DRLinFluids—An open-source python platform of coupling Deep Reinforcement Learning and OpenFOAM

Qiulei Wang
School of Civil and Environmental Engineering
Harbin Institute of Technology
Shenzhen 518055, China

Lei Yan
School of Civil and Environmental Engineering
Harbin Institute of Technology
Shenzhen 518055, China

Gang Hu *
School of Civil and Environmental Engineering
Harbin Institute of Technology
Shenzhen 518055, China

Chao Li
School of Civil and Environmental Engineering
Harbin Institute of Technology
Shenzhen 518055, China

Yiqing Xiao
School of Civil and Environmental Engineering
Harbin Institute of Technology
Shenzhen 518055, China

Hao Xiong
School of Mechanical Engineering and Automation
Harbin Institute of Technology
Shenzhen 518055, China

Jean Rabault
Information Technology Department
Norwegian Meteorological Institute
Oslo, Norway

Bernd R. Noack
School of Mechanical Engineering and Automation
Harbin Institute of Technology
Shenzhen 518055, China

ABSTRACT

We propose an open-source python platform for applications of Deep Reinforcement Learning (DRL) in fluid mechanics. DRL has been widely used in optimizing decision-making in nonlinear and high-dimensional problems. Here, an agent maximizes a cumulative reward with learning a feedback policy by acting in an environment. In control theory terms, the cumulative reward would correspond to the cost function, the agent to the actuator, the environment to the measured signals and the learned policy to the feedback law. Thus, DRL assumes an interactive environment or, equivalently, control plant. The setup of a numerical simulation plant with DRL is challenging and time-consuming. In this work, a novel python platform, named DRLinFluids is developed for this purpose, with DRL for flow control and optimization problems in fluid mechanics. The simulations employ OpenFOAM as popular, flexible Navier-Stokes solver in industry and academia, and Tensorforce or Tianshou as widely used versatile DRL packages. The reliability and efficiency of DRLinFluids are demonstrated for two wake stabilization benchmark problems. DRLinFluids significantly reduces the application effort of DRL in fluid mechanics and is expected to greatly accelerates academic and industrial applications.

Keywords Deep Reinforcement Learning · OpenFOAM · Flow control · Computational fluid mechanics · DRLinFluids

---

*Corresponding author: hugang@hit.edu.cn

1 Introduction

In recent years, reinforcement learning (RL), one of the machine learning paradigms, has been developed rapidly with the fast development of computer hardware and corresponding theories. Reinforcement learning is derived initially from behavioral psychology, where organisms frequently implement strategies to seek advantages and avoid disadvantages. The term *reinforcement* was first introduced in 1927 by a Russian physiologist Ivan Pavlov to describe the phenomenon that specific stimuli motivate organisms more inclined to adopt certain particular strategies [Pavlov, 1927]. The stimulus that reinforces behavior is called reinforcer, and the strategy change resulting from reinforcer is called reinforcement learning. Inspired by behavioral psychology, Sutton proposed a formal framework for reinforcement learning [Sutton, 1984, 1988]. The basic idea of RL is that the agent observes the environment, makes decisions, and gets rewards, as shown in Figure 1. The presence of deep learning significantly accelerates the broad applications of reinforcement learning in various fields [LeCun et al., 2015], which has been integrated very effectively into reinforcement learning with the development of deep learning [Mnih et al., 2015]. The value function, policy, and model (state transition function and reward function) were approximated by a deep neural network [Li, 2018], which forms deep reinforcement learning (DRL). A number of novel DRL algorithms were sequentially proposed, such as model-free RL methods, including A2C/A3C [Mnih et al., 2016], PPO [Schulman et al., 2017], DDPG [Lillicrap et al., 2019], SAC [Haarnoja et al., 2018], DQN [Mnih et al., 2013], TD3 [Fujimoto et al., 2018] and model-based RL methods including World Models [Ha and Schmidhuber, 2018], I2A [Weber et al., 2018], MBMF [Nagabandi et al., 2018], MBVE [Feinberg et al., 2018], AlphaZero [Silver et al., 2017] and so on.
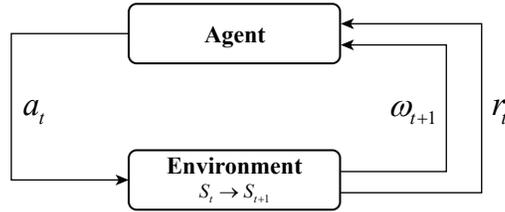
Figure 1: Interaction of agent and environment in reinforcement learning

To date, DRL has been applied widely in the field of fluid mechanics and structure engineering, such as shape optimization, active flow control, and structural reliability assessment and maintenance. Novati et al. [2017] investigated synchronized swimming between two self-propelled swimmers. A deep reinforcement learning algorithm was used to control swimming posture to reduce energy consumption by exploiting the vortex structure and finally improve swimming efficiency. Based on the above study, Verma et al. [2018] highlighted the influence of the past environmental state on the future state, i.e., not following first-order Markov property, using recurrent neural network to consider this effect and accelerate the learning process. The results show that energy savings could be achieved by properly utilizing the wake generated by the forward swimmer. Ma et al. [2018] proposed a DRL-based 2D coupled control system for fluids and rigid bodies, allowing a controller to drive nozzles to move on the boundary of the computational domain and to jet flow into the rigid body to accomplish a series of tasks, such as keeping the rigid body balanced, playing a two-player ping-pong game, and driving the rigid body to hit specified points on a wall in sequence. Lee et al. [2018] used the DoubleDQN method to identify a series of pillars to solve the inertial flow sculpting problem, of which the DRL model has a success rate of 90% in 200,000 episodes. Xiang et al. [2020] applied deep reinforcement learning to structural reliability assessment. using the sampling space and existing samples as states, the agent gives a set of actions along the limit state surface as the next experimental points. The results showed the advantages of the DRL-based approach with highly nonlinear problems. Wei et al. [2020] proposed a DRL-based framework for obtaining optimal structural maintenance strategies. The framework outputs maintenance strategies by observing bridge structures and components.

Active flow control combined with DRL may use a synthetic jet (momentum injection) or moving surface as actuation. Rabault et al. [2019] introduced DRL to active flow control for the first time by using a DRL agent model to control a pair of jets located on the upper and lower sides of a two-dimensional cylinder at low Reynolds number (Re=100). The interactive environment was built by open-source finite-element framework FEniCS [Logg et al., 2012], applying the PPO algorithm to learn the jet actuation strategy and achieving significant drag reduction. Considering that the solution speed of the CFD interactive environment dramatically limits the overall progress under moderate Reynolds number conditions, a further study by Rabault and Kuhnle [2019] performed a simultaneous training method with multiple simulation environments to shorten the training time. Ren et al. [2021] applied the lattice-Boltzmann method (LBM) to establish a CFD environment with weak turbulence conditions by raising the Reynolds number to 1000. Similar to Rabault et al. [2019]'s study, the jet actuators were deployed on the upper and lower sides of the cylinder. The results show that the DRL agent could find an effective feedback law and achieve a drag reduction of more than 30%. On the

other hand, Paris et al. [2021] introduced the S-PPO-CMA method, a novel DRL algorithm, to optimize the sensor layout and investigated the efficiency and robustness of the identified control strategy. The proposed algorithm optimizes the sensor layout and reduces the number of sensors while keeping state-of-the-art performance. In another study, Tang et al. [2020] used four synthetic jets symmetrically located on the upper and lower sides of a cylinder for active flow control. Xu et al. [2020] set up two small rotating cylinders behind the main cylinder obliquely with a Reynolds number of 240, and the rotational speed is controlled by a DRL agent. Fan et al. [2020] experimentally verified the effectiveness and feasibility of wake stabilization by two DRL-controlled rotating control cylinders. In addition, DRL has been applied to a number of chaotic active flow control tasks, including the one-dimensional Kuramoto-Sivashinski equation [Bucci et al., 2019], the one-dimensional falling fluid flow [Belus et al., 2019], and the 2D Rayleigh-Benard convection [Beintema et al., 2020].

The shape optimization of bluff bodies has also experienced substantial progress in past decades. Garnier et al. [2021], inspired by Meliga et al. [2014], placed a two-dimensional square cylinder in a medium Reynolds number flow field and attached a small control cylinder at a fixed position around the cylinder. Both the traditional optimization strategies and deep reinforcement learning PPO/A3C algorithm were tested to find the optimal position of the control cylinder such that the overall drag force of the two cylinders is lower than the single one. To speed up the training process, they first trained the DRL agent at a very low Reynolds number flow (Re=10) and then applied the transfer learning technique to transfer the trained agent model to cases with higher Reynolds numbers. Viquerat et al. [2021] applied DRL for the first time to direct shape optimization, using the Bezier curves technique to generate two-dimensional models with the lift-to-drag ratio of the model's aerodynamic shape as rewards, and the DRL algorithm is able to learn and generate a wing-like optimal shape without any a prior knowledge. Li et al. [2021] introduced a so-called *tacit* domain knowledge, such as transfer learning and meta learning, into naive DRL algorithms to obviously accelerate the process of training, which is very significant for the time-consuming interactive tasks, such as CFD simulation.

As exemplified above, DRL has been extensively and successfully applied to a variety of fluid mechanics problems, and considerable efforts are now being devoted to coupling state-of-the-art CFD solvers and DRL control algorithms, and deploying these effectively on High Performance Computing (HPC) systems [Vinuesa et al., 2022]. In the aforementioned studies, the interactive CFD environments almost relied on self-programming partial differential equation (PDE) solvers based on the Lattice Boltzmann method or Navier-Stokes equation. The employed in-house solvers require significant effort and experience to achieve the efficiency and reliability of well-developed open-source CFD solvers such as OpenFOAM, in particular for complex flow fields. In contrast, OpenFOAM has been developed for more than ten years, and its accuracy and robustness have been demonstrated by substantial research [Ashton and Skaperdas, 2019, Mack and Spruijt, 2013, Robertson et al., 2015, Suvanjumrat, 2017]. Many computational fluid mechanics courses employ OpenFOAM, nowadays. Moreover, OpenFOAM is also increasingly used in industry.

Currently, there is no general and mature platform for simplifying the application of DRL in OpenFOAM simulations. This paper proposes an open-source python platform, DRLinFluids, for coupling DRL and OpenFOAM based on reliable DRL packages, including Tensorforce [Kuhnle et al., 2017] and Tianshou [Weng et al., 2021], and OpenFOAM [Jasak et al., 2007]. DRLinFluids is a flexible and scalable platform to utilize DRL in the field of computational fluid mechanics, even in continuum mechanics. Two case studies of active flow control of bluff bodies are conducted to successfully demonstrate the feasibility and reliability of DRLinFluids. We release DRLinFluids under an open source license on Github (see Appendix), and we expect that this open-source platform will greatly simplify and accelerate the relevant research and application of DRL in fluid mechanics, bluff-body aerodynamics, and wind engineering.

2 Methodology

2.1 OpenFOAM

OpenFOAM is originally developed by Jasak et al. [2007], which is an object-oriented C++ library for complex physics simulations, including structural and CFD analysis. To achieve this purpose efficiently and flexibly, OpenFOAM establishes and solves the corresponding partial differential equations (PDE) of different continuum mechanics physical phenomena with complicated models. Benefiting from the features of object-oriented programming (OOP), especially inheritance, polymorphism, and encapsulation, users can create their own solvers easily for a specific need due to the user-friendly syntax for describing partial differential equations. Besides, OpenFOAM has the following capabilities [Jasak, 2009, Chen et al., 2014]: (1) the ability to handle unstructured polyhedral meshes; (2) widely and reliable out-of-the-box solvers for different fields, such as computational and conjugate heat transfer, combustion, chemical reactions, multiphase flows, mass transfer, particle methods, lagrangian particles tracking, financial analysis and so on; (3) open-source with no payment and a liberal license. Figure 2 illustrates a typically basic directory structure for an OpenFOAM case, which comprises the bare minimum of files required by running a simulation.
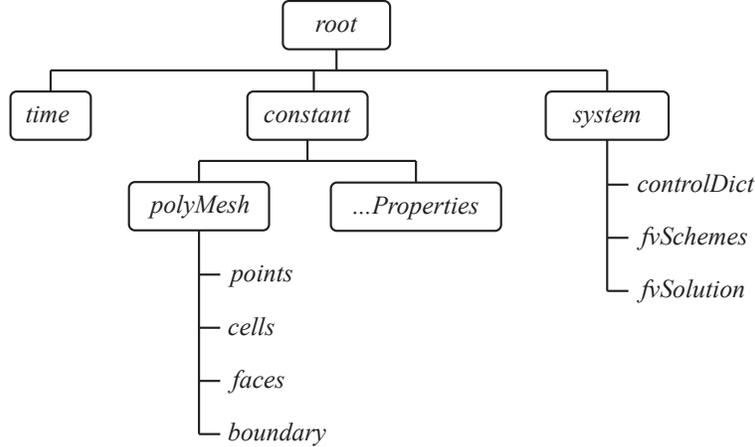
Figure 2: Directory structure of an OpenFOAM case.

The aforementioned features lead OpenFOAM to an ideal solution for research and development, which is well suited as an interactive environment for reinforcement learning in complex CFD problems.

2.2 Reinforcement Learning

Modern reinforcement learning algorithms are mainly based on the Markov Decision Process (MDP), which can be roughly divided into two categories, including model-based and model-free. Model-based methods can predict future states and rewards through environmental models, thereby helping agents to plan better. However, the environment model assumptions may be too idealized to reflect the internal mechanism. On the other hand, dynamic models of environments in the real world are always complex, e.g. the dimension of the state space is too high, and cannot be represented explicitly, making models often unavailable Brunton and Noack [2015]. In contrast, the model-free method does not require the construction of an environment model. The agent interacts directly with the environment and improves its policy performance based on the samples obtained from the exploration. Since the model-free method neither cares about the environment model nor needs to learn the model itself, there is no problem of inaccurate environment fitting, which means that it is relatively easier to implement and train. Model-free methods can be further divided into value-based, policy-based, and combined algorithms.

The value-based methods optimize the action-value function $Q^\pi(s, a)$. The optimal strategy π^* can be obtained by selecting the action corresponding to the maximum value function

$$\pi^* = \mathrm{argmax}_\pi Q^\pi \tag{1}$$

The advantage of the value-based method is that the sampling efficiency is relatively high, the estimated variance of the value function is small, and it is not easy to fall into a local optimum; the disadvantage is that it usually cannot deal with continuous action space problems, and the final strategy is usually a deterministic strategy rather than a probability distribution. In addition, the ϵ-greedy strategy and the max operator in algorithms such as deep Q-networks are prone to overestimation problems.

The policy-based method directly optimizes the policy, and maximizes the cumulative reward by iteratively updating the policy. Compared with value-based methods, policy-based methods have the advantages of simple policy parameterization, fast convergence, and are suitable for continuous or high-dimensional action spaces. As the foundations and main aspects of common DRL algorithms have been discussed extensively in both the DRL-centric literature [Arulkumaran et al., 2017, Schulman et al., 2017, Lillicrap et al., 2019, Haarnoja et al., 2018], and the fluid mechanics literature [Garnier et al., 2021, Rabault et al., 2020], we refer the reader curious of more details about the DRL algorithms in itself to the corresponding references.

2.3 OpenFOAM and Reinforcement Learning coupling schemes

At present, there are many well-known deep reinforcement learning frameworks, such as DeeR [François Lavet et al., 2016], Dopamine [Castro et al., 2018], ELF [Tian et al., 2017], OpenAI baselines [Dhariwal et al., 2017], rllab [Duan et al., 2016], Keras-RL [Plappert, 2016], Coach [Caspi et al., 2017], TF-Agents [Guadarrama et al., 2018], Acme

Table 1: Summary of the built-in low-level RL class.

OpenFoam DRL Class
Attributes
- Sensor location and identifier
- Controller location and identifier
- Controller Type and Output (action) Range
- Environmental parameters (e.g. inflow conditions, fluid properties, etc.)
Method
- Calling OpenFOAM operations, adding, deleting, and changing OpenFOAM configuration files
- DRL algorithms with Tensorforce as the backend
Functionality
- Flexible and easy to call OpenFOAM as a reinforcement learning interaction environment
- Suitable for a wide range of DRL Agent output (actions) types, such as jet velocity, rotating column speed, etc.
- Wide range of application scenarios, combined with OpenFOAM programming, to implement DRL on any PDE problem (e.g. CFD, finance, etc.)
- Reinforcement learning algorithms provided by Tensorforce, Tianshou, or custom algorithm.

[Hoffman et al., 2020]. Most of them are developed based on Python, with PyTorch and TensorFlow as automatic gradient solvers. Tensorforce [Kuhnle et al., 2017] and Tianshou [Weng et al., 2021] packages are used as an alternative ready-to-use backend to provide a wide range of reinforcement learning algorithms in the DRLinFluids platform. Besides, custom RL methods can also be easily integrated into the DRLinFluids platform by inheriting the built-in DRLinFluids's DRL class.

To ensure DRLinFluids flexible and scalable, a low-level RL class and high-level training and evaluation platform are defined separately. The object-oriented programming (OOP) paradigm are adopted to design the RL class to improve code reusability, which organizes platform design around data and objects.

The RL class in DRLinFluids is defined as a standard Gym environment to complete more flexible functions. Gym [Brockman et al., 2016] is a toolkit for developing and comparing reinforcement learning algorithms. It makes no assumptions about the structure of agents and is compatible with any numerical computation library, such as TensorFlow [Abadi et al., 2015] or PyTorch [Paszke et al., 2019]. The architecture of DRLinFluids are shown as Figure 3.
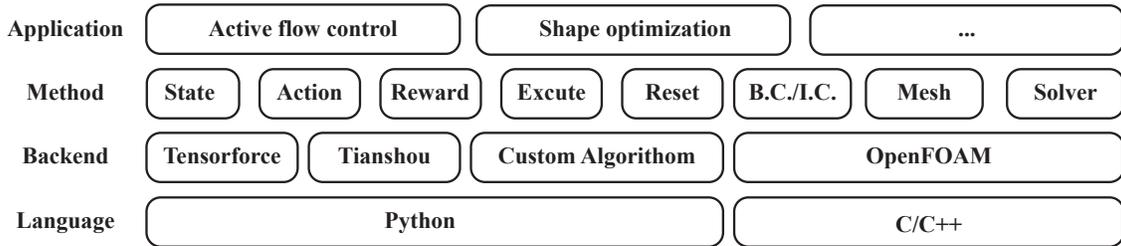
Figure 3: Architecture of the DRLinFluids reinforcement learning platform.

2.3.1 Low-level RL class

The RL-OpenFOAM class takes the following parameters as input: (1) the environment parameters, including the vanilla OpenFOAM case, root path, parallel parameters, simulation dimensions, etc.; (2) state parameters, such as the type and location of probes; (3) agent parameters, such as the range of actions, interaction period between adjacent actions, and (4) learning parameters, such as actuator velocity. Member function includes execute and reset. The details of the class are shown in Table 1.

2.3.2 High-level Platform

To further simplify the application of DRL coupled with OpenFOAM based on the RL-OpenFOAM class, this section introduces Tensorforce and Tianshou as the backend algorithm, and builds a high-level reinforcement learning training and evaluation platform. Generally, the procedure of the reinforcement learning shown in Figure 4 can be described as the following steps:

1. pre-defining a vanilla OpenFOAM simulation, including initial/boundary conditions, mesh, discrete schemes (fvSchemes), solver settings (fvSolution), and organizing these files according to the directory hierarchy specified by OpenFOAM.

2. importing and inheriting the RL-OpenFOAM class from DRLinFluids, overwriting the reward function.

3. setting relevant RL parameters and passing them to the RL-OpenFOAM object or instance.

4. training, evaluating, saving, and deploying the model in practice.

To avoid restriction introduced by OpenFOAM, this platform does not use built-in utilities. Instead, the RegExp (regular expression) and template engine are adopted to interact between OpenFOAM and Python backend according to the file input/output (I/O) stream. RegExp is a widely used method in text processing, which specifies a search pattern and uses it to search or substitute matching content. Generally, A number of ways exist to exchange or share data, such as serialization in shared memory or file I/O stream. The latter is a simple but effective method for exchanging data, writing data from memory to hard disk, and reading from hard disk to memory. It is not advisable to use this method in the HPC-like tasks because the file I/O is likely a speed bottleneck in the whole system. According to a previous CFD-based RL study [Rabault and Kuhnle, 2019] and our experience which will be further discussed in section 3, the CFD simulation part accounts for over 99% of the total computation time. Therefore, the I/O stream is highly suitable for CFD-based RL tasks, in which CFD simulation is very time-consuming compared to the communication in the I/O stream, even agent policy optimization.
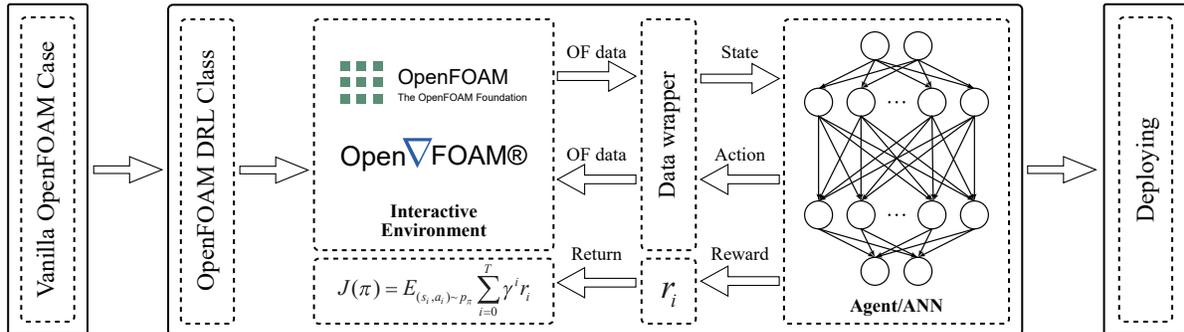
Figure 4: Procedure of a reinforcement learning instance transferred from a vanilla OpenFOAM case in DRLinFluids.

The interaction details inside DRLinFluids are shown in Figure 5, which is also a standard procedure from scratch. The DRLinFluids will call the OpenFOAM solver to generate an initial environment with states at the beginning. Then the loop of interaction between agent and environment, and policy network optimization are carried out step by step.

3 Case studies

To demonstrate the feasibility and reliability of DRLinFluids, two case studies of flow control of bluff bodies are tested based on DRLinFluids. The first case is an adaptation of a circular cylinder flow control conducted by Rabault et al. [2019]. The second case is a square cylinder flow control by jet flow at trailing edges of two sides. The successes of DRLinFluids in these two cases well prove the feasibility and reliability of DRLinFluids. The details of these case studies are documented below.

3.1 Case 1: Active flow control of a 2D circular cylinder

Flow around a circular cylinder has been studied extensively in the past century. When flow passes a circular cylinder, a series of vortexes shed from its two sides and form the well-known Kármán vortex street. The vortex shedding results in increased aerodynamic drag and unsteady lift forces acting on the cylinder, both of which are normally not desirable.

Figure 5: Detail procedures inside DRLinFluids.

In this study, two synthetic jets are located on the upper and lower side of the circular cylinder and controlled by DRL algorithm in order to minimize the drag and lift forces of the cylinder in an energy-efficient way. The whole process of this study is conducted based on DRLinFluids.

3.1.1 Numerical Simulation Model

The simulation environment is built based on OpenFOAM, and unstructured meshes are adopted for the CFD simulation. The mesh consists of 16200 triangular elements with increased resolution near the cylinder. The diameter of the cylinder is D. The length and width of the computational domain are $L = 22D$ and $B = 4.1D$ respectively as shown in Figure 6(a). Similar to the benchmark case [Schäfer et al., 1996], the cylinder is located slightly away from the centre of the computational domain, in order to trigger the vortex shedding. The distance from the outlet to the leeward side of the cylinder is $19.5D$, so that the wake can fully develop. Using a non-dimensional time step $dt = 5 \times 10^{-4}$, the CFL number ($\frac{dt|U|}{\delta x}$, $|U|$ is the modulus of the velocity vector in a grid unit, δx is the grid length towards the velocity direction) is less than 1, which meets the accuracy requirements.

The inlet boundary Γ_{in} uses the velocity inlet boundary condition, and the inflow velocity is set to a parabolic channel flow profile with a mean velocity magnitude of U. The corresponding Reynolds number Re is 100 ($Re = \frac{UD}{\nu}$, ν is the kinematic viscosity). An outflow boundary condition Γ_{out} is set for the outlet of the domain. The non-slip wall boundary condition Γ_W is applied on the top and bottom of the channel, as well as on the cylinder surfaces (see Figure 6(b)).

The flow control action is performed by two jet holes (Γ_1 and Γ_2) at two sides of cylinder. A parabolic velocity distribution with a jet width of $\omega = 10°$ is imposed at these two jet holes, as shown in Figure 6(b). The direction of the jet actuators is perpendicular to the cylinder wall, which implies that the drag reduction results from the indirect flow control, rather than direct injection of momentum. Meanwhile, the sum of the velocities of the jets is zero, i.e. $V_{\Gamma_1} = -V_{\Gamma_2}$.

The probes sensing the state for this case can utilize either velocity or pressure probes. For this case, only velocity probes are setup and they are arranged in two circles around the cylinder and also in the wake region with a total of 149 probes, as shown in Figure 6(a). These probes monitor the horizontal and vertical components of the velocity field, and hence the agent can perceive detailed information of the flow field and learn from vortex shedding patterns.

Figure 6: Description of the numerical setup adapted from Schäfer et al. [1996]: (a) unsteady pressure field around the circular cylinder after flow initialization without active control. The locations of the 149 velocity probes are indicated by the black dots. The locations of 2 jets are indicated by the red dots; (b) details of the cylinder and the jet actuators.

3.1.2   Formulation of optimizing flow control

The feedback stabilization of a circular cylinder is a typical optimization problem and leads reduced drag and vanishing lift forces. The use of DRL algorithms is essentially to find a control law that optimizes a given reward function in order to minimize drag and lift forces. This control law can make the reward function show a certain upward trend when it is properly trained. The proposed reward function is constructed as a combination of drag and lift coefficients in this study. The reward function is set as follows:

$$r_t = (C_D)_{T_0} - (C_D)_T - 0.1|(C_L)_T| \tag{2}$$

where $(C_D)_{T_0} = 3.205$ is the mean drag coefficient of the circular cylinder without flow control, $(\cdot)_T$ representing the sliding average over the duration of one jet flow control period T corresponding to active flow control cylinder.

The reward function in Equation (2) has been verified to be better than simply using the instantaneous drag coefficient, i.e. $r_t = -(C_D)_T$. First, only considering the change in drag coefficient could cause the lift to fluctuate too much, the algorithm tends to make the control value run to the extreme value. Furthermore, it has been proved that adding lift can improve the learning speed and stability. Second, the drag reduction is defined as $\Delta C_D = (C_D)_{T_0} - (C_D)_T$ relative to uncontrolled flow. This equation is designed to be as positive as possible, and as the training effect gets better, the reward function tends to increase while the drag value decreases. Meanwhile, it is necessary to keep changes in the drag (ΔC_D) and in the lift ($\Delta C_L = 0.1|(C_L)_T|$) the same order of magnitude, so that the algorithm does not only focus on drag reduction and ignore changes in lift during training. After taking the above points into account, the reinforcement learning algorithms could try to maximize r_t, by reducing both drag and lift during training.

The jet actuation is smoothed to ensure a continuous changing control signal without excessive lift fluctuations due to a sudden change in the jet velocity. To this end, the control action is then adjusted from one non-dimensional time step in the simulation to the next by

$$V_{\Gamma_1(t)} = V_{\Gamma_1(t-1)} + \alpha[a - V_{\Gamma_1(t-1)}] \tag{3}$$

where $\alpha = 0.1$ is a numerical parameter determined by trial and error, $V_{\Gamma_1(t)}$ and $V_{\Gamma_1(t-1)}$ is the jet flow velocity adopted at non-dimensional time instant t and $t-1$ respectively, and a is one jet flow velocity for the current 50 time steps proposed by the DRL agent.

3.1.3   DRL algorithm configuration and simulation results

The DRL framework uses the DRLinFluids framework, and the agent is trained by using the Proximity Policy Optimization (PPO) method, which is loaded from Tensorforce platform [Kuhnle et al., 2017]. The PPO method is a

policy gradient algorithm, and its training alternates between sampling data interactively with the environment and optimizing a surrogate objective function by using stochastic gradient ascent. It is one of the most commonly used algorithms in DRL [Schulman et al., 2017]. The ANN architecture consists of two dense layers with 512 fully connected neurons, the input layer receives data from velocity probes, and the output layer produces a jet velocity. The training loop of DRL is shown in Figure 4. Training is performed using a parallel environment, where the initial state of each environment is obtained by performing simulations without active control until a fully developed wake with a Karman vortex street. The state of the flow field at this point is stored and used as the starting point for subsequent learning episodes. At the beginning of the learning process, the PPO agent interacts and updates the ANN hyperparameters every 50 time steps. This numerical trick and the smoothing strategy described in Eq. (2) together ensure a continuous control signal, which is crucial for a successful DRL flow control process.
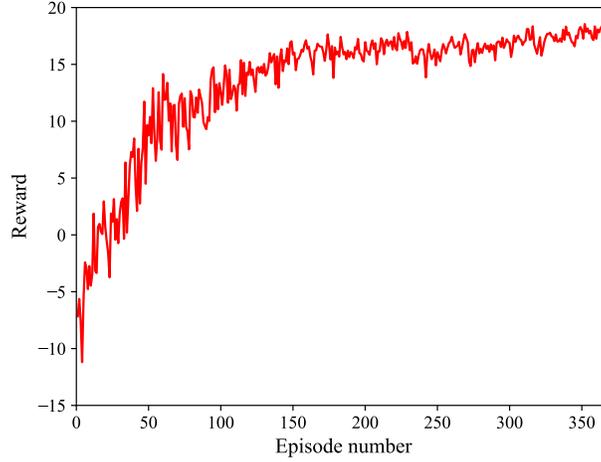
Figure 7: Reward curve from parallel environment of DRL-based active flow control for circular cylinder.

Using those methods, and choosing an episode duration period $T_{\max} = 2.5$ (corresponds to 5000 numerical non-dimensional time steps, i.e. 100 jet control flow set by the agent), the PPO agent is able to learn a control strategy through parallelizing the data sampling from the epochs. In Figure 7, it is a single training environment reward development process corresponding to 4 environments parallel training. After approximately 200 epochs, the reward is closed to the maxima. Fully stabilized control strategy is obtained with further epochs so that the purpose of tuning the policy would be achieved. Furthermore, it has been proved that no more improvement is obtained if the network is allowed to train for a longer duration.

As shown in Figure 8, the artificial neural network has been successfully trained by the PPO algorithm in the DRLinFluids framework and is able to perform active flow control to continuously reduce drag and suppress lift. As shown in Figure 8(a), in the absence of actuation, the drag coefficient C_D oscillates periodically around a mean value. The mean value of drag coefficient is 3.205, and its standard deviation value is 0.0236. The standard deviation value of the lift coefficient is 0.74. With DRL-based active flow control (denoted by DRL-cylinder), the mean drag coefficient of the cylinder is reduced to 2.95, corresponding to a drag reduction of approximately 8%, which is the optimal drag reduction attainable in this configuration, as discussed in Rabault et al. [2019]. In addition, the fluctuation of the drag coefficient is also significantly suppressed, corresponding to the standard deviation value of 0.0028, only 11.86% of that the unforced cylinder. On the other hand, in Figure 8(b), the lift coefficient fluctuates at the very early stage, while the fluctuation is almost fully suppressed when a proper control strategy is found after a certain number of iterations of training. The standard deviation value of the lift coefficient is significantly reduced to 0.032, only 4.32% that of the plain cylinder, which implies that the vortex shedding is almost fully suppressed.

Power spectrum analyses of C_D and C_L of the cylinder with and without active flow control are conducted and presented in Figure 8(c) and (d). An obvious peak can be observed from the power spectrum curves for both C_D and C_L of the plain cylinder, which implies regular vortex containing considerable energy and shedding from the plain cylinder. In contrast, the peaks disappear in the power spectrum curves of C_D and C_L of the cylinder with active flow control. That is to say, the regular vortex shedding has been destroyed by the jet flow.

Figure 9 presents the instantaneous flow field of the circular cylinder with or without active control. The effect of the controlled jet flow on the mitigation of aerodynamic force of the cylinder could be further elaborated in terms of the flow pattern. In Figure 9(a), an alternative vortex shedding pattern is clear for the plain cylinder as expected. As is
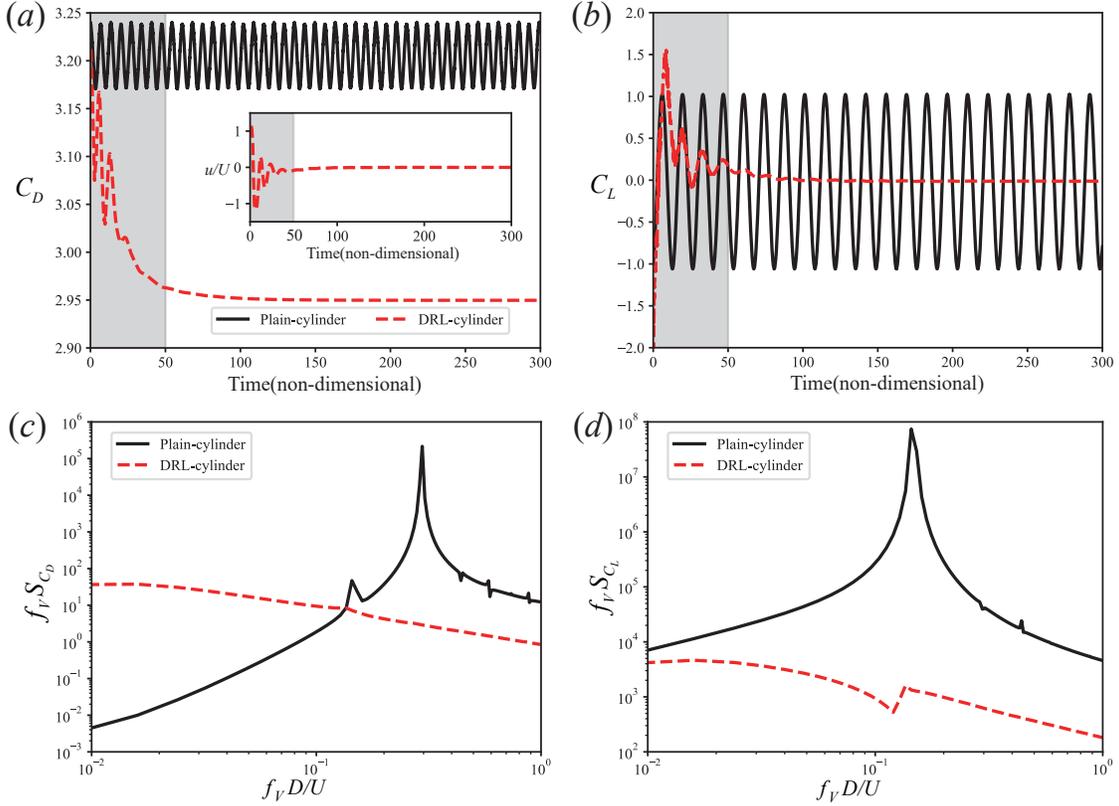
Figure 8: (a) Time-resolved value of drag coefficient C_D for the cylinder without (Plain-cylinder) and with (DRL-cylinder) active flow control, and corresponding normalized velocity flow rate of the jet flow Γ_1; (b) Time-resolved value of the lift coefficient C_L for the cylinder without (Plain-cylinder) and with (DRL-cylinder) active flow control; (c) power spectral density of the drag coefficient C_D during the period of non-dimensional time ranging from 50 to 300; (d) power spectral density of the lift coefficient C_L during the period of non-dimensional time ranging from 50 to 300.
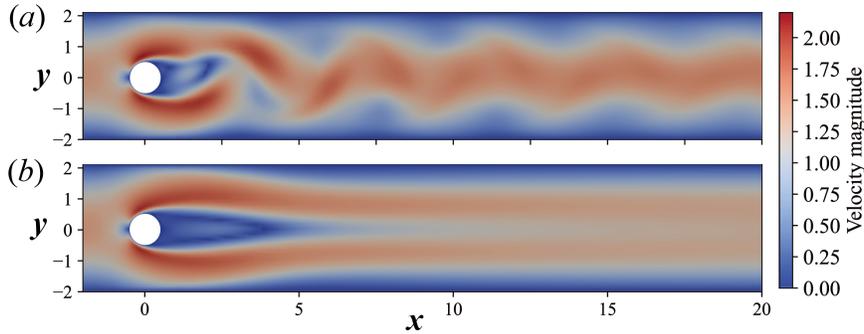
Figure 9: Comparison of representative snapshots of the velocity field around the cylinder (a) without and (b) with active flow control.

well known, this alternative vortex shedding pattern directly results in the fluctuations in both drag and lift coefficients which shown in Figure 8. In terms of the evolution of both force coefficients and the jet velocity shown in Figure 8, the flow control process can be divided into two stages. First, the DRL agent modifies the vortex shedding by imposing a relatively large jet flow in the period of non-dimensional time ranging from 0 to approximately 50, which reduces drag and lift considerably. Second, after this unstable control stage, the jet flow fluctuates weakly and gradually approaches zero. Only a very weak jet flow is required in the later stage when the vortex shedding is stabilized. The corresponding ratio of the jet speed to the incoming wind speed u/U is only 0.58%. In Figure 9(b), with small actuation, the alternative vortex shedding is suppressed, leading to the reduction in the fluctuations of both C_D and C_L. In the mean time, an

elongated re-circulation zone forms in the near wake. This elongated zone is associated with increasing pressure in the wake and hence reduces the drag force of the cylinder. Apparently, the DRL agent has learned an effective and efficient strategy for controlling the jet flow for minimizing the drag and lift of the cylinder.

3.2 Case 2: Active flow control of a 2D square cylinder

A cylinder with a square cross section has also been widely used in practical engineering such as tall buildings, bridge pylons. The square cylinder is also a classical bluff body suffering from strong flow-induced instability. Vortex shedding from the square cylinder induces alternative force acting on the cylinder, and this force is able to cause strong crosswind vibrations and hence potential structural failure. As a result, mitigation of aerodynamic forces of square cylinders has attracted substantial research in the past. In this study, two synthetic jets are located on the trailing edges of two sides of the square cylinder and intelligently controlled by DRL in order to minimize its drag and lift in an energy-efficient manner. The entire training process is implemented using DRLinFluids.

3.2.1 Numerical Simulation Model

The length/width of the square cylinder is D, and the length and width of the computational domain are $L = 40\,D$ and $B = 20\,D$, respectively, as shown in Figure 10(a). The distance from the outlet to the leeward face of the square cylinder is $29.5\,D$, thus minimizing the effect of the outflow condition on vortex shedding. Structural meshes with 23125 quadrilateral elements are adopted for the CFD simulation, and they are refined around the square cylinder. A non-dimensional time step $dt = 5 \times 10^{-4}$ is adopted. The corresponding CFL number is less than unity, which meets typical accuracy requirements.
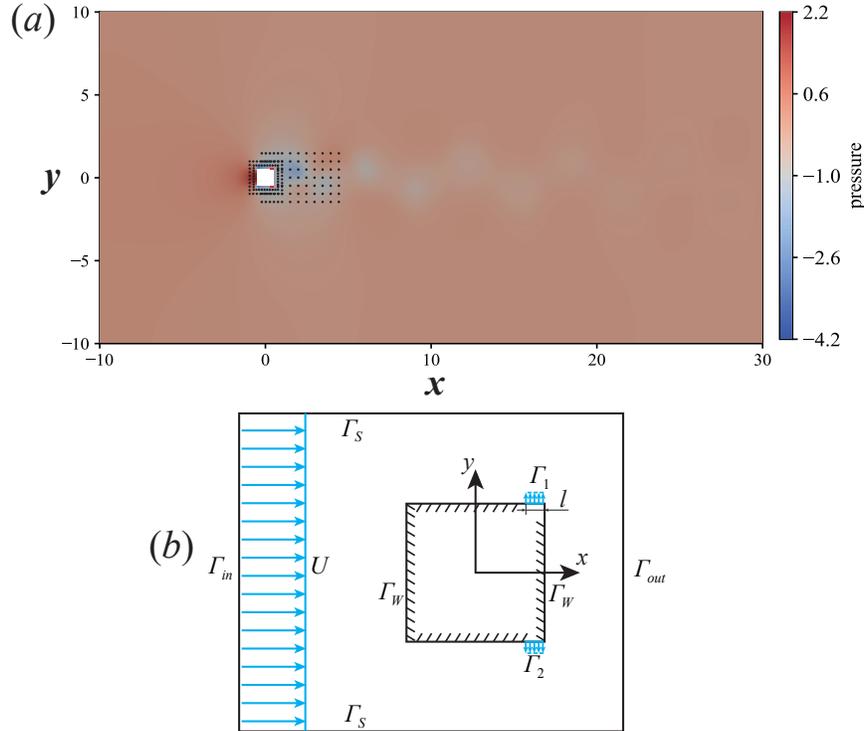
Figure 10: Description of the numerical setup: (a) unsteady pressure field around the square cylinder after flow initialization without active control. The location of the velocity probes is indicated by the black dots. The location of the control jets is indicated by the red dots; (b) details of the square cylinder and the jet actuators.

At the inlet boundary Γ_{in} uniform flow with velocity U is imposed. The corresponding Reynolds number Re is 100. An outflow boundary condition Γ_{out} is set for the outlet of the domain. A symmetrical boundary condition Γ_S is applied to the top and bottom of the domain, and a non-slip wall boundary condition Γ_W is adopted for the surface of the square cylinder (see Figure 10(b)).

The flow control action is implemented by two jet flows (Γ_1 and Γ_2) located on the trailing edges of two sides of the square cylinder. A uniform flow velocity distribution with a jet width $l = D/25$ is applied at these two jet holes, as shown in Figure 10(b). Similar to the first case study, the sum of the jet velocities is zero, i.e. $V_{\Gamma_1} = -V_{\Gamma_2}$. The state of the DRL model is set as flow velocities around the cylinder monitored by a total of 147 velocity probes distributed in the vicinity of the cylinder as shown in Figure 10(a).

3.2.2 DRL algorithm configuration and simulation results

The DRLinFluids framework is used as the DRL framework. The agent is trained using the Soft Actor-Critic (SAC) method, which is loaded from Tianshou Deep Reinforcement Learning Library[Weng et al., 2021]. SAC is an off-policy actor-critic DRL algorithm based on a maximum entropy reinforcement learning theory. The actor's goal is to maximize expected reward as well as maximize entropy, and succeed in reaching the desired value while acting as randomly as possible. Since it is an off-policy algorithm, training can be performed efficiently with limited samples. More details can be found in Haarnoja et al. [2018]. The ANN architecture also consists of two dense layers of 512 fully connected neurons, identical to that in the first case study. The input layer receives data from those velocity probes, and the output layer is the jet velocity. Since the vortex shedding frequency of square cylinder is higher, the SAC agent interacts and updates the ANN hyperparameters every 26 time steps. Training is performed using 5 parallel environments, each with an initial state of a fully developed stable wake and used as a starting point for subsequent training.
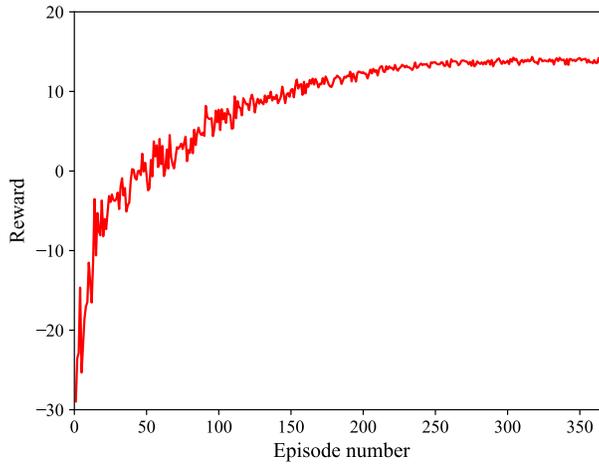
Figure 11: Reward curve from parallel environment of DRL-based active flow control for square cylinder.

Through choosing an episode duration period $T_{\max} = 1.3$ (corresponds to 2600 numerical non-dimensional time steps, i.e. 100 jet flow set by the agent), the SAC agent is able to learn a control strategy through parallelizing the data sampling from these epochs. In Figure 11, it is a single training environment reward development process corresponding to 5 environments parallel training. After typically approximately 250 epochs, the reward reaches its peak value, demonstrating that a stabilized control strategy is obtained. In fact, considerable computation time is spent in the flow simulation. Such a simple and fast simulation setup makes reproduction of the results easily.

As shown in Figure 12, it is clear that the artificial neural network has been successfully trained by the SAC algorithm in the DRLinFluids framework to intelligently control the jet flow in order to minimize the aerodynamic forces of the square cylinder. For the plain square cylinder, the mean drag coefficient is 1.549, and its standard deviation value is 0.0073. The standard deviation value of the lift coefficient is 0.227. With the active flow control, the mean drag coefficient of the cylinder is reduced to 1.337, corresponding to a drag reduction of approximately 13.7%. In addition, the standard deviation of the drag coefficient is decreased to 0.0061. More importantly, the standard deviation of the lift coefficient is significantly reduced to 0.0134, only 5.9% of that the plain square cylinder. The suppression of the lift force is desirable for the mitigation of flow-induced instability of the square cylinder.

Power spectrum analyses of C_D and C_L of the square cylinder with and without active flow control are conducted and presented in Figure 12(c) and (d). An obvious peak value can be observed from the power spectrum curves for both C_D and C_L of the plain cylinder, which shows regular vortex containing considerable energy and shedding from the plain square. On the contrary, with active flow control, the peak is eliminated in the power spectrum curves of C_D of the square cylinder. It shows that the regular vortex shedding has been fully disrupted by the jet actuation.
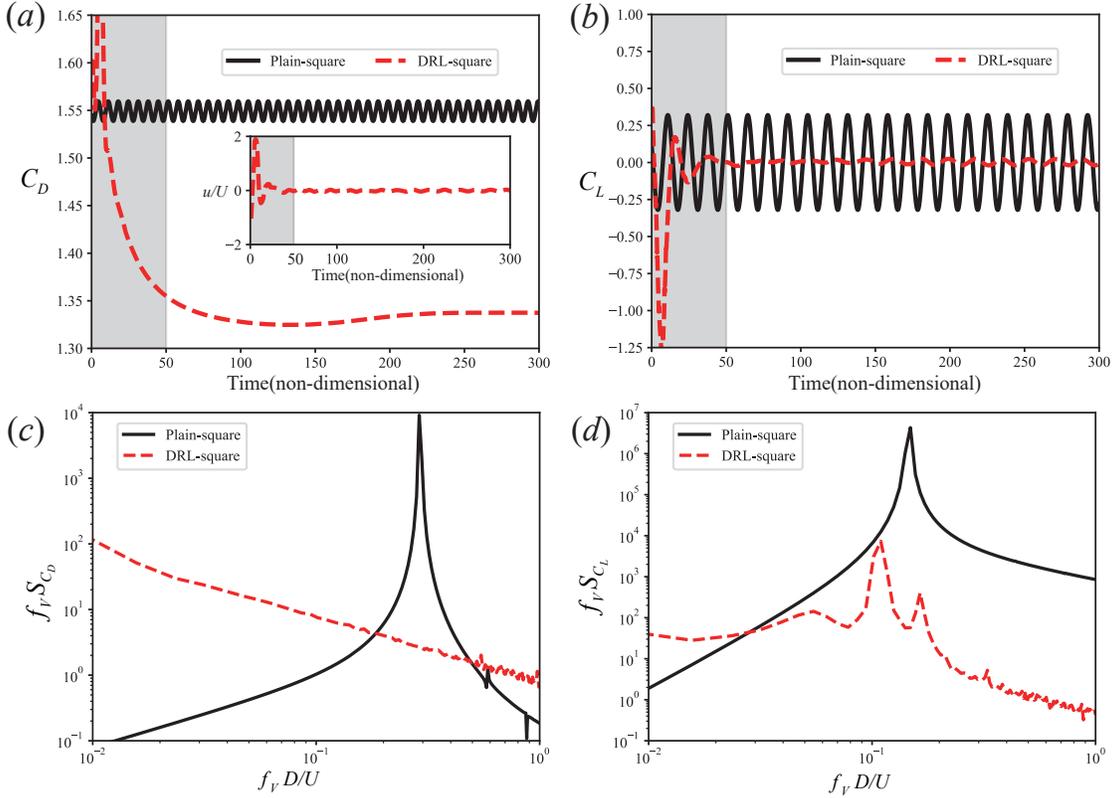
Figure 12: (a) Time-resolved value of the drag coefficient C_D of the square cylinder without (Plain-square) and with (DRL-square) active flow control, and corresponding normalized velocity flow rate of the control Γ_1; (b) Time-resolved value of the lift coefficient C_L of the square case without (Plain-square) and with (DRL-square) active flow control; (c) power spectral of the drag coefficient C_D during the period of non-dimensional time ranging from 50 to 300. (d) power spectral of the lift coefficient C_L during the period of non-dimensional time ranging from 50 to 300.

Two successive states can be observed from the time histories of aerodynamic forces and jet velocities shown in Figure 12(a). First, the DRL agent performs relatively large instantaneous jet to greatly reduce drag during the period of non-dimensional time ranging from 0 to approximately 50. Afterward, the controlled jet flow is stable with a very small velocity. The corresponding ratio of the peak jet velocity to the incoming wind speed u/U is only 3.443%. Apparently, the DRL agent has found an effective and efficient solution to maintain a small drag and a very weak lift for the square cylinder.

Figure 13 shows the instantaneous flow field of the square cylinder with or without active flow control. Similarly, the effect of the controlled jet flow on the mitigation of aerodynamic force of the square cylinder could be elaborated in terms of the flow pattern. In Figure 13(a), flow separates at the windward corners of the cylinder, and forms large-scale vortex in the near wake of the cylinder. The large-scale vortex alternatively sheds at two sides of the square cylinder, forming a clear vortex street. This alternative vortex shedding pattern directly results in the fluctuations in both drag and lift coefficients shown in Figure 12. In contrast, the vortex shedding for the square cylinder with active flow control occurs far away from the cylinder. That is to say, the effect of the alternative vortex shedding on the fluctuation of aerodynamic forces is indirect and hence very weak. That is why the fluctuation of both drag and lift coefficients is suppressed dramatically. On the other hand, the reduction in the mean drag coefficients results from the elongated wake. The elongated wake implies a reduced curvature of the shear layer and corresponds to an increased pressure at the rearward cylinder side. This increased pressure causes a smaller drag for the cylinder with active flow control.

4 Conclusion and future work

This study has developed a flexible, scalable, and user-friendly platform, DRLinFluids, for connecting reinforcement learning with fluid mechanics, and two case studies have been tested for validation. DRLinFluids uses one of most
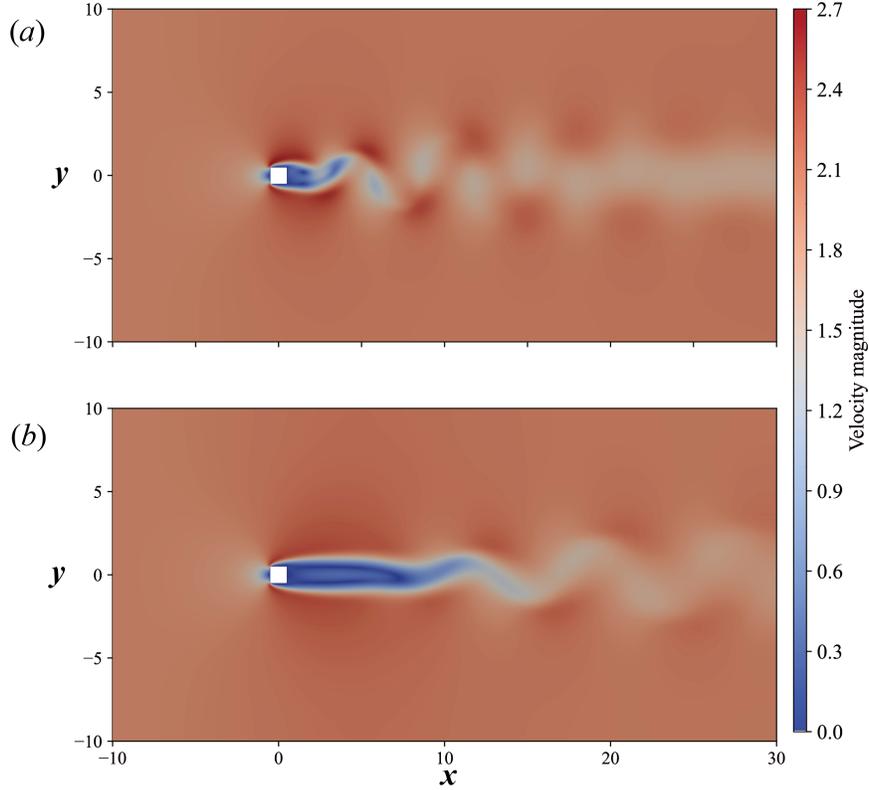
Figure 13: Comparison of representative snapshots of the velocity field around the square cylinder (a) without and (b) with active flow control.

used open-source CFD software, OpenFOAM, as the environment to interact with the agent in different reinforcement learning algorithms from a built-in or a custom DRL backend. To share data and call different modules flexibly, DRLinFluids leverages regular expressions to handle OpenFOAM dictionary files.

The use of DRLinFluids successfully carries out a jet actuation control policy, which remarkably reduces the drag and lift force acting on the circular and square cylinders, and suppresses vortex shedding in their wake region. To find the optimal feedback control, the policy network is continuously updated during interaction to maximize the reward function, which takes both drag and lift coefficients into consideration. A significant drag reduction by up to about 8% and 13.6% occurs for circular cylinder and square cylinder respectively. Besides, the lift fluctuations are also suppressed to a very low level and the regular vortex shedding has been destroyed by the jet flow. It is noteworthy that similar control strategies are obtained despite two different reinforcement learning algorithms being adopted in the two cases. The jet velocity starts at a relatively large jet flow, fluctuates weakly, and eventually reaches a very low level. The corresponding ratio of the jet velocity to the incoming wind velocity is only 0.58% and 3.443% for the circular cylinder and square cylinder cases respectively. Apparently, the DRL agent has found an effective and efficient solution to maintain a small drag and a very weak lift for the cylinders. It shows that under a low Reynolds number ($Re = 100$), active flow control developed based on DRLinFluids can effectively eliminate the adverse effect of vortex shedding on drag and lift of both circular and square cylinders with only a small amount of hyper-parameter tuning required.

These two validation cases demonstrate the effectiveness and reliability of DRLinFluids in DRL-based active flow control. We demonstrate the scaling of the training to both parallelization of the CFD simulation itself through the OpenFOAM framework, and the higher-level parallelization of the DRL episodes gathering by using several environments in parallel. This double scaling is critical for further applications, where increasingly complex flows are considered. More generally, since OpenFOAM has powerful application functions in CFD, we believe that DRLinFluids have a much broader application outlook in the field of fluid mechanics, which can bridge the gap between cutting-edge DRL algorithms and traditional fluid mechanics research or rapid implementation of DRL applications in the relevant industries. DRLinFluids is, therefore, an important step towards the application of DRL control to full-scale, realistic engineering configurations. Next steps in the development of DRLinFluids include the development of 3D active flow control benchmarks, testing on large scale supercomputers, and the development of automated methodologies to define

invariants and symmetry-aware strategies that take advantage of the structure of the dynamic system to control, which are the key elements needed to apply DRL on increasingly complex flow control problems.

5 Acknowledgement

Appendix

The DRLinFluids framework is released under an open source license on Github, at the following URL: `https://github.com/venturi123/DRLinFluids` [NOTE: the page is currently an empty placeholder, and the actual code will be released upon publication of the manuscript in the peer review literature]. We invite all users to further discuss and ask for help directly on Github, through the issue system, and we commit to helping develop a community around the DRLinFluids framework by providing in-depth documentation and help to new users.

References

Ivan P. Pavlov. *Conditioned reflexes: an investigation of the physiological activity of the cerebral cortex*. Oxford University Press, London. A PHYSIOLOGICAL MODEL OF PHOBIC ANXIETY A, 1927.

Richard Stuart Sutton. *Temporal credit assignment in reinforcement learning*. PhD thesis, University of Massachusetts Amherst, 1984.

Richard S. Sutton. Learning to predict by the methods of temporal differences. *Machine Learning*, 3(1):9–44, 8 1988. ISSN 1573-0565. doi:10.1007/BF00115009.

Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.

Yuxi Li. Deep reinforcement learning: An overview. *arXiv:1701.07274 [cs]*, 11 2018. URL `http://arxiv.org/abs/1701.07274`. arXiv: 1701.07274.

Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning*, pages 1928–1937. International Conference on Machine Learning, PMLR, 6 2016. URL `http://proceedings.mlr.press/v48/mniha16.html`. ISSN: 1938-7228.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv:1707.06347 [cs]*, 8 2017. URL `http://arxiv.org/abs/1707.06347`. arXiv: 1707.06347.

Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv:1509.02971 [cs, stat]*, 7 2019. URL `http://arxiv.org/abs/1509.02971`. arXiv: 1509.02971.

Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. PMLR, 2018.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv:1312.5602 [cs]*, 12 2013. URL `http://arxiv.org/abs/1312.5602`. arXiv: 1312.5602.

Scott Fujimoto, Herke van Hoof, and David Meger. Addressing function approximation error in actor-critic methods. *arXiv:1802.09477 [cs, stat]*, 10 2018. URL `http://arxiv.org/abs/1802.09477`. arXiv: 1802.09477.

David Ha and Jürgen Schmidhuber. World models. *arXiv:1803.10122 [cs, stat]*, 3 2018. doi:10.5281/zenodo.1207631. URL `http://arxiv.org/abs/1803.10122`. arXiv: 1803.10122.

Théophane Weber, Sébastien Racanière, David P. Reichert, Lars Buesing, Arthur Guez, Danilo Jimenez Rezende, Adria Puigdomènech Badia, Oriol Vinyals, Nicolas Heess, Yujia Li, Razvan Pascanu, Peter Battaglia, Demis Hassabis, David Silver, and Daan Wierstra. Imagination-augmented agents for deep reinforcement learning. *arXiv:1707.06203 [cs, stat]*, 2 2018. URL `http://arxiv.org/abs/1707.06203`. arXiv: 1707.06203.

Anusha Nagabandi, Gregory Kahn, Ronald S. Fearing, and Sergey Levine. Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7559–7566. 2018 IEEE International Conference on Robotics and Automation (ICRA), 5 2018. doi:10.1109/ICRA.2018.8463189. ISSN: 2577-087X.

Vladimir Feinberg, Alvin Wan, Ion Stoica, Michael I. Jordan, Joseph E. Gonzalez, and Sergey Levine. Model-based value estimation for efficient model-free reinforcement learning. *arXiv:1803.00101 [cs, stat]*, 2 2018. URL `http://arxiv.org/abs/1803.00101`. arXiv: 1803.00101.

David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, and Demis Hassabis. Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv:1712.01815 [cs]*, 12 2017. URL `http://arxiv.org/abs/1712.01815`. arXiv: 1712.01815.

Guido Novati, Siddhartha Verma, Dmitry Alexeev, Diego Rossinelli, van Wim M. Rees, and Petros Koumoutsakos. Synchronisation through learning for two self-propelled swimmers. *Bioinspiration & Biomimetics*, 12(3):036001, 3 2017. ISSN 1748-3190. doi:10.1088/1748-3190/aa6311. publisher: IOP Publishing.

Siddhartha Verma, Guido Novati, and Petros Koumoutsakos. Efficient collective swimming by harnessing vortices through deep reinforcement learning. *Proceedings of the National Academy of Sciences*, 115(23):5849–5854, 6 2018.

Pingchuan Ma, Yunsheng Tian, Zherong Pan, Bo Ren, and Dinesh Manocha. Fluid directed rigid body control using deep reinforcement learning. *ACM Transactions on Graphics*, 37(4):96:1–96:11, 7 2018. ISSN 0730-0301. doi:10.1145/3197517.3201334.

Xian Yeow Lee, Aditya Balu, Daniel Stoecklein, Baskar Ganapathysubramanian, and Soumik Sarkar. Flow shape design for microfluidic devices using deep reinforcement learning. *arXiv:1811.12444 [cs, stat]*, 11 2018. URL `http://arxiv.org/abs/1811.12444`. arXiv: 1811.12444.

Zhengliang Xiang, Yuequan Bao, Zhiyi Tang, and Hui Li. Deep reinforcement learning-based sampling method for structural reliability assessment. *Reliability Engineering & System Safety*, 199:106901, July 2020. ISSN 0951-8320. doi:10.1016/j.ress.2020.106901.

Shiyin Wei, Yuequan Bao, and Hui Li. Optimal policy for structure maintenance: A deep reinforcement learning framework. *Structural Safety*, 83:101906, March 2020. ISSN 0167-4730. doi:10.1016/j.strusafe.2019.101906.

Jean Rabault, Miroslav Kuchta, Atle Jensen, Ulysse Réglade, and Nicolas Cerardi. Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control. *Journal of Fluid Mechanics*, 865:281–302, 4 2019. ISSN 0022-1120, 1469-7645. doi:10.1017/jfm.2019.62.

Anders Logg, Kent-Andre Mardal, and Garth Wells. *Automated solution of differential equations by the finite element method: The FEniCS book*, volume 84. Springer Science & Business Media, 2012.

Jean Rabault and Alexander Kuhnle. Accelerating deep reinforcement learning strategies of flow control through a multi-environment approach. *Physics of Fluids*, 31(9):094105, 9 2019. ISSN 1070-6631. doi:10.1063/1.5116415. publisher: American Institute of Physics.

Feng Ren, Jean Rabault, and Hui Tang. Applying deep reinforcement learning to active flow control in weakly turbulent conditions. *Physics of Fluids*, 33(3):037121, 3 2021. ISSN 1070-6631. doi:10.1063/5.0037371. publisher: American Institute of Physics.

Romain Paris, Samir Beneddine, and Julien Dandois. Robust flow control and optimal sensor placement using deep reinforcement learning. *Journal of Fluid Mechanics*, 913, 2021.

Hongwei Tang, Jean Rabault, Alexander Kuhnle, Yan Wang, and Tongguang Wang. Robust active flow control over a range of reynolds numbers using an artificial neural network trained through deep reinforcement learning. *Physics of Fluids*, 32(5):053605, 5 2020. ISSN 1070-6631. doi:10.1063/5.0006492. publisher: American Institute of Physics.

Hui Xu, Wei Zhang, Jian Deng, and Jean Rabault. Active flow control with rotating cylinders by an artificial neural network trained by deep reinforcement learning. *Journal of Hydrodynamics*, 32:254–258, 2020. publisher: Springer.

Dixia Fan, Liu Yang, Zhicheng Wang, Michael S. Triantafyllou, and George Em Karniadakis. Reinforcement learning for bluff body active flow control in experiments and simulations. *Proceedings of the National Academy of Sciences*, 117(42):26091–26098, 10 2020. ISSN 0027-8424, 1091-6490. doi:10.1073/pnas.2004939117.

16

Michele Alessandro Bucci, Onofrio Semeraro, Alexandre Allauzen, Guillaume Wisniewski, Laurent Cordier, and Lionel Mathelin. Control of chaotic systems by deep reinforcement learning. *Proceedings of the Royal Society A*, 475(2231):20190351, 2019.

Vincent Belus, Jean Rabault, Jonathan Viquerat, Zhizhao Che, Elie Hachem, and Ulysse Reglade. Exploiting locality and translational invariance to design effective deep reinforcement learning control of the 1-dimensional unstable falling liquid film. *AIP Advances*, 9(12):125014, 2019.

Gerben Beintema, Alessandro Corbetta, Luca Biferale, and Federico Toschi. Controlling rayleigh–bénard convection via reinforcement learning. *Journal of Turbulence*, 21(9-10):585–605, 2020.

Paul Garnier, Jonathan Viquerat, Jean Rabault, Aurélien Larcher, Alexander Kuhnle, and Elie Hachem. A review on deep reinforcement learning for fluid mechanics. *Computers & Fluids*, 225:104973, 7 2021. ISSN 0045-7930. doi:10.1016/j.compfluid.2021.104973.

Philippe Meliga, Edouard Boujo, Gregory Pujals, and François Gallaire. Sensitivity of aerodynamic forces in laminar and turbulent flow past a square cylinder. *Physics of Fluids*, 26:26,104101, 2014. doi:10.1063/1.4896941. publisher: American Institute of Physics.

Jonathan Viquerat, Jean Rabault, Alexander Kuhnle, Hassan Ghraieb, Aurélien Larcher, and Elie Hachem. Direct shape optimization through deep reinforcement learning. *Journal of Computational Physics*, 428:110080, 3 2021. ISSN 0021-9991. doi:10.1016/j.jcp.2020.110080.

Shaopeng Li, Reda Snaiki, and Teng Wu. A knowledge-enhanced deep reinforcement learning-based shape optimizer for aerodynamic mitigation of wind-sensitive structures. *Computer-Aided Civil and Infrastructure Engineering*, 36 (6):733–746, 2021.

Ricardo Vinuesa, Oriol Lehmkuhl, Adrian Lozano-Durán, and Jean Rabault. Flow control in wings and discovery of novel approaches via deep reinforcement learning. *Fluids*, 7(2):62, Feb 2022. ISSN 2311-5521. doi:10.3390/fluids7020062. URL http://dx.doi.org/10.3390/fluids7020062.

Neil Ashton and Vangelis Skaperdas. Verification and Validation of OpenFOAM for High-Lift Aircraft Flows. *Journal of Aircraft*, 56(4):1641–1657, 2019. doi:10.2514/1.C034918.

A. Mack and M. P. N. Spruijt. Validation of OpenFoam for heavy gas dispersion applications. *Journal of Hazardous Materials*, 262:504–516, November 2013. ISSN 0304-3894. doi:10.1016/j.jhazmat.2013.08.065.

E. Robertson, V. Choudhury, S. Bhushan, and D. K. Walters. Validation of OpenFOAM numerical methods and turbulence models for incompressible bluff body flows. *Computers & Fluids*, 123:122–145, December 2015. ISSN 0045-7930. doi:10.1016/j.compfluid.2015.09.010.

Chakrit Suvanjumrat. Implementation and Validation of OpenFOAM for Thermal Convection of Airflow. *Engineering Journal*, 21(5):225–241, September 2017. ISSN 0125-8281. doi:10.4186/ej.2017.21.5.225.

Alexander Kuhnle, Michael Schaarschmidt, and Kai Fricke. Tensorforce: a tensorflow library for applied reinforcement learning, 2017. URL https://github.com/tensorforce/tensorforce. original-date: 2017-03-19T16:24:22Z.

Jiayi Weng, Huayu Chen, Dong Yan, Kaichao You, Alexis Duburcq, Minghao Zhang, Hang Su, and Jun Zhu. Tianshou: A highly modularized deep reinforcement learning library. *arXiv preprint arXiv:2107.14171*, 2021.

Hrvoje Jasak, Aleksandar Jemcov, and Zeljko Tukovic. Openfoam: A c++ library for complex physics simulations. In *International workshop on coupled methods in numerical dynamics*, volume 1000, page 1–20. IUC Dubrovnik Croatia, 2007.

Hrvoje Jasak. OpenFOAM: Open source CFD in research and industry. *International Journal of Naval Architecture and Ocean Engineering*, 1(2):89–94, December 2009. ISSN 2092-6782. doi:10.2478/IJNAOE-2013-0011.

Goong Chen, Qingang Xiong, Philip J. Morris, Eric G. Paterson, Alexey Sergeev, and Y. Wang. OpenFOAM for computational fluid dynamics. *Notices of the AMS*, 61(4):354–363, 2014.

S. L. Brunton and B. R. Noack. Closed-loop turbulence control: Progress and challenges. *Appl. Mech. Rev.*, 67(5): 050801:01–48, 2015.

Kai Arulkumaran, Marc Peter Deisenroth, Miles Brundage, and Anil Anthony Bharath. Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine*, 34(6):26–38, 2017. doi:10.1109/MSP.2017.2743240.

Jean Rabault, Feng Ren, Wei Zhang, Hui Tang, and Hui Xu. Deep reinforcement learning in fluid mechanics: A promising method for both active flow control and shape optimization. *Journal of Hydrodynamics*, 32(2):234–246, 2020.

Vincent François Lavet et al. Deer, 2016. URL https://deer.readthedocs.io/.

Pablo Samuel Castro, Subhodeep Moitra, Carles Gelada, Saurabh Kumar, and Marc G. Bellemare. Dopamine: A research framework for deep reinforcement learning. *arXiv:1812.06110*, 2018. URL `http://arxiv.org/abs/1812.06110`.

Yuandong Tian, Qucheng Gong, Wenling Shang, Yuxin Wu, and C. Lawrence Zitnick. Elf: An extensive, lightweight and flexible research platform for real-time strategy games. In *Advances in neural information processing systems*, page 2656–2666, 2017.

Prafulla Dhariwal, Christopher Hesse, Oleg Klimov, Alex Nichol, Matthias Plappert, Alec Radford, John Schulman, Szymon Sidor, Yuhuai Wu, and Peter Zhokhov. Openai baselines, 2017. URL `https://github.com/openai/baselines`.

Yan Duan, Xi Chen, Rein Houthooft, John Schulman, and Pieter Abbeel. Benchmarking deep reinforcement learning for continuous control. *arXiv:1604.06778 [cs]*, 5 2016. URL `http://arxiv.org/abs/1604.06778`. 01221 arXiv: 1604.06778.

Matthias Plappert. keras-rl, 2016. URL `https://github.com/keras-rl/keras-rl`.

Itai Caspi, Gal Leibovich, Gal Novik, and Shadi Endrawis. Reinforcement learning coach, 12 2017. URL `https://doi.org/10.5281/zenodo.1134899`.

Sergio Guadarrama, Anoop Korattikara, Oscar Ramirez, Pablo Castro, Ethan Holly, Sam Fishman, Ke Wang, Ekaterina Gonina, Neal Wu, Efi Kokiopoulou, Luciano Sbaiz, Jamie Smith, Gábor Bartók, Jesse Berent, Chris Harris, Vincent Vanhoucke, and Eugene Brevdo. Tf-agents: A library for reinforcement learning in tensorflow, 2018. URL `https://github.com/tensorflow/agents`.

Matt Hoffman, Bobak Shahriari, John Aslanides, Gabriel Barth-Maron, Feryal Behbahani, Tamara Norman, Abbas Abdolmaleki, Albin Cassirer, Fan Yang, Kate Baumli, Sarah Henderson, Alex Novikov, Sergio Gómez Colmenarejo, Serkan Cabi, Caglar Gulcehre, Tom Le Paine, Andrew Cowie, Ziyu Wang, Bilal Piot, and Nando de Freitas. Acme: A research framework for distributed reinforcement learning. *arXiv preprint arXiv:2006.00979*, 2020. URL `https://arxiv.org/abs/2006.00979`.

Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv:1606.01540 [cs]*, 6 2016. URL `http://arxiv.org/abs/1606.01540`. arXiv: 1606.01540.

Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL `http://tensorflow.org/`. Software available from tensorflow.org.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL `https://proceedings.neurips.cc/paper/2019/hash/bdbca288fee7f92f2bfa9f7012727740-Abstract.html`. [Online; accessed 2022-04-16].

Michael Schäfer, Stefan Turek, Franz Durst, Egon Krause, and Rolf Rannacher. Benchmark computations of laminar flow around a cylinder. In *Flow simulation with high-performance computers II*, pages 547–566. Springer, 1996.