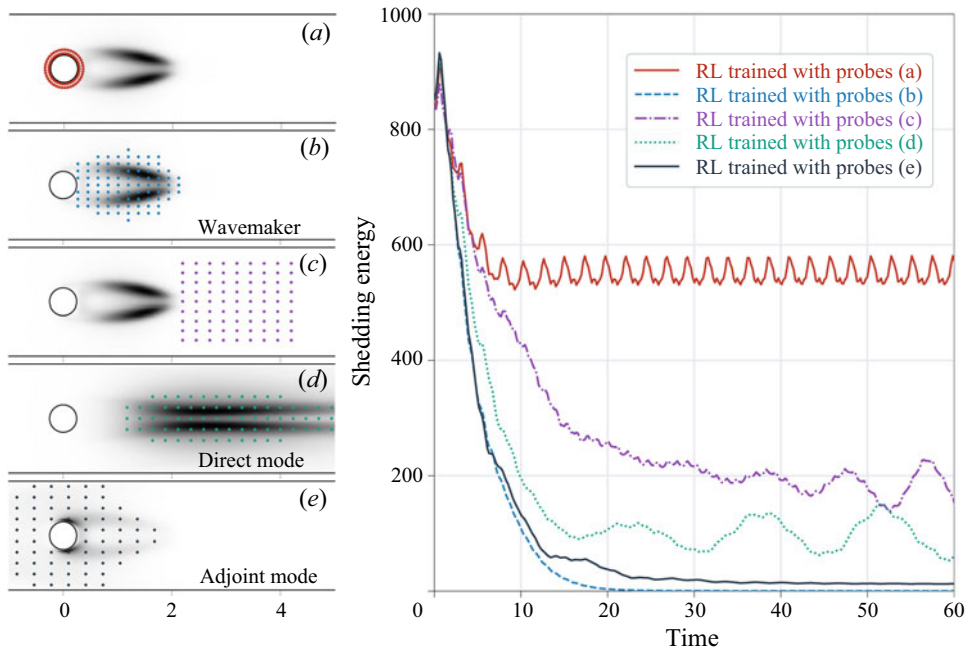


Figure 19. Probes covering the wavemaker zone lead to the best control of the confined cylinder wake ( $\beta = 0.25$  at  $Re = 150$ ).

essential changes in the flow. In the following, we take a heuristic approach to evaluate the placement strategy of the probes by harnessing the results of the structure sensitivity.

As the most sensitive region to flow perturbations, the wavemaker region may be the ideal monitoring position for the probe placement. Strykowski & Sreenivasan (1990) showed that the wavemaker is similar to the region where one can place a small control cylinder to suppress the vortex shedding. We hypothesise that the most effective probes should be placed to cover (at least partially) the wavemaker domain in RL-based control of vortex suppression in confined cylinder wakes. We place the probes according to the wavemaker region calculated using the SFD base flow. In the following, three sets of tests have been conducted as shown in figures 19–21, which are explained below.

First, in figure 19, we would like to demonstrate that the probes should be placed to cover the wavemaker region. To verify this, we consider three kinds of probes: (1) probes upstream of the wavemaker; (2) probes inside the wavemaker; (3) probes downstream of the wavemaker. As shown in figure 19, placing the probes upstream of the wavemaker significantly reduces control effectiveness (even though they are very close to the cylinder); while placing the probes downstream of the wavemaker cannot train a stable control policy either, even though the RL performance is better than the previous case. Probes placed inside the wavemaker have an effective and stable RL control. Because the wavemaker is obtained by overlapping the direct mode and adjoint mode, we further investigate the control effectiveness with probes being placed to cover the direct mode and adjoint mode, respectively. As shown in figure 19, the RL-based control with probes covering the direct mode is ineffective, while the control with probes covering the adjoint mode is effective and the performance is almost close to those covering the wavemaker. This may arise from the fact that the adjoint mode covers most of the wavemaker

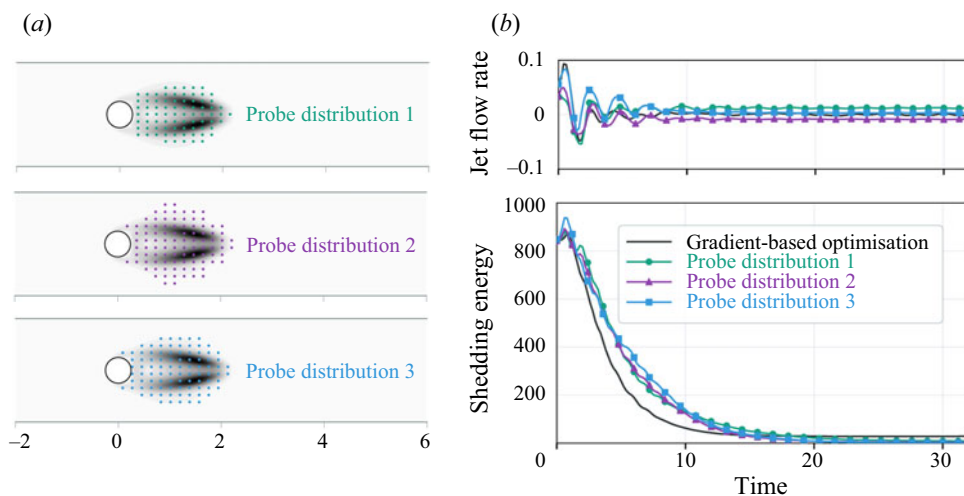


Figure 20. Even with slight differences in the placement, probes covering the wavemaker zone can always lead to effective and stable RL control of the confined cylinder wake ( $\beta = 0.25$  at  $Re = 150$ ). The optimal control policy solved by the gradient-based optimisation achieves better short-term performance and RL agents lead to a better performance in the long term.

domain while the direct mode is located further downstream. Although the direct mode describes the spatial structure of vortex shedding and most residual non-stationarities in the controlled wake flows are located in the region of the direct mode, it is more effective for RL-based control to monitor flow changes in the core domain of instability.

Next, in [figure 20](#), we would like to demonstrate that a small difference in the probe distribution will not affect the general good function of RL as long as the main part of the wavemaker is covered, which will confirm the robustness of its performance. All the control policies also display the similar features found in the last section; that is, large flow rates of the synthetic jets are used in the beginning to modify the mean flow structure and much smaller time-varying jet flow rates are required to maintain the status. In the same figure, we have also added a result of a gradient-based optimisation method in reducing the shedding energy. This consideration is to understand how the RL-based policy compares with other control methods. In the gradient-based optimisation method, the flow rates at different time steps are discretised as independent design variables with a control horizon of 32 time units. The objective function of the control design optimisation is the total shedding energy in the control time-horizon. The gradient of the objective function with respect to the flow rates is solved by using the finite difference method. More details are provided in [Appendix C](#). It can be seen in [figure 20](#) that placing the probes in the wavemaker region achieves a control performance that is close to that solved by the gradient-based optimisation. The control policy found by gradient-based optimisation leads to more reductions of the shedding energy in the short-term ( $t \in [5, 15]$ ), because this strategy can achieve the smallest total shedding energy in the chosen control horizon (meaning that the controlled solution based on the gradient-based optimisation method is optimal in the sense that the total energy with  $t \in [0, 32]$  is minimum). Nevertheless, it cannot be guaranteed that other solutions may exist that have locally smaller shedding energies than the policy found by the gradient-based optimisation. For example, our RL-based controls have slightly better long-term performances (for instance, look at the results at  $t = 30$  in the figure), which can be more preferable in practice.

The wavemaker domain of the confined cylinder wake changes with the blockage ratio or the Reynolds number. We investigate the influence of different blockage ratios on RL-based control performance at  $Re = 150$ . As shown in figure 22 of Appendix D, increasing the blockage ratio pushes the wavemaker to the downstream region, as we have discussed above. We place probes for RL control of confined wakes with different blockage ratios at the corresponding wavemaker regions. With the RL-based control, the shedding energies are all reduced to a small value. This means that regardless of the blockage ratio, as long as the probes are placed to cover the wavemaker region, the RL-control policy is able to effectively control the confined flow to damp the vortex shedding (at least for the Reynolds number investigated,  $Re = 150$ ). In practice, if the probes are fixed in a series of control tests with different parameters (such as  $\beta$  or  $Re$ ), the optimal placement of the probes should be determined to be suitable for all the parameters (the trend of how the flow properties change with these flow parameters has been analysed in the previous sections).

Finally, we consider the cases where the number of probes is insufficient. We choose to train the RL agent using ten probes. The probe distributions and control performances are shown in figure 21. Distributions 1 ~ 3 are evenly placed in the wavemaker region, and Distributions 4 ~ 5 are placed in the most sensitive part (black). We can see that even though all the five placements successfully reduce the shedding energy, difference exists in their performance. The best performance belongs to Distribution 1, where the shedding energy has been abated to almost zero when  $t > 20$ . This result indicates that Distribution 1 should be close to the global optimal solution in the RL (in which the shedding energy is zero). However, Distributions 2,3, even though they are also trying to cover the wavemaker region as fully as possible, somehow perform poorly owing to the insufficient number of probes. Note that, compared with figure 20, the results of Distributions 2,3 should be considered to be worse. Finally, when clustering all the available probes in the most sensitive region (black colour) as in Distributions 4,5, the RL control does not necessarily yield good performance. All these results indicate that when insufficient probes are used, the performance of the RL control is scattered depending on the placement strategy of the probes. Our heuristic approach indicate that Distribution 1 is the best choice; however, this fortuitous result may not be carried over to other situations. Based on this result, future works can consider coupling the proposed criterion of probe placement (to cover the wavemaker region) with an optimisation method with the former providing a general good choice of the initial placement of the probes and the latter fine-tuning the selection of the probes by the optimisation of the probes, exemplifying the combination of the *a priori* knowledge on flow physics and the power of RL algorithm. For example, the modified RL developed by Paris *et al.* (2021) can be used to select the optimal subset from the probes which are initially covering or not covering the wavemaker region (and we hypothesise that the probes covering the wavemaker region will be selected when vortex shedding is to be abated).

## 6. Conclusions

In this work, we have studied the linear stability and flow sensitivity of confined wake flows in a range of  $\beta$  and  $Re(\leq 200)$ . The main objective of the work is to understand further the reinforcement learning (RL) algorithm in controlling a complex flow system and showcase how some information of the flow physics can help us in designing and applying the RL in the flow control (more specifically, to suppress the vortex shedding by actively changing the flow rates of two synthetic jets on the cylinder).

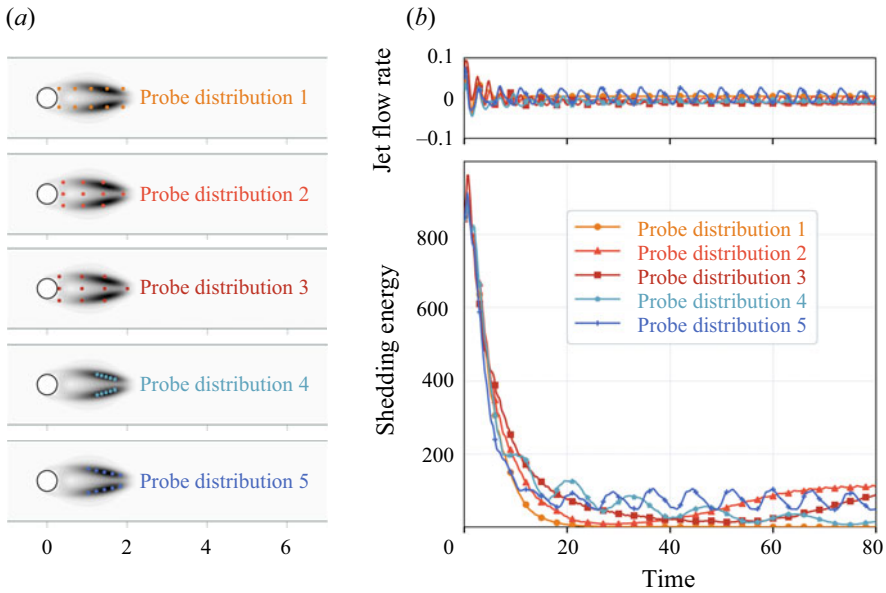


Figure 21. RL control of the confined cylinder wake ( $\beta = 0.25$  at  $Re = 150$ ) using ten probes. Different distributions of probes lead to a significant divergence in the control performance.




Figure 22. RL control performance of cylinder wakes with different  $\beta$  at  $Re = 150$ . The probes are placed based on the wavemaker-based criterion.

### 6.1. Flow stability and sensitivity of confined wake flows

In the stability study of the confined cylinder wakes, the neutral curve in the  $Re$ - $\beta$  plane is determined. In a range of  $Re \in [0, 200]$ , the confined cylinder wake with  $\beta < 0.7$  has one critical  $Re$ , where the wake becomes unstable beyond this point and the vortex shedding starts. For confined cylinder wakes with  $\beta > 0.7$ , other than the main recirculation bubble downstream the cylinder, two additional recirculation bubbles may develop close to the confinement walls further downstream of the cylinder with the increasing effect of confinement. This is similar to the observations by Sahin & Owens (2004) and the recirculation bubbles on the confinement walls become larger with the increase of  $Re$ . This is associated with a second critical point, after which the vortex shedding is suppressed for confined cylinder wakes with  $\beta > 0.7$ . Vortex shedding of confined cylinder wakes

with different  $\beta$  and  $Re$  is studied using the global linear stability analyses. Similar to the findings by Maurel *et al.* (1995), Noack *et al.* (2003) in unconfined cylinder wakes, we find that the oscillating wakes of confined cylinders saturates when the time-averaged flow is marginally stable. The frequencies solved by nonlinear DNS agree well with the results of the global linear analysis based on the time-mean flow, rather than the SFD base flow. Thus the mean flow can provide a good profile to predict the shedding frequency of confined cylinder wakes. The relationship between the SFD base flow and the saturated mean flow is studied by using global linear stability and DMD. Similar to the unconfined flow past a cylinder (Barkley 2006), the evolution of the unstable SFD base flow in the confined cylinder wake is a nonlinear saturation of oscillations, and the selection of vortex shedding amplitude and frequency is based on the marginal stability of the mean flow.

The wavemaker regions of confined cylinder wake flows with various  $\beta$  and  $Re$  are obtained by performing structural sensitivity analyses. Similar to unconfined cylinder wakes (Giannetti & Luchini 2007; Marquet *et al.* 2008), the wavemaker region of the confined wake flow is located downstream of the cylinder, which means the downstream domain is more important in terms of flow sensitivity and should be monitored in the suppression of wake vortex shedding. We found that with the increase of  $\beta$ , a longer and wider recirculation zone is developed and the wavemaker region expands correspondingly and moves downstream. Increasing  $Re$  also leads to longer recirculation zones and the wavemaker region is pushed downstream. The results imply that to efficiently suppress the vortex shedding in confined cylinder wakes, perturbations further downstream from the cylinder should be monitored when  $\beta$  or  $Re$  increases because the most sensitive region (wavemaker) is located further downstream.

## 6.2. *RL-based control of confined wake flows*

In the second part of this work, we used RL-based control to suppress the vortex shedding in the cylinder wake. It is found that using the sum of kinetic energy as the reward function can suppress the oscillation to some extent but does not fully damp the shedding. We define a reward function based on the shedding energy relative to the SFD base flow and this reward gives rise to a more effective RL control. With the shedding energy being reduced to approximately zero, it can be concluded that vortex shedding suppression of the confined cylinder wake flow is realised by modifying the confined wake flow to a status similar to the SFD base flow. Similar results have been observed in the adjoint-based optimal control of an unconfined cylinder wake by Flinois & Colonius (2015). We found that the RL-based control tends to spend large energy in the beginning to fast modify the mean flow structure. No matter when the control starts, the trained RL agent is adaptive to the shedding phase and can capture the ‘right’ timing, i.e. at the moment of the vortex rolling-up, to excite large flow rates. Afterwards, much less effort with significantly small jet flow rates is required to maintain the stabilised flow. Nevertheless, consistent active control using the RL agent is still necessary; otherwise, the stabilised flow will become unstable and the vortex starts shedding again. This is because the controlled flow is the unstable SFD base flow.

For confined cylinder wakes with  $\beta \leq 0.5$ , increasing  $Re$  brings more challenges to the vortex shedding suppression owing to the rise of flow instability (from the modal stability analysis, we know that for smaller  $\beta \leq 0.5$ , the flow becomes more unstable when  $Re$  increases). In this circumstance, we used a stability-enhanced reward function to embed the flow instability (evaluated by an instability penalty) into the RL reward. With the stability-enhanced reward and the wavemaker-based probe placement, vortex shedding

of confined cylinder wakes with different  $\beta$  at higher  $Re$  can be further suppressed with different degrees of success.

In the end, we find that placing the velocity probes covering the wavemaker is preferable in RL-based control of confined cylinder wakes to suppress the vortex shedding. As mentioned earlier, the wavemaker region is a region where the flow is most sensitive to variation. Our results show that placing the probes to cover the wavemaker region yields a better performance of the RL than placing them elsewhere. This can be interpreted as the flow information (as the environment and state components in the RL framework) is more accurately detected by the probes if they are placed in this manner, so that the action can be more efficient in controlling the flow. Furthermore, the robustness has also been confirmed: as long as the main part of the wavemaker is covered, small differences in the placement of probes do not affect the general good performance of RL. When the probes are properly placed, we also find that the RL-control policy can outperform a gradient-based optimisation method (optimised in a certain time-horizon) in the long run. When insufficient probes are used, the performance of the probe distributions considered in this work is scattered. Our heuristic approach is able to identify a good distribution as an initial strategy for the placement of the probes, but more systematic approaches should also be adopted, such as that in Paris *et al.* (2021) to obtain more desirable results. Combining this heuristic result with the optimisation method thus may be promising.

In this study, the best policy trained in a fixed episode number (500) is used, which has shown a good convergence because there is only one independent actuator. For policy-gradient RL methods such as PPO, the policy network integrates over both state and action spaces, and increasing the number of actuators/actions may require much more training episodes. Thus, the exploitation performance may be significantly decreased within a given learning budget, especially for cases using computationally expensive three-dimensional simulations. In this circumstance, the deterministic policy gradient method (Silver *et al.* 2014) might be more efficient because the policy merely integrates over the state space. Furthermore, proper use of the knowledge of flow physics in a control agent can simplify its structure without losing performance (see the phase control in the drag reduction of a bluff body in Pastoor *et al.* 2008, which is mostly based on the understanding of decoupling shear layer development and wake processes). More research on the coupling of fluid mechanics and RL-based control is needed to improve efficiency, effectiveness and robustness in real-world complex applications.

**Acknowledgements.** We acknowledge the computational resources provided by the National Supercomputing Centre of Singapore.

**Funding.** The work is supported by a Tier 2 grant from the Ministry of Education, Singapore (R-265-000-661-112).

**Declaration of interest.** The authors report no conflict of interest.

**Author ORCIDs.**

① Jichao Li <https://orcid.org/0000-0003-2541-2557>;

② Mengqi Zhang <https://orcid.org/0000-0002-8354-7129>.

## Appendix A. Open-source code

The code for the RL-based flow control in this study is open-source as a GitHub repository: [https://github.com/npuljc/RL\\_control\\_Nek5000](https://github.com/npuljc/RL_control_Nek5000). It is developed by referring to the open-source RL-based control code of Rabault *et al.* (2019), so they have the same structure and both are based on the PPO agent implemented in Tensorforce.

The main difference is the simulation environment, which is based on a SEM solver Nek5000 (Version 19.0, Fischer *et al.* 2017) in our repository. Another difference is that the reward function in our code is defined by the shedding energy (5.2) and there is an option to use the stability-enhanced reward (5.4). More details can be found in the Readme file of the repository.

### Appendix B. Dynamic mode decomposition

In our analysis, we have also used the dynamic mode decomposition (DMD) to further probe the flow instability when the conventional linear stability analysis is hard to apply. When the flow rate of the synthetic jet is zero (i.e. the control is off), the conventional linear stability analysis can be applied properly. However, when the control is on and the synthetic jet flow rates are varying, it turns out to be difficult to apply the conventional linear stability analysis and in this case we use DMD.

As first proposed by Schmid (2010), Rowley *et al.* (2009), DMD is a data-based method in the analysis of the time evolution of fluid flows. We briefly summarise the algorithm and refer the reader to Schmid (2010) for the complete theory. For a series of flow snapshots  $V_1^{N_{\text{snapshot}}} = \mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{N_{\text{snapshot}}}$  that is generated with a time interval of  $\Delta t$ , one assumes that a linear mapping  $\mathcal{A}$  connects the flow field  $\mathbf{v}_i$  to the subsequent flow field  $\mathbf{v}_{i+1}$ , that is,  $\mathbf{v}_{i+1} = \mathcal{A}\mathbf{v}_i$ . Then, we have  $\mathcal{A}V_1^{N_{\text{snapshot}}-1} = V_1^{N_{\text{snapshot}}-1}\mathcal{S} + \mathbf{r}e_{N_{\text{snapshot}}-1}^T$ , where  $\mathbf{r}$  is the residual vector and  $e_{N_{\text{snapshot}}-1} \in \mathbb{R}^{N_{\text{snapshot}}-1}$  is a  $(N_{\text{snapshot}} - 1)$ th unit vector. The eigenvalues of  $\mathcal{S}$  then approximate some of the eigenvalues of  $\mathcal{A}$ . To improve the robustness, we use an implementation based on eigenvalue decomposition of a ‘full’ matrix  $\hat{\mathcal{S}}$ , which is related to  $\mathcal{S}$  via a similarity transformation:

$$\hat{\mathcal{S}} = UV_2^{N_{\text{snapshot}}}W\Sigma^{-1}, \tag{B1}$$

where  $U$ ,  $W$  and  $\Sigma$  are obtained by performing a singular value decomposition of  $V_1^{N_{\text{snapshot}}-1} = U\Sigma W^H$ . The dynamic modes are  $\Phi_i = Uy_i$ , where  $y_i$  is the  $i$ th eigenvector of  $\hat{\mathcal{S}}$ , i.e.

$$\hat{\mathcal{S}}y_i = \mu_i y_i. \tag{B2}$$

The growth rates and frequencies of the modes can be obtained by the logarithmic mapping of corresponding DMD eigenvalues. For the  $i$ th mode with eigenvalue  $\mu_i$ , the the frequency  $f_i = \text{Im}(\ln \mu_i)/2\pi\Delta t$  and growth rate  $g_i = \text{Re}(\ln \mu_i)/\Delta t$ .

### Appendix C. Gradient-based optimisation of the control policy for synthetic jet flow rates

For reference, gradient-based optimisation is used to solve the optimal control policy in a given control horizon of 32 time units, which equals to two episode in RL-based control. The sequential least-squares programming (SLSQP) algorithm implemented in pyOptSparse (Perez, Jansen & Martins 2012; Wu *et al.* 2020) (<https://github.com/mdolab/pyoptsparse>) is used to minimise the total shedding energy in the control horizon. The shedding energy is evaluated using (5.2) with a reference of the SFD base flow. The finite difference method implemented in pyOptSparse is used to solve the gradient of the shedding energy with respect to each control variable. As mentioned above, the synthetic jet flow rate can be adjusted every 0.2 time unit and this means that we have 160 independent control variables in the optimisation. To reduce the computational cost

	Functions	Quantity	Description
minimise	$\sum_{t=0}^{t=32} e_s^{(t)}$	1	Total shedding energy in a control horizon of 32 time units
with respect to	$k_i$	80	Cubic spline knots for jet flow rates

Table 2. Optimisation problem statement of the synthetic jet flow rate control policy.

in gradient evaluations, the 160 jet flow rates are parametrised by a cubic spline with 80 evenly distributed knots. Referring to the RL-based control policies, a cubic spline with 80 knots can provide enough degrees of freedom for the control of synthetic jet flow rates in this problem. Thus, the gradient-based optimisation is actually subject to 80 design variables. We summarise this optimisation problem in [table 2](#).

#### Appendix D. RL control of cylinder wakes with different $\beta$

To see if the wavemaker-based probe placement criterion works with the change of the blockage ratio, we investigate the influence of different blockage ratios on RL-based control performance at  $Re = 150$ , and probes for RL control are adjusted based on the corresponding wavemaker region of the confined wake. The results are shown in [figure 22](#).

#### REFERENCES

- AKERVIK, E., BRANDT, L., HENNINGSON, D.S., HOEPFFNER, J., MARXEN, O. & SCHLATTER, P. 2006 Steady solutions of the Navier–Stokes equations by selective frequency damping. *Phys. Fluids* **18** (6), 068102.
- ANAGNOSTOPOULOS, P., ILIADIS, G. & RICHARDSON, S. 1996 Numerical study of the blockage effects on viscous flow past a circular cylinder. *Intl J. Numer. Meth. Fluids* **22** (11), 1061–1074.
- ARNOLDI, W.E. 1951 The principle of minimized iterations in the solution of the matrix eigenvalue problem. *Q. Appl. Maths* **9** (1), 17–29.
- BARKLEY, D. 2006 Linear analysis of the cylinder wake mean flow. *Europhys. Lett. (EPL)* **75** (5), 750–756.
- BARKLEY, D., BLACKBURN, H.M. & SHERWIN, S.J. 2008 Direct optimal growth analysis for timesteppers. *Intl J. Numer. Meth. Fluids* **57** (9), 1435–1458.
- BEINTEMA, G., CORBETTA, A., BIFERALE, L. & TOSCHI, F. 2020 Controlling Rayleigh–Bénard convection via reinforcement learning. *J. Turbul.* **21** (9–10), 585–605.
- BELUS, V., RABAULT, J., VIQUERAT, J., CHE, Z., HACHEM, E. & REGLADE, U. 2019 Exploiting locality and translational invariance to design effective deep reinforcement learning control of the 1-dimensional unstable falling liquid film. *AIP Adv.* **9** (12), 125014.
- BOUJO, E. & GALLAIRE, F. 2014 Controlled reattachment in separated flows: a variational approach to recirculation length reduction. *J. Fluid Mech.* **742**, 618–635.
- BRANDT, L., SIPP, D., PRALITS, J.O. & MARQUET, O. 2011 Effect of base-flow variation in noise amplifiers: the flat-plate boundary layer. *J. Fluid Mech.* **687**, 503–528.
- BRUNTON, S.L., NOACK, B.R. & KOUMOUTSAKOS, P. 2020 Machine learning for fluid mechanics. *Annu. Rev. Fluid Mech.* **52** (1), 477–508.
- BRUNTON, S.L., PROCTOR, J.L. & KUTZ, J.N. 2016 Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proc. Natl Acad. Sci.* **113** (15), 3932–3937.
- BUDANUR, N.B. & CVITANOVIĆ, P. 2017 Unstable manifolds of relative periodic orbits in the symmetry-reduced state space of the Kuramoto–Sivashinsky system. *J. Stat. Phys.* **167** (3), 636–655.
- CHEN, J.-H., PRITCHARD, W.G. & TAVENER, S.J. 1995 Bifurcation for flow past a cylinder between parallel planes. *J. Fluid Mech.* **284**, 23–41.
- CORBETT, P. & BOTTARO, A. 2001 Optimal control of nonmodal disturbances in boundary layers. *Theor. Comput. Fluid Dyn.* **15** (2), 65–81.
- COUTANCEAU, M. & BOUARD, R. 1977 Experimental determination of the main features of the viscous flow in the wake of a circular cylinder in uniform translation. Part 1. Steady flow. *J. Fluid Mech.* **79** (2), 231–256.



- DELAUNAY, Y. & KAIKTSIS, L. 1999 Active control of cylinder wakes: use of base suction and blowing. *ESAIM: Proc.* **7**, 104–119.
- DURAISAMY, K., IACCARINO, G. & XIAO, H. 2019 Turbulence modeling in the age of data. *Annu. Rev. Fluid Mech.* **51** (1), 357–377.
- ERIKSSON, L.E. & RIZZI, A. 1985 Computer-aided analysis of the convergence to steady state of discrete approximations to the euler equations. *J. Comput. Phys.* **57** (1), 90–128.
- FAN, D., YANG, L., WANG, Z., TRIANTAFYLLOU, M.S. & KARNIADAKIS, G.E. 2020 Reinforcement learning for bluff body active flow control in experiments and simulations. *Proc. Natl Acad. Sci.* **117** (42), 26091–26098.
- FISCHER, P.F., LOTTES, J.W. & KERKEMEIER, S.G. 2017 Nek5000 Version 17.0. Argonne National Laboratory, Illinois. Available at: <https://nek5000.mcs.anl.gov>.
- FLINOIS, T.L.B. & COLONIUS, T. 2015 Optimal control of circular cylinder wakes using long control horizons. *Phys. Fluids* **27** (8), 087105.
- GAO, C., ZHANG, W., KOU, J., LIU, Y. & YE, Z. 2017 Active control of transonic buffet flow. *J. Fluid Mech.* **824**, 312–351.
- GARNIER, P., VIQUERAT, J., RABAULT, J., LARCHER, A., KUHNLE, A. & HACHEM, E. 2021 A review on deep reinforcement learning for fluid mechanics. *Comput. Fluids* **225**, 104973.
- GIANNETTI, F. & LUCHINI, P. 2007 Structural sensitivity of the first instability of the cylinder wake. *J. Fluid Mech.* **581**, 167–197.
- GLEZER, A. 2011 Some aspects of aerodynamic flow control using synthetic-jet actuation. *Phil. Trans. R. Soc. A: Math. Phys. Engng Sci.* **369** (1940), 1476–1494.
- HERVÉ, A., SIPP, D., SCHMID, P.J. & SAMUELIDES, M. 2012 A physics-based approach to flow control using system identification. *J. Fluid Mech.* **702**, 26–58.
- HILL, D. 1992 A theoretical approach for analyzing the restabilization of wakes. In *30th Aerospace Sciences Meeting and Exhibit. AIAA Paper* 1992-67.
- HUERRE, P. & MONKEWITZ, P.A. 1990 Local and global instabilities in spatially developing flows. *Annu. Rev. Fluid Mech.* **22**, 473–537.
- JACKSON, C.P. 1987 A finite-element study of the onset of vortex shedding in flow past variously shaped bodies. *J. Fluid Mech.* **182** (-1), 23.
- JAMESON, A. 1988 Aerodynamic design via control theory. *J. Sci. Comput.* **3** (3), 233–260.
- JONSSON, E., RISO, C., LUPP, C.A., CESNIK, C.E., MARTINS, J.R. & EPUREANU, B.I. 2019 Flutter and post-flutter constraints in aircraft design optimization. *Prog. Aerosp. Sci.* **109**, 100537.
- KANARIS, N., GRIGORIADIS, D. & KASSINOS, S. 2011 Three dimensional flow around a circular cylinder confined in a plane channel. *Phys. Fluids* **23** (6), 064106.
- KIM, J. & BEWLEY, T.R. 2007 A linear systems approach to flow control. *Annu. Rev. Fluid Mech.* **39** (1), 383–417.
- KUHNLE, A., SCHAARSCHMIDT, M. & FRICKE, K. 2017 Tensorforce: a tensorflow library for applied reinforcement learning. Web page.
- LEONTINI, J.S., THOMPSON, M.C. & HOURIGAN, K. 2010 A numerical study of global frequency selection in the time-mean wake of a circular cylinder. *J. Fluid Mech.* **645**, 435–446.
- LING, J., KURZAWSKI, A. & TEMPLETON, J. 2016 Reynolds averaged turbulence modelling using deep neural networks with embedded invariance. *J. Fluid Mech.* **807**, 155–166.
- LOISEAU, J.-C., NOACK, B.R. & BRUNTON, S.L. 2018 Sparse reduced-order modelling: sensor-based dynamics to full-state estimation. *J. Fluid Mech.* **844**, 459–490.
- LUCHINI, P. & BOTTARO, A. 2014 Adjoint equations in stability analysis. *Annu. Rev. Fluid Mech.* **46** (1), 493–517.
- MACEDA, G.Y.C., LI, Y., LUSSEYRAN, F., MORZYŃSKI, M. & NOACK, B.R. 2021 Stabilization of the fluidic pinball with gradient-enriched machine learning control. *J. Fluid Mech.* **917**, A42.
- MARQUET, O., SIPP, D. & JACQUIN, L. 2008 Sensitivity analysis and passive control of cylinder flow. *J. Fluid Mech.* **615**, 221–252.
- MAUREL, A., PAGNEUX, V. & WESFREID, J.E. 1995 Mean-flow correction as non-linear saturation mechanism. *Europhys. Lett. (EPL)* **32** (3), 217–222.
- NOACK, B.R., AFANASIEV, K., MORZYŃSKI, M., TADMOR, G. & THIELE, F. 2003 A hierarchy of low-dimensional models for the transient and post-transient cylinder wake. *J. Fluid Mech.* **497**, 335–363.
- NOACK, B.R. & ECKELMANN, H. 1994 A global stability analysis of the steady and periodic cylinder wake. *J. Fluid Mech.* **270**, 297–330.
- PARIS, R., BENEDDINE, S. & DANDOIS, J. 2021 Robust flow control and optimal sensor placement using deep reinforcement learning. *J. Fluid Mech.* **913**, A25.
- PASTOOR, M., HENNING, L., NOACK, B.R., KING, R. & TADMOR, G. 2008 Feedback shear layer control for bluff body drag reduction. *J. Fluid Mech.* **608**, 161–196.

- PEPLINSKI, A., SCHLATTER, P., FISCHER, P.F. & HENNINGSON, D.S. 2014 Stability tools for the spectral-element code NEK5000: application to jet-in-crossflow. In *Spectral and High Order Methods for Partial Differential Equations - ICOSAHOM 2012* (ed. M. Azañez, H. El Fekih & J.S. Hesthaven), pp. 349–359. Springer International Publishing.
- PEREZ, R.E., JANSEN, P.W. & MARTINS, J.R.R.A. 2012 pyOpt: a Python-based object-oriented framework for nonlinear constrained optimization. *Struct. Multidiscipl. Optim.* **45** (1), 101–118.
- PIER, B. 2002 On the frequency selection of finite-amplitude vortex shedding in the cylinder wake. *J. Fluid Mech.* **458**, 407–417.
- PIER, B. & HUERRE, P. 2001 Nonlinear self-sustained structures and fronts in spatially developing wake flows. *J. Fluid Mech.* **435**, 145–174.
- PROVANSAL, M., MATHIS, C. & BOYER, L. 1987 Bénard-von Kármán instability: transient and forced regimes. *J. Fluid Mech.* **182**, 1–22.
- RABAUULT, J., KUCHTA, M., JENSEN, A., RÉGLADE, U. & CERARDI, N. 2019 Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control. *J. Fluid Mech.* **865**, 281–302.
- RABAUULT, J. & KUHNLE, A. 2019 Accelerating deep reinforcement learning strategies of flow control through a multi-environment approach. *Phys. Fluids* **31** (9), 094105.
- RABAUULT, J., REN, F., ZHANG, W., TANG, H. & XU, H. 2020 Deep reinforcement learning in fluid mechanics: a promising method for both active flow control and shape optimization. *J. Hydrodyn.* **32** (2), 234–246.
- RAISSI, M., BABAEI, H. & GIVI, P. 2019a Deep learning of turbulent scalar mixing. *Phys. Rev. Fluids* **4**, 124501.
- RAISSI, M., PERDIKARIS, P. & KARNIADAKIS, G. 2019b Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707.
- RAISSI, M., WANG, Z., TRIANTAFYLLOU, M.S. & KARNIADAKIS, G.E. 2018 Deep learning of vortex-induced vibrations. *J. Fluid Mech.* **861**, 119–137.
- REN, F., RABAUULT, J. & TANG, H. 2021 Applying deep reinforcement learning to active flow control in weakly turbulent conditions. *Phys. Fluids* **33** (3), 037121.
- ROWLEY, C.W., MEZIĆ, I., BAGHERI, S., SCHLATTER, P. & HENNINGSON, D.S. 2009 Spectral analysis of nonlinear flows. *J. Fluid Mech.* **641**, 115–127.
- VON RUEDEN, L., *et al.* 2021 Informed machine learning - a taxonomy and survey of integrating prior knowledge into learning systems. *IEEE Trans. Knowledge Data Engng.* <https://doi.org/10.1109/TKDE.2021.3079836>
- SAAD, Y. 1980 Variations on arnoldi method for computing eigenelements of large unsymmetric matrices. *Linear Algebra Appl.* **34**, 269–295.
- SAHN, M. & OWENS, R.G. 2004 A numerical investigation of wall effects up to high blockage ratios on two-dimensional flow past a confined circular cylinder. *Phys. Fluids* **16** (5), 1305–1320.
- SCHÄFER, M., TUREK, S., DURST, F., KRAUSE, E. & RANNACHER, R. 1996 Benchmark computations of laminar flow around a cylinder. In *Notes on Numerical Fluid Mechanics*, pp. 547–566. Vieweg+Teubner Verlag.
- SCHMID, P.J. 2010 Dynamic mode decomposition of numerical and experimental data. *J. Fluid Mech.* **656**, 5–28.
- SCHULMAN, J., WOLSKI, F., DHARIWAL, P., RADFORD, A. & KLIMOV, O. 2017 Proximal policy optimization algorithms. Preprint. [arXiv:1707.06347](https://arxiv.org/abs/1707.06347).
- SILVER, D., LEVER, G., HEES, N., DEGRIS, T., WIERSTRA, D. & RIEDMILLER, M. 2014 Deterministic policy gradient algorithms. In *International Conference on Machine Learning*, pp. 387–395. PMLR.
- SIPP, D. & LEBEDEV, A. 2007 Global stability of base and mean flows: a general approach and its applications to cylinder and open cavity flows. *J. Fluid Mech.* **593**, 333–358.
- SREENIVASAN, K., STRYKOWSKI, P. & OLINGER, D. 1987 Hopf bifurcation, landau equation, and vortex shedding behind circular cylinders. In *Forum on Unsteady Flow Separation*, vol. 1, pp. 1–13. ASME.
- STRYKOWSKI, P.J. & SREENIVASAN, K.R. 1990 On the formation and suppression of vortex ‘shedding’ at low Reynolds numbers. *J. Fluid Mech.* **218**, 71–107.
- TANG, H., RABAUULT, J., KUHNLE, A., WANG, Y. & WANG, T. 2020 Robust active flow control over a range of Reynolds numbers using an artificial neural network trained through deep reinforcement learning. *Phys. Fluids* **32** (5), 053605.
- TARTAKOVSKY, A.M., MARRERO, C.O., PERDIKARIS, P., TARTAKOVSKY, G.D. & BARAJAS-SOLANO, D. 2020 Physics-informed deep neural networks for learning parameters and constitutive relationships in subsurface flow problems. *Water Resour. Res.* **56** (5), e2019WR026731.

*RL-based control of confined wakes with stability analyses*

- TEZUKA, A. & SUZUKI, K. 2006 Three-dimensional global linear stability analysis of flow around a spheroid. *AIAA J.* **44** (8), 1697–1708.
- THEOFILIS, V. 2011 Global linear instability. *Annu. Rev. Fluid Mech.* **43** (1), 319–352.
- TURTON, S.E., TUCKERMAN, L.S. & BARKLEY, D. 2015 Prediction of frequencies in thermosolutal convection from mean flows. *Phys. Rev. E* **91** (4), 043009.
- VERMA, S., NOVATI, G. & KOUMOUTSAKOS, P. 2018 Efficient collective swimming by harnessing vortices through deep reinforcement learning. *Proc. Natl Acad. Sci.* **115** (23), 5849–5854.
- VIQUERAT, J., MELIGA, P. & HACHEM, E. 2021 A review on deep reinforcement learning for fluid mechanics: an update. arXiv:2107.12206v2.
- WILLIAMSON, C.H.K. 1996 Vortex dynamics in the cylinder wake. *Annu. Rev. Fluid Mech.* **28** (1), 477–539.
- WU, N., KENWAY, G., MADER, C.A., JASA, J. & MARTINS, J.R.R.A. 2020 pyoptsparse: a Python framework for large-scale constrained nonlinear optimization of sparse systems. *J. Open Source Softw.* **5** (54), 2564.
- XU, H., ZHANG, W., DENG, J. & RABAULT, J. 2020 Active flow control with rotating cylinders by an artificial neural network trained by deep reinforcement learning. *J. Hydrodyn.* **32** (2), 254–258.
- YANG, X. & ZEBIB, A. 1989 Absolute and convective instability of a cylinder wake. *Phys. Fluids A: Fluid Dyn.* **1** (4), 689–696.
- ZENG, K. & GRAHAM, M.D. 2021 Symmetry reduction for deep reinforcement learning active control of chaotic spatiotemporal dynamics. *Phys. Rev. E* **104**, 014210.