



City Research Online

City, University of London Institutional Repository

Citation: Cavallazzi, G. M., Guastoni, L., Vinuesa, R. & Pinelli, A. (2024). Deep Reinforcement Learning for the Management of the Wall Regeneration Cycle in Wall-Bounded Turbulent Flows. *Flow, Turbulence and Combustion*, doi: 10.1007/s10494-024-00609-4

This is the published version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/34092/>

Link to published version: <https://doi.org/10.1007/s10494-024-00609-4>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk



Deep Reinforcement Learning for the Management of the Wall Regeneration Cycle in Wall-Bounded Turbulent Flows

Giorgio Maria Cavallazzi¹ · Luca Guastoni^{2,3} · Ricardo Vinuesa³ · Alfredo Pinelli¹

Received: 12 August 2024 / Accepted: 29 October 2024
© The Author(s) 2024

Abstract

The *wall cycle* in wall-bounded turbulent flows is a complex turbulence regeneration mechanism that remains not fully understood. This study explores the potential of deep reinforcement learning (DRL) for managing the wall regeneration cycle to achieve desired flow dynamics. To create a robust framework for DRL-based flow control, we have integrated the *StableBaselines3* DRL libraries with the open-source direct numerical simulation (DNS) solver *CaNS*. The DRL agent interacts with the DNS environment, learning policies that modify wall boundary conditions to optimise objectives such as the reduction of the skin-friction coefficient or the enhancement of certain coherent structures' features. The implementation makes use of the message-passing-interface (MPI) wrappers for efficient communication between the Python-based DRL agent and the DNS solver, ensuring scalability on high-performance computing architectures. Initial experiments demonstrate the capability of DRL to achieve drag reduction rates comparable with those achieved via traditional methods, although limited to short time intervals. We also propose a strategy to enhance the coherence of velocity streaks, assuming that maintaining straight streaks can inhibit instability and further reduce skin-friction. Our results highlight the promise of DRL in flow-control applications and underscore the need for more advanced control laws and objective functions. Future work will focus on optimising actuation intervals and exploring new computational architectures to extend the applicability and the efficiency of DRL in turbulent flow management.

Keywords Flow control · Drag reduction · Direct numerical simulation · Deep reinforcement learning

1 Introduction

The wall cycle (Jiménez and Pinelli 1999) in wall-bounded turbulent flows is a turbulence regenerating mechanism whose dynamics are still not fully understood. Understanding and controlling the wall-cycle is a technological challenge primarily pursued to either reduce

Luca Guastoni, Ricardo Vinuesa, and Alfredo Pinelli have contributed equally to this work.

Extended author information available on the last page of the article

the skin-friction coefficient or to achieve a desired impact on flow dynamics (e.g. enhancement of heat and mass transfer).

Active flow control techniques are among the most promising methodologies to manipulate the near wall flow for technological benefits because of their ability to work with closed-loop feedback and handle off-design conditions. However, their complexity currently limits real-world application. One such technique is based on the use of Streamwise Travelling Waves (STW) of spanwise velocity (Quadrio et al. 2009), which involves applying a sinusoidal wave with a time-dependent pulsation as a boundary condition for the spanwise velocity w at the wall. The use of this method can lead to skin-friction coefficient C_f reduction by up to 45% (Quadrio et al. 2009). However, the definition of a particular STW leading to the highest drag reduction requires a computationally expensive analysis of the parametric space at hand, defined by the frequency and spatial wavenumber of the travelling wave

Deep Reinforcement Learning (LeCun et al. 2015; Li 2018) offers a potential breakthrough in fluid mechanics applications, although its use in this area is still in its early stages compared to traditional supervised learning approaches (Garnier et al. 2021; Rabault et al. 2019). Two primary categories of research contributions have emerged in this field: the first one focuses on using DRL for turbulence modelling (Novati et al. 2021), while the second applies DRL to active flow control. Our present work falls into the latter category, aiming to explore how DRL can be used to control complex flow systems dynamically and efficiently (Guastoni et al. 2023).

A few literature contributions looked at the DRL control of the movement of an object within a fluid flow. For instance, some studies represent a fish swimming in turbulent flow or within a school of fish, where the reward function is designed to maximise the swimming efficiency (Biferale et al. 2019; Verma et al. 2018). Another group of studies leverages DRL to control the dynamics of the fluid itself. One prominent example of this is found in Rabault et al. (2019), where the flow around a cylinder is controlled using jets perpendicular to the main flow direction. This setup has since been adopted as a benchmark in numerous studies (Rabault and Kuhnle 2019; Paris et al. 2021; Tang et al. 2020; Fan et al. 2020; Ren et al. 2021; Xu et al. 2020), highlighting both the growing interest in DRL for active flow control and the importance of developing benchmark cases to serve as a foundation for future research.

In this study, we use DRL as an active flow control technique for wall-bounded turbulent flows, a significant increase in complexity for two key reasons. Firstly, the two-dimensional (2D) cylinder case presents a well-defined vortex shedding pattern at a specific frequency, which dominates the flow. Suppressing this frequency offers a direct venue for controlling the wake. However, at higher Reynolds numbers, a different strategy is employed, where the agent energises the boundary layer on the cylinder's surface to induce drag crisis (Varela et al. 2022). In contrast, wall-bounded flows, such as those examined here, are inherently multi-scale phenomena, with no clear dominant frequencies or mechanisms to target for drag reduction. Indeed, turbulent channel flows represent an inherently three-dimensional (3D) situation populated by turbulent structures, featuring a nonlinear, chaotic, and challenging benchmark for testing DRL algorithms and control strategies.

Recent efforts have begun to investigate similar flow cases, such as in Sonoda et al. (2023), where DRL was applied to a standard channel flow, and in Zeng et al. (2022), where a Couette flow was controlled using two streamwise parallel slots. In the latter case, the DRL agent was trained on a reduced-order model before being applied to the full problem. Another example, closely related to the present contribution, the work of Guastoni et al. (2023) and Wälchli et al. (2024), where the traditional opposition control method (i.e.

blowing and suction from the wall) was improved by using a DRL algorithm, specifically deep deterministic policy gradient (DDPG). The findings demonstrated that DRL could achieve up to 30% drag reduction, surpassing the performance of the classical opposition control technique (Choi et al. 1994).

Finally, it is worth mentioning that DRL has been applied to a variety of other control tasks, spanning simple one-dimensional falling-fluid instabilities (Belus et al. 2019; Vignon et al. 2023), convection problems (Beintema et al. 2020), chaotic turbulent combustion systems (Bucci et al. 2019), and a range of engineering applications (Henry de Frahan et al. 2022; Zheng et al. 2021; Vinuesa et al. 2022).

In fluid dynamics-related contexts (Vignon et al. 2023), DRL typically involves the interaction of an agent with a physical environment, as illustrated in Fig. 1. The agent, implemented as a neural network, begins with no prior knowledge of the underlying physical phenomena. Using on-policy algorithms like Proximal Policy Optimization (PPO) (Schulman et al. 2017), the agent receives an observation—a partial and potentially noisy representation of the environment’s full state. The state represents the complete set of information describing the environment, but is often not fully accessible to the agent. The observation, which is a subset of this state, is used as input to the policy. The agent then receives a reward, a scalar value that indicates how favourable the observation is towards achieving the desired outcome. Based on the observation and the reward, the agent selects an action (or set of actions), which modifies the environment, leading to a new state in the next step.

When the environment (e.g. the DNS of a turbulent flow) receives a new set of actions, the physics of the flow are simulated to produce a new observation (e.g. a sample of the velocity field) and its associated reward. This observation updates the environment’s state, and the agent uses the collected experience to adjust the neural network’s weights based on the current reward and the chosen policy. These updated weights improve the agent’s decision-making process for future actions. The agent’s new actions then interact with the environment, leading to another state transition, which triggers the next step. Multiple steps form an episode, and the agent continuously collects experiences over many episodes, using an explicit optimisation step to update the policy and maximise cumulative rewards.

In this work, an agent is trained using the PPO algorithm (Schulman et al. 2017) to gather observations from a Direct Numerical Simulation (DNS) (*the environment*) of a plane, turbulent channel flow and apply a time-dependent *action* that dynamically modifies the wall boundary conditions, ultimately pursuing a *policy* that maximises a given the reward.

We will be looking at two different rewards: one is related to the minimisation of C_f , the other to the promotion of highly elongated velocity streaks.

This second choice is motivated by the intention of understanding how the topology of coherent structures impacts the development of the wall regeneration cycle. This involves promoting specific topologies of the flow, which helps infer the causality between the control law and its effect on the turbulent wall cycle by linking the straightening of streaks to the reduction of C_f . Further efforts will focus on this direction by testing alternative objective functions that target specific features of the wall cycle.

From an implementation perspective, this work involved integrating the *StableBases3* DRL libraries (Raffin et al. 2021) with the open-source DNS solver *CaNS* (Costa 2018). The final interface is based on an ad hoc MPI implementation, where a Python code invokes the execution of the DNS code on a set of processors and manages the communication with the DRL library. While the most common choice is resorting to Python codes both for the DNS and the DRL routines [see for example Vignon et al. (2023)], similar

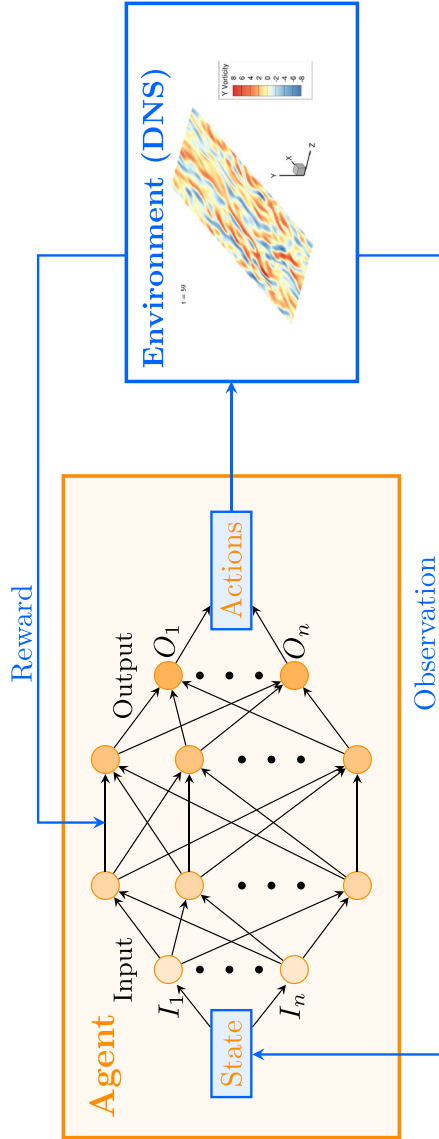


Fig. 1 Schematic representation of the interaction between the DRL agent and the DNS environment

unconventional couplings become necessary when the 3D turbulent DNS requires more tailored codes not to waste computational resources, like in the work of Guastoni et al. (2023), Suárez et al. (2024) and Font et al. (2024). Alternatives that are simpler to implement have been explored, such as DNS codes being asked to save data on disk, rather than keeping it in memory, with bash scripts orchestrating the various calls of the DRL training and the DNS code to produce new data. This method has been initially implemented with a reliable DNS in-house developed solver, SUSA (Monti et al. 2019, 2020), but it has later been discontinued, due to the high inefficiency of the coupling itself.

The initial results indicate the need for more advanced control laws to fully exploit the tools developed in these early experiments. Currently, DRL agents are trained to find a control law based on a single dynamic parameter, but better policies could emerge from more complex control laws involving multiple dynamic parameters.

2 Methodology

2.1 Deep Reinforcement Learning

As mentioned, in the present context, DRL-based flow control consists in the interaction of an agent with a turbulent channel flow modeled through a DNS, as depicted in Fig. 1.

The agent aims to learn a policy $\pi_\theta(a|s)$ that represents the probability of taking action a in state s , given neural network parameters θ . The interaction with the environment generates a trajectory $\{(s_t, a_t, r_t, s_{t+1})\}$ in the state-action space based on the current policy. The agent then evaluates the advantage function A_t , which measures how much better a particular action is compared to the average action taken in a given state. A_t is typically defined as:

$$A_t = Q(s_t, a_t) - V(s_t) \quad (1)$$

where $Q(s_t, a_t)$ is the action-value function, and $V(s_t)$ is the value function.

However, in Proximal Policy Optimization (PPO), the agent does not have access to the true values of $Q(s_t, a_t)$ or $V(s_t)$. Instead, PPO approximates the advantage function using the critic network (value function network) and the sampled returns from the environment. Specifically, the advantage function is approximated by:

$$A_t \approx R_t - V(s_t) \quad (2)$$

where R_t is the sampled return, and $V(s_t)$ is the value function estimated by the critic.

The policy parameters θ are then updated to minimise the following loss function:

$$L(\theta) = \mathbb{E}_t [L^{CLIP}(\theta) - c_1 L^{VF}(\theta) + c_2 S[\pi_\theta](s_t)], \quad (3)$$

where

$$L^{CLIP}(\theta) = \mathbb{E}_t [\min(r_t(\theta)A_t, \text{clip}(p_t(\theta), 1 - \epsilon, 1 + \epsilon)A_t)], \quad (4)$$

and

$$L^{VF}(\theta) = \mathbb{E}_t [(V_\theta(s_t) - V_t^{\text{target}})^2]. \quad (5)$$

Here, $p_t(\theta) = \pi_\theta(a_t|s_t)/\pi_{\theta_{old}}(a_t|s_t)$ is the probability ratio, and ϵ is a hyperparameter controlling the clipping range. V_t^{target} is typically computed as the discounted sum of rewards plus the bootstrapped value of the next state. The terms c_1 and c_2 are coefficients that balance the contributions of the value function loss and the entropy bonus, respectively. The entropy bonus is defined as:

$$S[\pi](s) = - \sum_a \pi(a|s) \log \pi(a|s), \quad (6)$$

encouraging exploration by discouraging premature convergence of the policy.

The policy outputs a probability distribution over possible actions, and during training, an action a is sampled from this distribution based on the predicted probabilities.

The input data from DNS can have any structure, as long as the neural network (NN) is built with matrices of compatible sizes for multiplication. Inputs can be array-like, and in such cases, the most common policy adopted is the Multi-Input Policy (MLP) (Lillicrap et al. 2019), which learns to weigh each input based on its relevance to the task, but it does not explicitly model any interdependencies between input features. However, in our specific application, it is important to preserve the spatial correlation of the input, as any combinations of sampled instantaneous quantities from a flow field are linked by the Navier–Stokes equations. Therefore, 2D, black-and-white, images rather than one-dimensional arrays of re-arranged data are fed to the neural network. This data consists of wall-parallel slices of streamwise velocity values (96×64) sampled at a given y^+ (i.e. $y^+ = 20$). When using an image as input, the colour of a given pixel is not necessarily uncorrelated with the colours of other pixels.

By using images, we can effectively exploit the features of convolutional neural networks (CNNs) (LeCun et al. 1989, 1998) without losing information. These policies could show a computational advantage to more standard fully connected layers, reducing computational requirements that otherwise could become prohibitive, and they are effective in fluid dynamics for inspecting turbulent fields as images. This approach is illustrated in the work of Guastoni et al. (2021), where DRL with CNNs was applied to another flow control technique, opposition control, achieving notable results in terms of drag reduction. A more schematic but detailed representation of how the training and the interaction with the DNS environment take place can be found in Algorithm 1.

The custom CNN built for this purpose and represented in Fig. 2 begins by applying a series of convolutional layers, which progressively learn to extract increasingly complex features from the input. In the initial stage, the network employs a convolutional layer with 48 filters, each using a 3×3 kernel. This layer scans the input, detecting basic patterns like edges and textures. After each convolution, a batch normalisation step stabilises the training process by normalising the output, and a ReLU (Rectified Linear Unit) activation function introduces non-linearity, enabling the network to learn more complex patterns.

As the data progresses through the network, it encounters additional convolutional layers, first with 64 filters and then with 96 filters, both using the same 3×3 kernel size. More layers will lead to more complete but simplified representation of turbulent snapshots corresponding to the observed state. Each convolutional layer is again followed by batch normalisation and a ReLU activation to ensure that the learning process remains robust and capable of capturing complex features.

Following the convolutional layers, the network transitions to fully connected layers. First, the high-dimensional output from the convolutional layers is flattened into a one-dimensional vector. This vector, now representing a combination of all the features

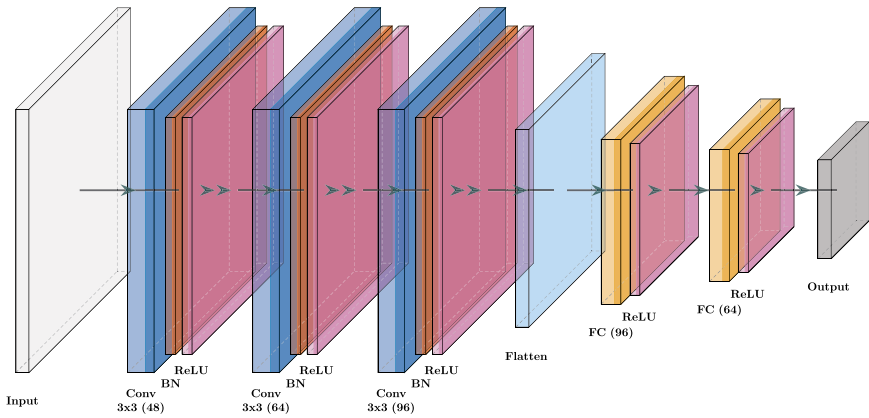


Fig. 2 Schematic representation of the CNN neural network. The dimensions are indicated in brackets, “BN” stands for batch normalisation and “FC” for fully connected layers

extracted, is passed through a series of fully connected layers. The first of these layers reduces the vector to 96 units, and the second further reduces it to 64 units, both using ReLU activations to retain the network’s ability to model complex relationships within the data.

Finally, the output of the last fully connected layer is passed through another linear layer, which reduces the feature vector to the desired output size specified by the user. The dimension of this output, which subsequently becomes the input for both the actor and critic networks, has been determined through empirical tests within the framework of this work and is the outcome of a trade-off argument. A value that is too small may fail to adequately represent complex data structures, while a larger dimension, although offering a better representation of turbulent fields, increases memory usage and extends learning time. This can slow down training to the extent that useful policies may not be learned within the given network architecture.

As far as the critic/value network is concerned, two hidden layers of 512 neurons each are used for both. Smaller sizes have been tried (256, 384) but they struggled to deal with the dimensionality of the input or with the number of features to be extracted (always 16).

2.2 DNS Set Up

As already mentioned, our *environment* is the DNS of a turbulent plane channel flow with modified wall boundary conditions. Thus, we solve the full 3D, unsteady Navier–Stokes (NS) equations, in their incompressible formulation:

$$\frac{\partial u_i}{\partial t} + u_j \frac{\partial u_i}{\partial x_j} = -\frac{1}{\rho} \frac{\partial p}{\partial x_i} + \nu \frac{\partial^2 u_i}{\partial x_j \partial x_j} \tag{7}$$

$$\frac{\partial u_i}{\partial x_i} = 0, \tag{8}$$

where the Einstein index notation, with summation over repeated indices, is used to represent the differential operators, $x_{\{i,j,k\}}$ with $i = [1, 2, 3]$ corresponds to the i^{th} direction in a three-dimensional space, u_i is the corresponding velocity component, p is the pressure, ρ the density and ν the kinematic viscosity of the flow.

To advance the equations in time we use two different solvers, *CaNS* (Costa 2018) and *SUSA* (Monti et al. 2019, 2020). They are both efficient solver for massively parallel direct numerical simulations of incompressible turbulent flows.

CaNS is a freely available code (Costa 2023) designed for simulating parallel turbulent flows. The code uses a second-order, staggered Cartesian finite grid pressure-correction scheme (Chorin 1968), advanced in time through a three-stage, low-memory Runge–Kutta scheme (Rai and Moin 1991). Additionally, the code offers an optional implicit treatment of viscous terms in the wall-normal direction. The pressure Poisson equation, solved at each Runge–Kutta stage, is handled using eigenfunction expansions, allowing the use of highly efficient FFT-based solvers (Cooley and Tukey 1965) for problems with various combinations of homogeneous pressure boundary conditions. *CaNS* incorporates a 2D pencil-like domain decomposition, enabling efficient massively parallel simulations with excellent strong-scaling performance across thousands of cores. Extensive validation cases are available in the code repository (Costa 2023), including turbulent Poiseuille flow, turbulent Couette flow, and turbulent square duct flow. The code supports Dirichlet or Neumann boundary conditions in one direction, periodic conditions in another, and a combination of Dirichlet, Neumann, or periodic conditions in the third direction.

SUSA is a second-order finite volume code (in both time and space) developed at City St George's, University of London. It has been extensively validated for both DNS and LES versions in a wide range of turbulent flows, with a full list available in Monti et al. (2020) and the references therein. The code supports a periodic direction (spanwise), while in the wall-normal plane, it adopts a body-fitted approach compatible with Cartesian, C, and O grids. *SUSA* uses a discretisation of the momentum and mass conservation equations on a cell-centered, co-located grid. Fluxes are approximated using a second-order central formulation, and the Rhie and Chow method (1983) is employed to prevent spurious pressure oscillations. The equations are advanced in time through a second-order semi-implicit fractional-step procedure, where the implicit Crank–Nicolson scheme is applied for wall-normal diffusive terms, and the explicit Adams–Bashforth scheme is used for all other terms. Further details can be found in Monti et al. (2020) and the references therein.

The baseline codes have been modified to account for time-dependent boundary conditions at the walls. Specifically, we have implemented modified conditions on the spanwise velocity components w to introduce the effect of the streamwise travelling waves at the two walls of a fully turbulent channel flow. The boundary condition on w reads as:

$$w(x, 0, z, t) = A \sin(\kappa_x x - \omega t), \quad (9)$$

where ω is a pulsation, A is an amplitude, and κ_x is a spatial wavenumber. The same conditions are applied at $y = 0$ and $y = 2H$, where H is the semi-height of the channel, as shown in Fig. 3.

All simulations have been performed at the same $Re_b = U_b h / \nu = 3170$ (where U_b is the bulk velocity, h is the semi-height of the channel and ν is the kinematic viscosity) which corresponds to $Re_\tau = u_\tau h / \nu = 200$ (where $u_\tau = \sqrt{\tau_w / \rho}$ is the friction velocity, $\tau_w = \mu(dU/dz)|_{z=0}$ being the wall-shear stress), using the same domain size (streamwise dimension $L_x/h = \pi$, spanwise dimension $L_z/h = \pi/3$), and the same number of grid nodes ($96 \times 100 \times 64$, along the streamwise x direction, the wall-normal y direction

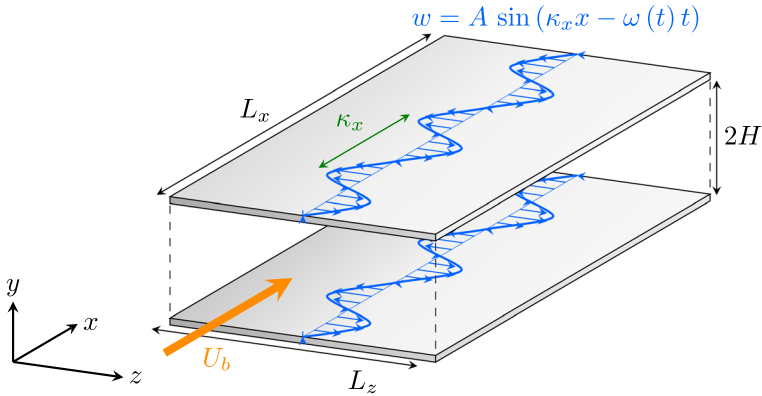


Fig. 3 Channel flow DNS setup. The STW boundary condition is applied equally at the two walls. Now ω is a function of t , so the phase-speed is not constant any more. This is still a class of Dirichlet BCs, that along with homogeneous Dirichlet for u and v and no-slip conditions are applied at the two walls. Transparent sides in this figure correspond to periodic boundary condition, in the streamwise and spanwise directions

and the spanwise z direction, respectively). Before considering DRL-actuated cases, both codes have been validated, using the aforementioned parameters, for both non-actuated channel flows and for an actuated STW case (Quadrio et al. 2009) with $\kappa_x^+ = 0.0063$ and $\omega^+ = 0.0195$ (Gatti and Quadrio 2016) (units with a + superscript are normalised with ν and u_τ). In Figs. 4 and 5, we report the spanwise velocity and the diagonal Reynolds stresses distribution for the STW actuated channel in the case of a drag reducing setup. The obtained results are in good agreement with the ones in Quadrio et al. (2009) (see figures 9 and 10 within the given reference). In the given figures, the averages are obtained in phase with the wall motion, i.e. computed accumulating data at $\xi = x - \omega/\kappa_x t$. Those averages, denoted by $\langle \cdot \rangle$, are computed using a 2D average over the spanwise direction and time. In particular, statistics have been accumulated over ≈ 150 periods of the STW after the initial numerical transient.

As far as all the DRL-controlled cases reported, the nominal value of $Re_\tau = 200$ (computed in the corresponding non-actuated case) results in a domain size, expressed in viscous units, of $L_x^+ = 628$, $L_y^+ = 200$, and $L_z^+ = 314$. These values were selected for several reasons: they are slightly larger than the minimum requirements for a self-sustained *minimal flow unit* channel flow (Jiménez and Moin 1991); they provide a numerical domain that accommodates at most two pairs of velocity streaks, simplifying the analysis of the effects of actuations on coherent structures; and they allow for fast evaluation of the environment step without imposing significant memory usage requirements. Regarding the grid size, Δx^+ and Δz^+ are well within the commonly recommended values (Moser et al. 1999), and the distance of the first internal grid point from the wall is consistently below $y^+ = 1$ for all considered cases.

2.3 Actuation and Objective Functions

The connection with DRL is embedded in ω , which becomes a time-dependent parameter (i.e., *the actuation*), whose behaviour now depends on time and varies according to the estimations of the *agent* to achieve a better reward. Most of the work carried out so far has

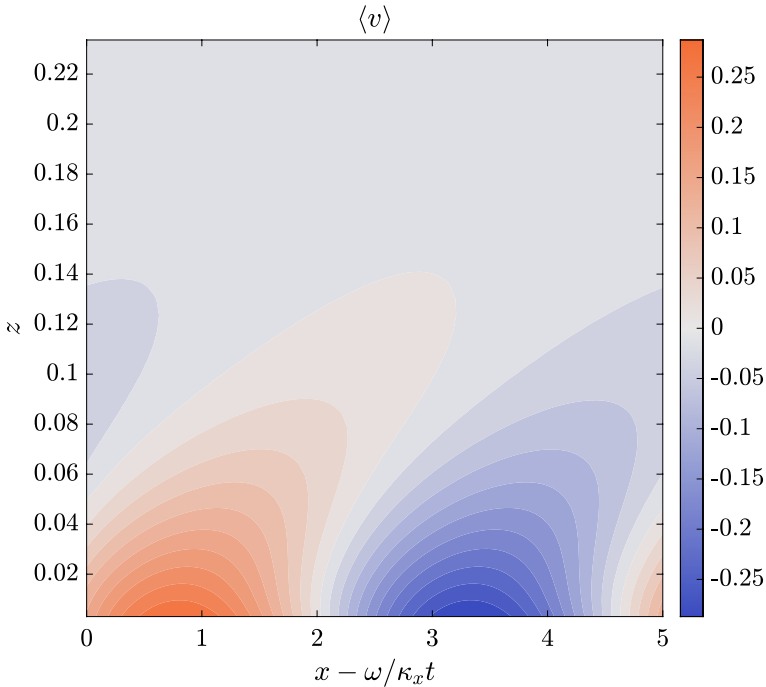


Fig. 4 Average of the spanwise velocity in phase with the wave in a region close to the wall. The values of $\langle v \rangle$ are reported in the colorbar. The vertical represents the wall-normal coordinate in outer units; the horizontal axis corresponds to the convective variable ξ in phase with the wave. The positive peak and its negative counterpart are well represented at the wall

focused on using DRL to find a time-law $\omega(t)$ to maximise a chosen objective function during an *episode*.

It is noted that using $\omega(t)$ as the actuation or the actual phase $\phi(t) = \omega(t) \cdot t$ in Eq. 9 represents an equivalent choice, as both approaches describe the same oscillatory behaviour.

To justify the choice of ω as an *actuator*, it is necessary to provide some background on the physical modifications that STWs introduce in the near-wall region. The following conceptual description closely follows the discussion offered by Quadrio et al. (2009). The mechanism induced by STWs involves two main effects: phase locking between the traveling waves and near-wall turbulent structures, and the interaction of the turbulent flow with the Stokes layer generated by the oscillatory, time-dependent condition on the spanwise velocity at the wall.

When the phase speed of the traveling waves, $c = \omega/k_x$, approaches the *convective velocity* of the near-wall turbulent structures, U_w (where $U_w \simeq 0.5U_b$, Kim and Hussain 1993), the near-wall flow structures synchronise with the wall waves. For STW with a fixed ω , as in Quadrio et al. (2009), this phase-locking mechanism leads to a maximum drag increase when $c = U_w \simeq 0.5U_b$. In our case, by allowing ω to vary in time, it is possible to adapt to the local convection speed of the structures, since $U_w = 0.5U_b$ is an average value representative of various local convection speeds.

Concerning the second mechanism, the traveling wave condition can be compared to the oscillating wall mechanism (i.e., the STW limit for $k_x = 0$). In the oscillating wall case,

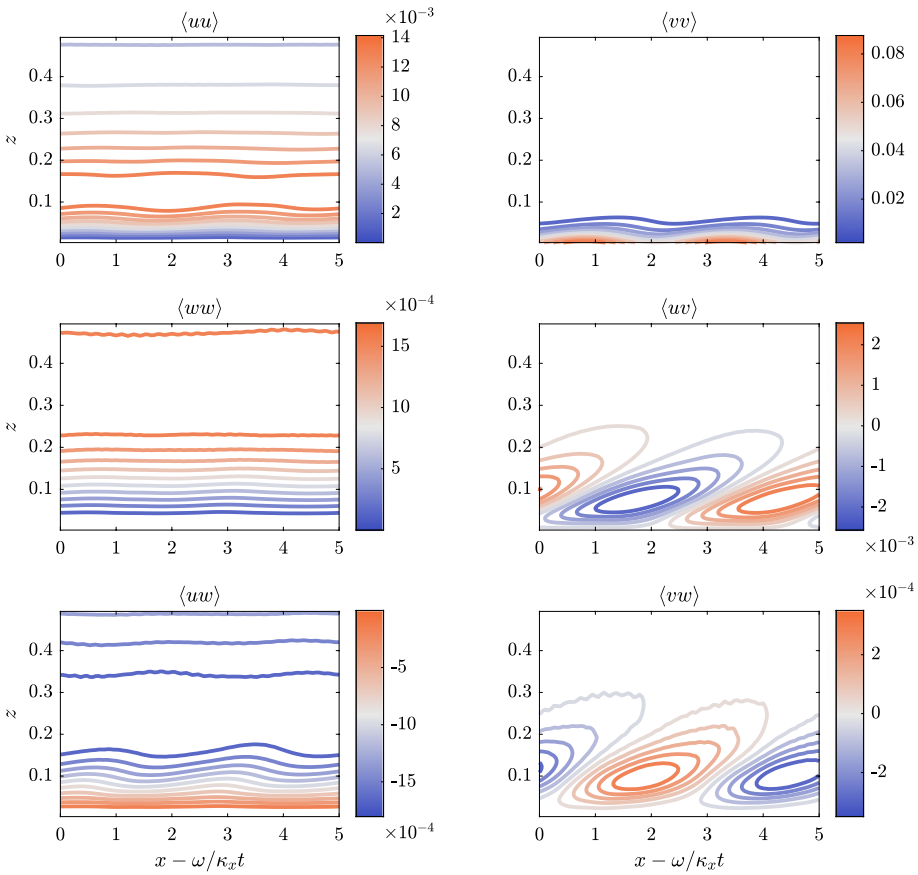


Fig. 5 Reynolds stresses $\langle uu \rangle$, $\langle vv \rangle$, $\langle ww \rangle$, $\langle uv \rangle$, $\langle uw \rangle$ and $\langle vw \rangle$ in a region close to the wall. Values are in outer units to allow for a direct comparison with figure 10 of Quadrio et al. (2009). The axes share the values of Fig. 4, with different colorbars for each quantity. Reynolds stresses involving spanwise velocity fluctuations feature patterns correlated with the wall motion; the others present a thick layer with small, but still visible, footprints of the STW

for small Stokes layer thicknesses (i.e., high frequencies), only a modest effect is exerted on the structures. A negligible effect also occurs for extremely low frequencies when the Stokes layer thickness exceeds the buffer layer size. Between these two extremes, there exists an optimal frequency ω_{opt} that maximises drag reduction by weakening the regeneration cycle, thus reducing turbulence intensities. This optimal value of ω (ω_{opt}), is related to the typical lifetime of coherent structures in the buffer layer. Similarly, for traveling waves, the flow embeds a Stokes layer, and its thickness again plays a fundamental role in determining turbulence suppression and drag reduction.

The important difference for STWs is that the Stokes layer thickness depends not only on ω but also on the wavelength k_x . The maximum interaction, resulting in turbulence suppression and drag reduction, occurs when $|\mathcal{L}_v k_x - \omega|$ approaches ω_{opt} . In this case, ω_{opt} is affected by how the traveling wave modifies the expected lifetime of the coherent structures (e.g., structures phase-locked with the traveling wave tend to survive longer). In the present case, when considering a minimal flow unit, it is expected that by varying ω as a function

of time, the Stokes layer thickness can be optimised (*modulated*) based on the actual flow configuration.

Another parameter that needs to be considered is the duration of an episode. To this end, one should consider the physical phenomenon at hand. In our particular case, it is known that STWs can affect the presence/position of the low-speed streaks in the logarithmic region close to the wall. The streaks are characterised by a life cycle of $\Delta t_{\text{cycle}}^+ = 80 - 100$ (see Jiménez and Pinelli 1999), thus an episode length $t_{\text{ep}}^+ = 160$ is suitable to capture, with some margin of uncertainty, at least one full regeneration cycle. We have also considered some cases with $t_{\text{ep}}^+ = 80$, finding that this time interval is not always sufficient to control a full cycle, and other cases with a larger time interval, i.e., $t_{\text{ep}}^+ = 320$ that did not present any significant variation in the results.

Within the duration of each episode, ω is set to change every $\Delta\tau$, assigning a smoothed piece-wise constant behaviour to avoid sharp changes in the boundary values across the actuations. Note that $\Delta\tau$ is the actuation period (also called a *step*), and its duration was initially set to 10 viscous time units to ensure a sufficient number of changes during an episode while not excessively interfering with the actuations. It was found that a smaller $\Delta\tau$ had less impact on the outer flow and led to unwanted oscillations in the flow statistics.

As previously discussed, the DRL agent needs to be trained to learn how to choose the next action and achieve an effective policy to optimise the reward r . The initial and most straightforward reward r for this problem is based on C_f :

$$r = \frac{C_{f,\text{max}} - C_f}{C_{f,\text{max}} - C_{f,\text{min}}} \quad (10)$$

In the above expression, $C_{f,\text{max}}$ and $C_{f,\text{min}}$ are, respectively, the maximum C_f that is considered (the one corresponding to an unactuated turbulent channel flow with no control at the wall) and the minimum one that is aimed for, corresponding to 65% of DR, an optimistic goal based on results achieved with some other DR techniques coupled with DRL, such as by Guastoni et al. (2023). This initial training with the C_f -based reward was carried out with the DNS-DRL coupling SUSA-Tensorforce via bash scripts (see details in Sect. 2.4). Another characteristic aspect of this initial study is that the observed space consisted of a 2D wall-parallel field of the streamwise velocity sampled at $y^+ = 1$. The rationale behind this choice was to provide data to the agent that was strongly correlated to the quantity that defined the reward, given that C_f is proportional to dU/dy at the wall. For the next studies, this value of y^+ was changed to acquire snapshots containing more information of the whole state of the DNS simulation, relaxing the constraint of the relation with the reward.

These values can be tuned to assign a high reward to realistically achievable flow configurations, but the tuning procedure is not trivial: if $C_{f,\text{min}}$ is too low, the network will not learn because all the policies lead to a low reward and are discarded, while if $C_{f,\text{min}}$ is too high, the network will learn how to reach a configuration that does not significantly reduce the drag. With this formulation, a small C_f implies that some drag reduction (DR) is obtained and that an associated temporal law for $\omega(t)$ has been achieved.

Alongside Eq. 10, we have considered a second type of reward that tackles the actual topology of the coherent flow structures populating the near-wall region. This choice is motivated by the aim of discovering the causality between the emergence and decay of coherent structures. We targeted a measure of the coherence of the velocity streaks as the cost function. In particular, following Doohan et al. (2021), we used DRL to maximise the energy content of the streamwise velocity fluctuations within a box localised on the walls.

This objective should be seen as an attempt to produce streaks that are more *rectilinear* and resilient to instabilities. In particular, the reward has been designed to maximise the ratio $\tilde{E}_{ks}/\tilde{E}_k$, where

$$\tilde{E}_{ks} = \int_0^{L_x} \int_0^{\delta_y} \int_0^{L_z} u'^2 dx dy dz. \quad (11)$$

$$\tilde{E}_k = \int_0^{L_x} \int_0^{\delta_y} \int_0^{L_z} (u'^2 + v'^2 + w'^2) dx dy dz. \quad (12)$$

In Eqs. 11 and 12, u' , v' and w' are the fluctuation of the streamwise, wall-normal and spanwise velocity respectively, δ_y is chosen to define a box whose height in friction units is set to $\delta_y^+ = 60$, a size that should be sufficient to host the wall velocity streaks, while not extending too much beyond the buffer zone.

Maximising the ratio

$$r = \tilde{E}_{ks}/\tilde{E}_k \quad (13)$$

should promote the longitudinal straightness of the streaks by favouring fluctuations in the streamwise direction. The rationale behind this experiment is related to the minimisation of the sinusoidal instability (Kawahara et al. 2003) of the streaks and its possible impact on the wall regeneration cycle.

One final matter to be discussed concerns the size of the DNS computational box. A large box would be computationally challenging for the training of the DRL network since it would involve a much larger dataset featuring complex and non-local flow interactions. This would lead to a long time required for completing the training. On the other hand, a box that is too small could lead to re-laminarisation of the flow for sizes below the minimal flow unit threshold, as indicated by Jiménez and Moin in their 1991 seminal paper (Jiménez and Moin 1991).

For this study, a box size slightly larger than the minimal flow unit (Jiménez and Moin 1991) was chosen. This choice is based on the hypothesis that a pair of velocity streaks contains all the required dynamics to sustain the wall regeneration cycle of turbulence, and that the DRL policy will be tailored to this condition.

2.4 DRL-DNS Interface

Coupling the DNS solver with the DRL interface occupied most of the time needed to carry out this work. Initially, a finite volume method solver previously mentioned, SUSAN, has been re-designed to interact with bash scripts that commanded when to write the files needed for the training, and to re-initialise the grid with the previous flow field with a new actuation parameter (if the new actuation belongs to the same episode). This is the setup that has been used, along with the Tensorforce libraries previously used by Vignon et al. (2023), for the cases where the reward was solely based on C_f , and the numerical setup for the episodes and the training is shown in Table 1 in the left column.

The DRL portion of the code of the cases where the streaks topology is discussed relies on the use of *StableBaselines3* library, an open source implementation of DRL based on *PyTorch* (Paszke et al. 2019). Despite not being optimal from a computational point of view, the use of Python is still the best option for a few reasons.

Table 1 Comparison of parameters used to train the network for reducing C_f in the SUSA simulation and promoting straighter streak topologies in the CaNS simulation

Parameter	SUSA	CaNS
Incompressible DNS solver	SUSA	CaNS
Simulation time t^+	80	160
Actuation time $\Delta\tau^+$	10	10
Cores per simulation	32 + 1	32 + 1
y^+ for collecting data	1	15
DRL agent policy	PPO	PPO
DRL library	Tensorforce	StableBaselines3
Python-Fortran interface	bash	MPI (mpi4py)
Input policy	Multi Input Policy	CNN
Actor/critic architecture	[512, 512], [512, 512]	[512, 512], [512, 512]

- *Strong community support* DRL is a machine learning methodology used in many different fields by both practitioners and developers. This vast base of users has produced many plug-and-play libraries available for public use, allowing, for example, engineers to interface their own software with sophisticated DRL libraries almost effortlessly. This enables researchers to focus on the aspects of the work closer to their area of expertise, easing and speeding up their research.
- *Cross-platform libraries* The existence of Python environments that are easy to set up on different machines allows users to avoid spending too much time compiling/installing libraries.
- *Libraries with Fortran/C++ bindings* The speed of compiled languages is still much higher than that of interpreted languages like Python. However, some Python-based packages such as NumPy (Harris et al. 2020) embed popular, highly efficient, pre-compiled libraries, such as the popular C++ *BLAS* and *LAPACK* libraries. Others are just C++ wrappers, such as *mpi4py* (Dalcin and Fang 2021).

The overall software platform developed for the present work can be separated into two parts: the first part contains the external calls that define the duration of the network training and the settings for the agent; the second part concerns the environment, consisting of a class made up of several functions that are called during an actuation to obtain new observations and rewards, and to simulate the fluid flow.

The training is managed using a script that allows the user to change any settings. Additionally, a feature has been added that allows building a CNN from scratch, avoiding reliance on automated algorithms to set matrix sizes based on the input dimensions and the number of actions. In the current version, the customisation of the CNN parameters has been limited to the basic requirements needed to carry out a training session. In the future, other features that will add further flexibility will be implemented.

The development of a specialised interface is essential to balance computational efficiency in DNS simulations with the capabilities of state-of-the-art, open-source machine learning libraries. It is particularly noted that:

- Compiled languages are more efficient in terms of speed and memory requirements than interpreted languages such as Python.

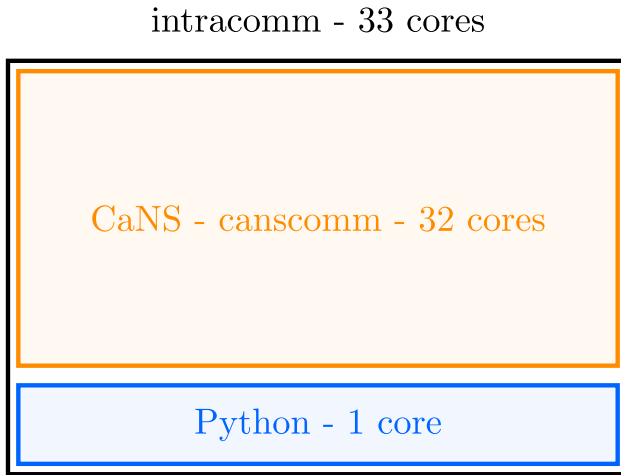


Fig. 6 Schematic of the MPI communicators for the DRL agent and CaNS with $N = 32$. The autonomy of the CaNS environment is preserved with the existence of the *canscomm* communicator, while *intracomm* guarantees freedom of communication between the DNS ranks and the DRL one

- Migrating a CFD code from one language to another, including the use of MPI-based calls, is a task that would require months, if not years. The effort would include programming, debugging, performance tuning, and extensive validation campaigns.
- A researcher with a background in turbulence and fluid dynamics may not have a sufficient skill set to implement efficient, state-of-the-art machine learning routines in a compiled language.

These reasons constitute the rationale behind the development of our DNS-DRL software platform. In particular, we have followed a simple approach for linking a Fortran/C++ DNS code to a Python DRL library. The basic idea concerns the execution of bash commands inside Python to launch the DNS environment for data production. These data are used by the DRL agent that will update the weights of the neural network defining the policy. This approach is robust and easy to implement, but it is case-specific and cannot be generalised. For example, if the amount of data that the environment has to provide to the agent is too big, the time for the Fortran I/O interface to write the file on the disk and for the Python I/O interface to read it may become not negligible, and even saturate a node. Moreover, for execution on large parallel architectures, the bash calls from Python must comply with the constraints of an HPC environment, such as requests to send a job in a shared queue and wait for its execution. In the present work, these issues have been overcome using C-MPI wrappers for Python, with the *mpi4py* package (Dalcin and Fang 2021). Specifically, with this implementation, the main DRL code spawns N processes with the DNS executable, allowing a DNS to take place in an environment that is aware of the parent process by which it was spawned. This is formalised in the code by modifying the MPI initialisation in the DNS code, as shown by the example provided in Fig. 6.

When the main code spawns the N DNS parallel processes with `mpi_comm_spawn` command, the DNS simulation is initialised with an MPI request to merge all the ranks and create an *intracomm* communicator that contains all of the processes with 0 corresponding to the agent and $procs = 1 \dots N$ to the processes allocated to the DNS. Alongside the

definition of a global set of processes, we also need to generate other two communicators, one encompassing all the DNS processes (i.e. the, *canscomm* communicator spanning $procs = 0 \dots N - 1$) and the other dedicated to the DRL agent. Unfortunately, collective communications in the overall communicator (i.e. *intracomm*) are not available yet within the *mpi4py* library, so when the DNS solver has to share data with the DRL agent, *mpi_gather* and *mpi_gatherv* are used to collect data on rank 0 of *canscomm* (which is also rank 1 of *intracomm*), followed by a simple *mpi_send* from rank 1 to 0 in *intracomm* to make the data available for the Python library that manages the agent.

Another important limitation of the actual *mpi4py* library, concerns the *mpi_comm_spawn* command that does not work on distributed resources, such as multiple nodes of an HPC environment. This sets a limit to the maximum number of ranks that can be used by the DNS code to $N - 1$, being N the number of cores available on a single node. This limitation is a serious handicap also for the maximum amount of memory that can be used.

From a more practical perspective, it should be mentioned that the training based on *StableBaselines3* and *CaNS-DNS* code is much more complex, in terms of subroutines and scripts to keep under control, than the one required for the previous implementation with *SUSA* and *Tensorforce*. The main difficulty being related with the management of the communication between the environment and the ML libraries that now takes place using MPI calls and not I/O written on hard disks. Although further improvements are possible, the new implementation offers important advantages in terms of computational efficiency.

3 Results

The initial training session, with the reward r based on the skin-friction coefficient C_f (see 10), focused on achieving an amount of skin-friction drag reduction (DR) comparable to the results by Quadrio et al. (2009). This goal was met, as shown in Fig. 7. However, the strategies developed were only effective for $t_{ep}^+ = t_{ep} U_b / H \approx 80$ time units. We are currently investigating the cause of this limitation, with the long-term objective of controlling the flow for an arbitrary duration.

Figures 8 and 9 show the time history of the actuation $\omega(t)$ applied during this first training session, reported in Fig. 10. It is noticeable that the values tend to be negative, unlike the best-performing ω obtained by Quadrio et al. Furthermore, when compared with their results, it seems that a mild improvement in DR is achieved. However, this conclusion should be taken with caution since the average field may not have reached statistical convergence across the episodes.

The subsequent experiment aimed at maximising the ratio $\tilde{E}_{ks} / \tilde{E}_k$, focusing on the enhancement of the distribution of the streamwise velocity fluctuations prioritising the presence of straight near-wall velocity streaks. The latter are characterised as coherent streamwise, alternating positive and negative fluctuations regions that mark the region adjacent to the wall. They typically span about 100^+ in the spanwise direction and extend roughly 1000^+ in the streamwise length. Their disruption and regeneration are key to the self-sustaining cycle of wall turbulence, as described by Jiménez and Pinelli (1999). During the cycle, periods of low skin-friction coefficient coincide with straight streaks undergoing viscous diffusion, while bursts and undulations of the streaks are indicative of periods of high skin-friction drag. Therefore, the motivation driving this series of experiments was based on the hypothesis that maintaining the streaks in a straight configuration could inhibit or reduce their instability and, indirectly, reduce skin friction.

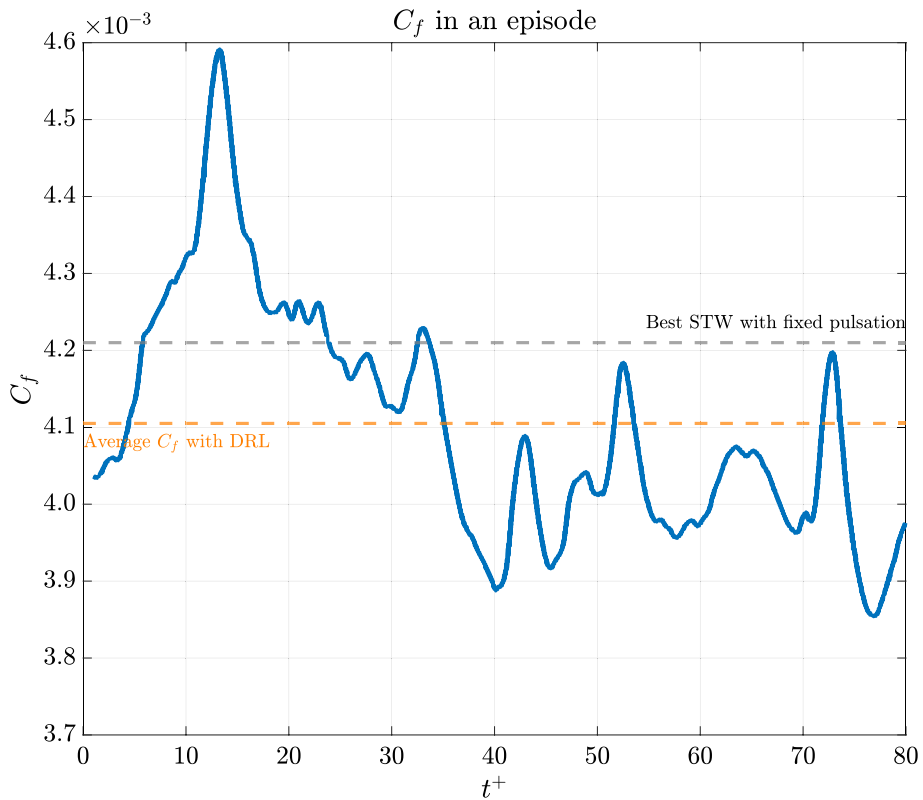


Fig. 7 Time evolution of C_f with the trained network, using 8 actuations per episode and implementing the best-performing strategy for ω

The exploration of the effects of this alternative reward was based on a set of modified parameters applied to both the DNS solver and the DRL libraries. Table 1, right column, summarises those modifications. In particular, we have now moved the plane used for data collection to 15^+ , while previously it was at $\approx 1^+$. This location should provide richer information compared to the previous choice that mainly reflects the state of the laminar sub-layer.

Moreover, $y^+ = 15$ is considered to be an optimal choice for observation planes in DR and flow control feedback (Hammond et al. 1998; Chung and Talha 2011). However, it should also be mentioned that strategies based on sampling at $y^+ = 1$ remain still appealing due to the possibility of validating with experiments using non-intrusive measurements of velocity and stresses at the wall.

Regarding the results obtained when pursuing the maximisation of the streaks energy content, from the time history of C_f during an episode, reported in Fig. 11, two conclusions can be drawn. Firstly, that we do obtain a significant *indirect* drag reduction although limited to 20% (against the 45% of Quadrio et al.). Secondly, that the episode duration seems to be sufficient to condition the short-term, transient behaviour of the flow, but not long enough to establish a consistent drag-reducing scenario.

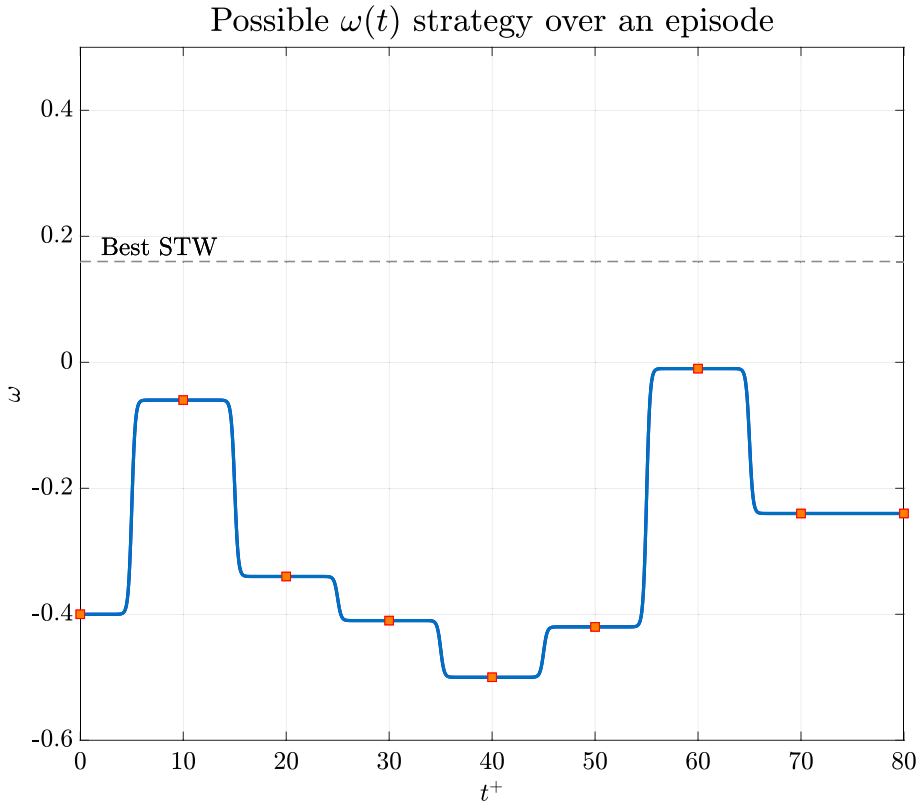


Fig. 8 Variation of ω during one of the episodes corresponding to the best C_f performance shown in Fig. 7

It is also noticed that the value of the reward approaches $r = 6.7$, as shown in Fig. 12. For this reward value, the DRL policy is not able to reach values of the ratio $\tilde{E}_{ks}/\tilde{E}_k$ greater than 42%, with the upper bound not known yet for this kind of control. This limitation may indicate the cause-and-effect relationships involved in the regeneration cycle and suggests that other integral quantities might be more effective in manipulating the flow behaviour (e.g., sweep events may not be controlled by the streaks' topology).

Nonetheless, the instantaneous streamwise velocity snapshots extracted on a wall-parallel plane at $y^+ = 20$ for the actuated and non-actuated cases, shown in Figs. 13 and 14, show that the maximisation of the energy content of the streamwise velocity fluctuations does have an impact on the flow topology. In particular, the actuated case features a smoother behaviour with less entanglement between low and high speed velocity streaks that present a higher, coherent intensity.

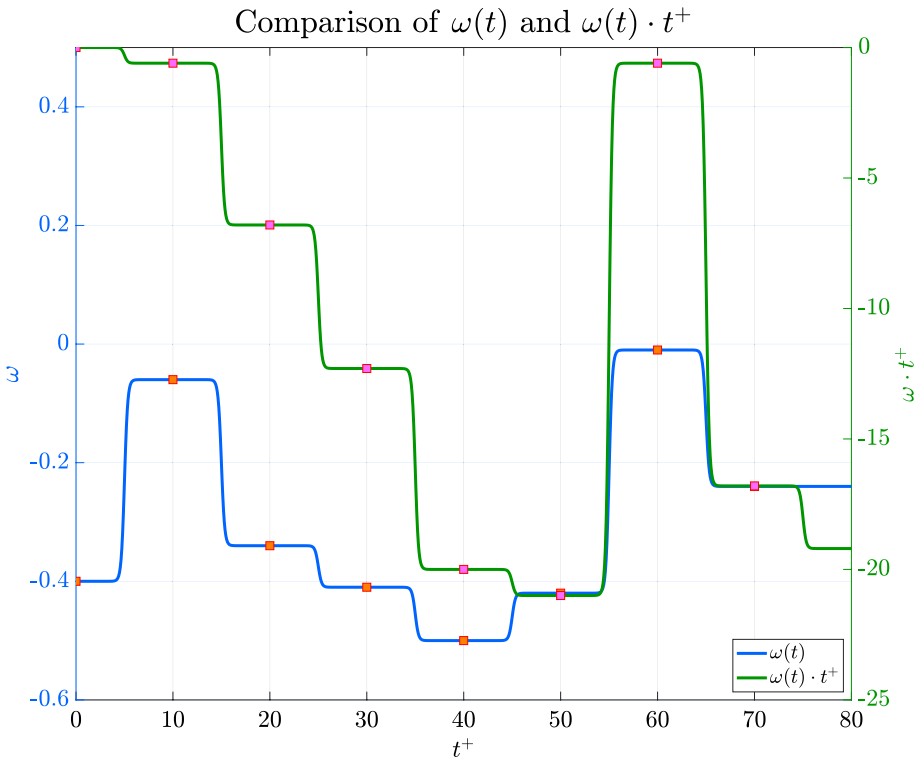


Fig. 9 Comparison between the behaviour of ω , in blue, and the value of the phase $\phi = \omega(t)t^+$, in green, with values reported on the right. It is possible to notice that the phase tends to get more and more negative during the episode, until a bump towards the end of it that makes it assume a similar value to the starting one. However, it never exceeds it since the values of ω remain negative

4 Discussion and Conclusion

The availability of a code featuring a clear, customisable interface enables the proposal of dynamic flow control techniques. Apart from the technological benefit, the same tool can be used to enhance or dampen the effect of various coherent structures embedded in the wall region, opening new avenues in understanding the wall cycle and its manipulation. Here, we presented a case that focused on the fluctuating energy distribution, but other rewards can be envisaged, such as the minimisation of the fourth quadrant events [i.e., *sweeps* (Hammond et al. 1998)] in the streamwise-wall-normal velocity fluctuations distribution.

Although we have presented only two physical realisations, the primary focus of the present contribution is on the implementation of DRL control techniques within the framework of direct numerical simulations. In the context of this work, several key features and choices need to be highlighted and further explored.

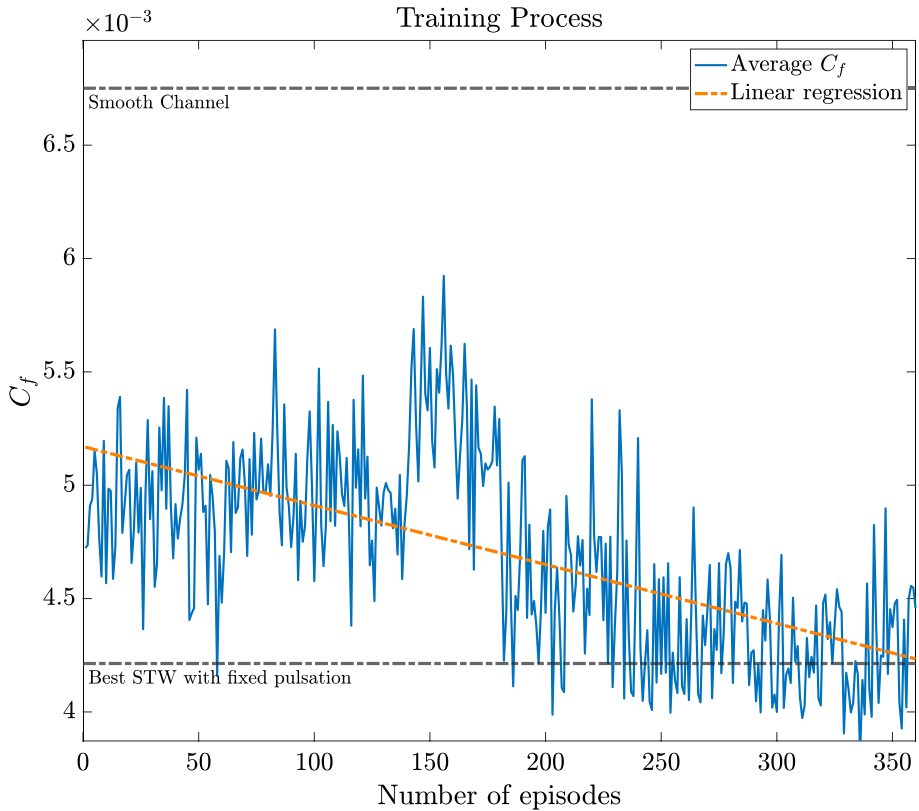


Fig. 10 Training process to minimise C_f . The horizontal axis corresponds to the number of episodes of the training, while on the vertical axis the C_f value, averaged over each episode, is shown. The network starts learning better strategies after 180 episodes, until it reaches a plateau, after which it is able to outperform the standard STW. The training is overall quite linear, never exceeding the C_f of a case with non-actuated wall condition

4.1 Actuation Duration

One of the most important parameters to set concerns the time period between consecutive $\omega(t)$ adjustments (see Fig. 8). Initially, the time interval between actuations $\Delta\tau^+$ was set to 10^+ . This choice corresponds to a quarter of the standard STW period when ω is fixed in time in the maximum DR configuration. It was assumed that a possible control law defined by the DRL algorithm would require an actuation update on a similar time scale. Note that $\Delta\tau^+ = 10$ implies up to four actuations on ω per period of the standard STW, thus giving sufficient relaxation time to the near wall turbulence.

Although we were able to find some interesting results, the choice of a piecewise constant actuation over a certain $\Delta\tau^+$ may result in a too restrictive control law that does not allow for the determination of the optimal $\omega(t)$.

The duration of the actuation period can be optimised based on physical considerations. We could start by defining an event \mathbb{E} as a general change in the topology of the structures within the logarithmic layer. The likelihood of this event occurring depends on the conditions at the wall a few time units prior. Ideally, this delay should be slightly shorter, but not

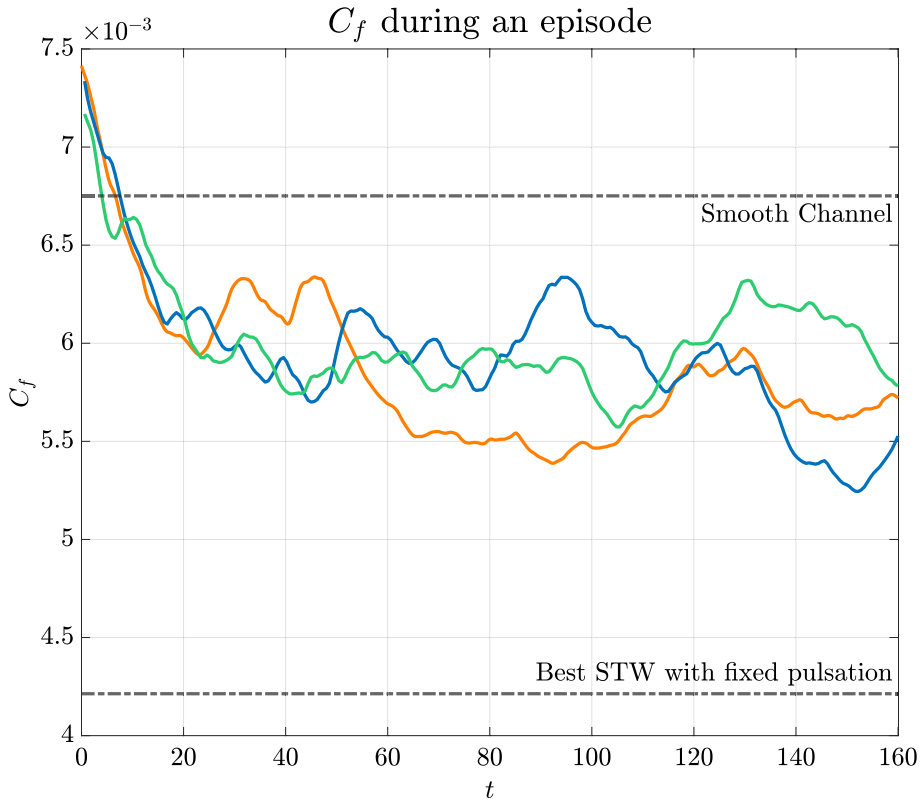


Fig. 11 Training of the DRL network—16 actuations per episode, with maximisation of streamwise velocity fluctuations as a reward (see Eq. 13). Different colours refer to different runs, which converge to the same value of dp/dx

too much shorter, than the $\Delta\tau^+$ chosen for the actuation period. This would ensure that the action imposed by the STW within the actuation period would have an impact on the close-to-the-wall structures. The agent would also learn faster because the correlation between its actions and the change in the environment would be more direct.

4.2 Hyper-Parameter Tuning

When tackling increasingly complex optimisation problems, it is important to tune the DRL hyper-parameters. *Hyper-parameters* are parameters that can be adjusted to modify the learning behaviour and the agent's features, such as the learning rate and batch size. These hyper-parameters can be adjusted depending on the agent's policy without altering the DNS environment. However, their impact on training cannot usually be assessed beforehand, necessitating iterative processes to develop and refine a meaningful policy. This process is time-consuming and difficult to estimate in advance. Future studies will include this tuning to avoid sub-optimal policies and ensure convergence toward a satisfactory reward.



Fig. 12 Reward computed as $\tilde{E}_{ks}/\tilde{E}_k$ per training step. The training reaches a plateau at a value higher than the standard STW

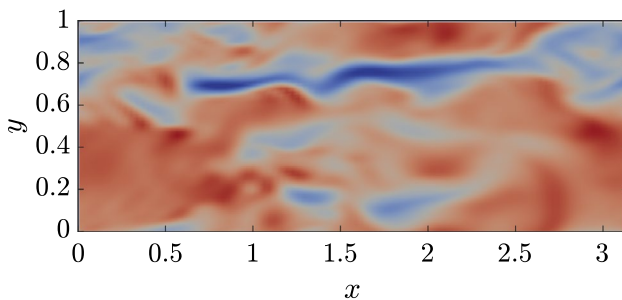


Fig. 13 $y^+ = 20$, non-actuated channel $\tilde{E}_{ks}/\tilde{E}_k = 33\%$

4.3 New Computational Architectures

The numerical simulation of turbulence still predominantly implies the use of CPU-based computer architectures for a series of reasons previously mentioned in the introduction. GPU-based solvers started to appear in the last two decades, with outstanding performances

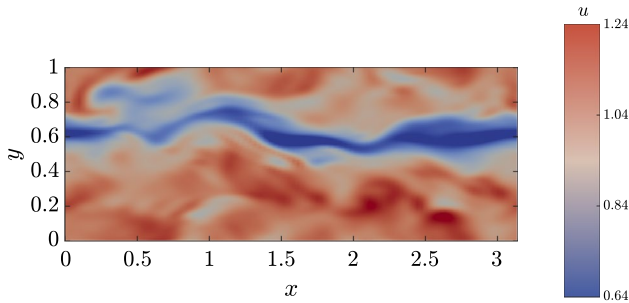


Fig. 14 $y^+ = 20$, $\omega(t)$ STW-DRL control $\tilde{E}_{ks}/\tilde{E}_k = 42\%$

when applied to fully explicit algorithms typical of compressible flow solvers where there is no need for the solution of a pressure Poisson equation. GPU computing for incompressible solvers is more recent although it has been proven to have the potential to exceed the performances of standard CPU-based solvers (Karp et al. 2022).

Since DRL libraries based on Python are already routinely used on GPUs architecture and the DNS environment is based on CaNS, that offers compiling instructions to work for GPUs architectures, we plan to migrate all the developed software on GPUs or hybrid GPU-CPU platforms. This should enhance the performances of the overall numerical methodology which will become of paramount importance when using a multi-agent approach in the framework of the previously introduced spanwise strips.

4.4 Conclusion

This study was motivated by the growing opportunities arising from recent algorithms advancements in deep reinforcement learning (DRL) and deep learning. The DRL methodology and related tools are promising candidates for improving flow control techniques by enabling actuations to be governed by policies derived from training a DRL agent. They can also be used to explore the underlying physics of wall turbulence by performing numerical experiments that enhance the presence and the interactions between flow structures. This task may be pursued more efficiently with other types of actuations different from the ones we are currently testing.

Appendix 1: Training Algorithm—Single Step

Algorithm 1 Schematised step of DRL training

```

Input: action
Output: observation, reward, flag
1 Check act; // act is the actuation number, starting from 0
2 if act == 0 then
3 | send CaNS START instructions;
4 else
5 | send CaNS CONTN instructions;
6 Map action  $\in [-1, 1] \mapsto [\omega_{min}, \omega_{max}]$ ; // This is needed to rescale  $\omega$ 
   within boundaries imposed by the user. action comes from
   StableBaselines3 routines used when the NN weights are updated.
7 Send  $\omega$  to CaNS;
8 Wait for NS eqs to be solved for  $t_{step}$ ; // time between actuations
9 Receive observations obs from CaNS;
10 Receive rewtemp from CaNS; // rewtemp is a quantity meaningful to
   judge the state obtained in the DNS, like  $C_f$ 
11 Convert numpy array obs into a grey-scale image; // Needed with CNNs
12 Compute reward rew =  $f(rew_{temp})$ ; //  $rew \in [0, 1]$ 
13 Update act = act + 1;
14 Update count = count + 1; // count keeps track of all the steps
   performed across all the episodes
15 if act ==  $n_{act}$  then
16 | flag = episode finished;
17 | if count ==  $n_{tot}$  then
18 | | send CaNS ENDED instructions; //  $n_{act}$  is the number of
   | | actuations in an episode,  $n_{tot}$  is the total number of
   | | timesteps to complete all the episodes
19 | else
20 | | check, possible I/O issues;
21 else
22 | flag = episode not finished;
23 | send CaNS CONTN instructions;

```

Funding R.V. was supported by ERC Grant No. 2021-CoG-101043998, DEEPCONTROL. The views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of European Union or European Research Council. G.C. would like to acknowledge the support of the EPSRC through their DTP studentships programme, grant number EP/W524608/1.

Declarations

Conflict of interest The authors report no conflict of interest.

Data Access Assessment The relevant data has been provided in the article, but further data can be provided upon request. The repository to access the DRL-DNS interface embedded in CaNS can be found at the following link: https://github.com/gmcavallazzi/CaNS_DRL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Beintema, G., Corbetta, A., Biferale, L., et al.: Controlling Rayleigh–Bénard convection via reinforcement learning. *J. Turbul.* **21**(9–10), 585–605 (2020)
- Belus, V., Rabault, J., Viquerat, J., et al.: 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 (2019)
- Biferale, L., Bonaccorso, F., Buzzicotti, M., et al.: Zermelo's problem: optimal point-to-point navigation in 2d turbulent flows using reinforcement learning. *Chaos Interdiscip. J. Nonlinear Sci.* **29**(10), 103138 (2019)
- Bucci, M.A., Semeraro, O., Allauzen, A., et al.: Control of chaotic systems by deep reinforcement learning. *Proc. R. Soc. A* **475**(2231), 20190351 (2019)
- Choi, H., Moin, P., Kim, J.: Active turbulence control for drag reduction in wall-bounded flows. *J. Fluid Mech.* **262**, 75–110 (1994). <https://doi.org/10.1017/S0022112094000431>
- Chorin, A.J.: Numerical solution of the Navier–Stokes equations. *Math. Comput.* **22**(104), 745–76 (1968). <https://doi.org/10.2307/2004575>
- Chung, Y.M., Talha, T.: Effectiveness of active flow control for turbulent skin friction drag reduction. *Phys. Fluids* **23**(2), 02510 (2011). <https://doi.org/10.1063/1.3553278>
- Cooley, J.W., Tukey, J.W.: An algorithm for the machine calculation of complex Fourier series. *Math. Comput.* **19**(90), 297–300 (1965). <https://doi.org/10.2307/2003354>
- Costa, P.: A FFT-based finite-difference solver for massively-parallel direct numerical simulations of turbulent flows. *Comput. Math. Appl.* **76**(8), 1853–1862 (2018). <https://doi.org/10.1016/j.camwa.2018.07.034>
- Costa, P.: Cans world. <https://github.com/CaNS-World/CaNS>, version 2.0 (2023)
- Dalcin, L., Fang, Y.L.L.: Mpi4py: status update after 12 years of development. *Comput. Sci. Eng.* **23**(4), 47–54 (2021). <https://doi.org/10.1109/MCSE.2021.3083216>
- Doohan, P., Willis, A.P., Hwang, Y.: Minimal multi-scale dynamics of near-wall turbulence. *J. Fluid Mech.* **913**, A8 (2021). <https://doi.org/10.1017/jfm.2020.1182>
- Fan, D., Yang, L., Wang, Z., et al.: Reinforcement learning for bluff body active flow control in experiments and simulations. *Proc. Natl. Acad. Sci.* **117**(42), 26091–26098 (2020)
- Font, B., Alcántara-Ávila, F., Rabault, J., et al.: Active flow control of a turbulent separation bubble through deep reinforcement learning. *J. Phys.: Conf. Ser.* **2753**(1), 01202 (2024). <https://doi.org/10.1088/1742-6596/2753/1/012022>
- Henry de Frahan, M.T., Wimer, N.T., Yellapantula, S., et al.: Deep reinforcement learning for dynamic control of fuel injection timing in multi-pulse compression ignition engines. *Int. J. Engine Res.* **23**(9), 1503–1521 (2022)
- Garnier, P., Viquerat, J., Rabault, J., et al.: A review on deep reinforcement learning for fluid mechanics. *Comput. Fluids* **225**, 10497 (2021). <https://doi.org/10.1016/j.compfluid.2021.104973>
- Gatti, D., Quadrio, M.: Reynolds-number dependence of turbulent skin-friction drag reduction induced by spanwise forcing. *J. Fluid Mech.* **802**, 553–58 (2016). <https://doi.org/10.1017/jfm.2016.485>
- Guastoni, L., Güemes, A., Ianiro, A., et al.: Convolutional-network models to predict wall-bounded turbulence from wall quantities. *J. Fluid Mech.* **928**, A2 (2021). <https://doi.org/10.1017/jfm.2021.812>
- Guastoni, L., Rabault, J., Schlatter, P., et al.: Deep reinforcement learning for turbulent drag reduction in channel flows. *Eur. Phys. J. E* **46**(4), 2 (2023). <https://doi.org/10.1140/epje/s10189-023-00285-8>

- Hammond, E.P., Bewley, T.R., Moin, P.: Observed mechanisms for turbulence attenuation and enhancement in opposition-controlled wall-bounded flows. *Phys. Fluids* **10**(9), 2421–242 (1998). <https://doi.org/10.1063/1.869759>
- Harris, C.R., Millman, K.J., van der Walt, S.J., et al.: Array programming with NumPy. *Nature* **585**(7825), 357–36 (2020). <https://doi.org/10.1038/s41586-020-2649-2>
- Jiménez, J., Moin, P.: The minimal flow unit in near-wall turbulence. *J. Fluid Mech.* **225**, 213–24 (1991). <https://doi.org/10.1017/S0022112091002033>
- Jiménez, J., Pinelli, A.: The autonomous cycle of near-wall turbulence. *J. Fluid Mech.* **389**, 335–35 (1999). <https://doi.org/10.1017/S0022112099005066>
- Karp, M., Massaro, D., Jansson, N., et al.: Large-scale direct numerical simulations of turbulence using GPUs and modern Fortran. *Int. J. High Perform. Comput. Appl.* **37**(5), 487–502 (2022)
- Kawahara, G., Jiménez, J., Uhlmann, M., et al.: Linear instability of a corrugated vortex sheet—a model for streak instability. *J. Fluid Mech.* **483**, 315–34 (2003). <https://doi.org/10.1017/S002211200300421X>
- Kim, J., Hussain, F.: Propagation velocity of perturbations in turbulent channel flow. *Phys. Fluids A* **5**(3), 695–706 (1993)
- LeCun, Y., Boser, B., Denker, J.S., et al.: Backpropagation applied to handwritten zip code recognition. *Neural Comput.* **1**(4), 541–55 (1989). <https://doi.org/10.1162/neco.1989.1.4.541>
- LeCun, Y., Bottou, L., Bengio, Y., et al.: Gradient-based learning applied to document recognition. *Proc. IEEE* **86**(11), 2278–232 (1998). <https://doi.org/10.1109/5.726791>
- LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**(7553), 436–44 (2015). <https://doi.org/10.1038/nature14539>
- Li Y.: Deep reinforcement learning: An overview. [arXiv:1701.07274](https://arxiv.org/abs/1701.07274) (2018)
- Lillicrap, T.P., Hunt, J.J., Pritzel, A., et al.: Continuous control with deep reinforcement learning. [arXiv:1509.02971](https://arxiv.org/abs/1509.02971) (2019)
- Monti, A., Omidyeganeh, M., Pinelli, A.: Large-eddy simulation of an open-channel flow bounded by a semi-dense rigid filamentous canopy: scaling and flow structure. *Phys. Fluids* **31**(6), 06510 (2019). <https://doi.org/10.1063/1.5095770>
- Monti, A., Omidyeganeh, M., Eckhardt, B., et al.: On the genesis of different regimes in canopy flows: a numerical investigation. *J. Fluid Mech.* **891**, A9 (2020). <https://doi.org/10.1017/jfm.2020.155>
- Moser, R.D., Kim, J., Mansour, N.N.: Direct numerical simulation of turbulent channel flow up to $Re\tau=590$. *Phys. Fluids* **11**(4), 943–945 (1999)
- Novati, G., de Laroussilhe, H.L., Koumoutsakos, P.: Automating turbulence modelling by multi-agent reinforcement learning. *Nat. Mach. Intell.* **3**(1), 87–9 (2021). <https://doi.org/10.1038/s42256-020-00272-0>
- Paris, R., Beneddine, S., Dandois, J.: Robust flow control and optimal sensor placement using deep reinforcement learning. *J. Fluid Mech.* **913**, A25 (2021)
- Paszke, A., Gross, S., Massa, F., et al.: PyTorch: an imperative style, high-performance deep learning library. In: Proceedings of the 33rd International Conference on Neural Information Processing Systems, vol. 721, Curran Associates Inc., Red Hook, NY, USA, pp. 8026–8037 (2019)
- Quadrio, M., Ricco, P., Viotti, C.: Streamwise-travelling waves of spanwise wall velocity for turbulent drag reduction. *J. Fluid Mech.* **627**, 161–17 (2009). <https://doi.org/10.1017/S0022112009006077>
- Rabault, J., Kuhnle, A.: Accelerating deep reinforcement learning strategies of flow control through a multi-environment approach. *Phys. Fluids* **31**(9), 094105 (2019)
- Rabault, J., Kuchta, M., Jensen, A., et al.: Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control. *J. Fluid Mech.* **865**, 281–30 (2019). <https://doi.org/10.1017/jfm.2019.62>
- Raffin, A., Hill, A., Gleave, A., et al.: Stable-baselines3: reliable reinforcement learning implementations. *J. Mach. Learn. Res.* **22**(268), 1–8 (2021)
- Rai, M.M., Moin, P.: Direct simulations of turbulent flow using finite-difference schemes. *J. Comput. Phys.* **96**(1), 15–53 (1991)
- Ren, F., Rabault, J., Tang, H.: Applying deep reinforcement learning to active flow control in weakly turbulent conditions. *Phys. Fluids* **33**(3), 037121 (2021)
- Rhie, C.M., Chow, W.L.: Numerical study of the turbulent flow past an airfoil with trailing edge separation. *AIAA J.* **21**(11), 1525–1532 (1983)
- Schulman, J., Wolski, F., Dhariwal, P., et al.: Proximal policy optimization algorithms. [arXiv:1707.06347](https://arxiv.org/abs/1707.06347) (2017)
- Sonoda, T., Liu, Z., Itoh, T., et al.: Reinforcement learning of control strategies for reducing skin friction drag in a fully developed turbulent channel flow. *J. Fluid Mech.* **960**, A3 (2023). <https://doi.org/10.1017/jfm.2023.147>

- Suárez, P., Alcántara-Ávila, F., Rabault, J., et al.: Flow control of three-dimensional cylinders transitioning to turbulence via multi-agent reinforcement learning. *arXiv:2405.17210* (2024)
- Tang, H., Rabault, J., Kuhnle, A., et al.: 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 (2020)
- Varela, P., Suárez, P., Alcántara-Ávila, F., et al.: Deep reinforcement learning for flow control exploits different physics for increasing Reynolds number regimes. In: *Actuators*, MDPI, p. 359 (2022)
- Verma, S., Novati, G., Koumoutsakos, P.: Efficient collective swimming by harnessing vortices through deep reinforcement learning. *Proc. Natl. Acad. Sci.* **115**(23), 5849–5854 (2018)
- Vignon, C., Rabault, J., Vasanth, J., et al.: Effective control of two-dimensional Rayleigh–Bénard convection: invariant multi-agent reinforcement learning is all you need. *Phys. Fluids* **35**(6), 06514 (2023). <https://doi.org/10.1063/5.0153181>
- Vignon, C., Rabault, J., Vinuesa, R.: Recent advances in applying deep reinforcement learning for flow control: perspectives and future directions. *Phys. Fluids* **35**(3), 03130 (2023). <https://doi.org/10.1063/5.0143913>. *arXiv: 2304.03181* [physics]
- Vinuesa, R., Lehmkuhl, O., Lozano-Durán, A., et al.: Flow control in wings and discovery of novel approaches via deep reinforcement learning. *Fluids* **7**(2), 62 (2022)
- Wälchli, D., Guastoni, L., Vinuesa, R., et al.: Drag reduction in a minimal channel flow with scientific multi-agent reinforcement learning. *J. Phys. Conf. Ser.* **2753**, 012024 (2024)
- Xu, H., Zhang, W., Deng, J., et al.: Active flow control with rotating cylinders by an artificial neural network trained by deep reinforcement learning. *J. Hydrodyn.* **32**(2), 254–258 (2020)
- Zeng, K., Linot, A.J., Graham, M.D.: Data-driven control of spatiotemporal chaos with reduced-order neural ODE-based models and reinforcement learning. *Proc. R. Soc. A* **478**(2267), 20220297 (2022)
- Zheng, C., Ji, T., Xie, F., et al.: From active learning to deep reinforcement learning: intelligent active flow control in suppressing vortex-induced vibration. *Phys. Fluids* **33**(6), 063607 (2021)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Giorgio Maria Cavallazzi¹ · Luca Guastoni^{2,3} · Ricardo Vinuesa³ · Alfredo Pinelli¹

✉ Giorgio Maria Cavallazzi
giorgio.cavallazzi@city.ac.uk

Luca Guastoni
luca.guastoni@tum.de

Ricardo Vinuesa
rvinuesa@mech.kth.se

Alfredo Pinelli
alfredo.pinelli.1@city.ac.uk

¹ Department of Engineering, City St George's, University of London, Northampton Square, London EC1V 0HB, UK

² School of Computation, Information and Technology, TU Munich, Boltzmannstr. 3, 85748 Garching, Germany

³ FLOW, Engineering Mechanics, KTH Royal Institute of Technology, 100 44 Stockholm, Sweden