Physics-informed Neural-operator Predictive Control for Drag Reduction in Turbulent Flows

Abstract

Turbulence control for wall friction reduction poses a significant challenge due to computational costs associated with modeling turbulent dynamics. In this work, we present a learning and control scheme termed physics-informed neural-operator-based predictive control (PINO-PC). Our method uses a Neural Operator framework that enables accurate learning and control of turbulent flows. It carries out predictive control where both the policy and the observer model for turbulence control are learned jointly. We show that PINO-PC outperforms prior model-free reinforcement learning methods with a stable policy learning procedure. We experiment with various challenging generalized scenarios where flows are of unseen high Reynolds numbers, and we find that our method achieves a drag reduction of 39.0% under a bulk-velocity Reynolds number of $15k$, outperforming previous fluid control methods by more than 32%.

Keywords: Drag reduction, Fluid control, Neural operators, Machine-learning-based control

1 Main

Turbulent flows are prevalent in many areas of science and engineering, such as atmospheric weather [1], ocean currents [2], and blood flow in arteries [3] and veins [4]. The turbulent flow is generally more unstable when compared to laminar flow and has a higher skin friction drag, which is caused by the friction of a fluid moving against a surface or a wall. Reducing such drag and controlling turbulent flows is necessary for various applications such as aerospace engineering, fluid transport, and biomedical devices [5], since it mitigates adverse effects associated with turbulence, such as increased energy consumption, reduced efficiency, and heightened mechanical stress on structures. By gaining a deeper understanding of turbulent flow dynamics and implementing effective control strategies, we can enhance performance, optimize design, and ensure the safety and reliability of various engineering systems and biological processes.

The standard approach to controlling turbulent channel flows [6–10] involves blowing and suction of fluid air at chosen positions along the wall to control the boundary velocity Fig 1a. There are both passive and active control approaches. Passive approaches do not have a feedback loop for obtaining information and employing them for controlling flows [11, 12], and hence, are generally inferior in drag reduction compared to active control methods [6–8], which have a feedback loop and dynamically reduce skin friction through suction and blowing at the wall.

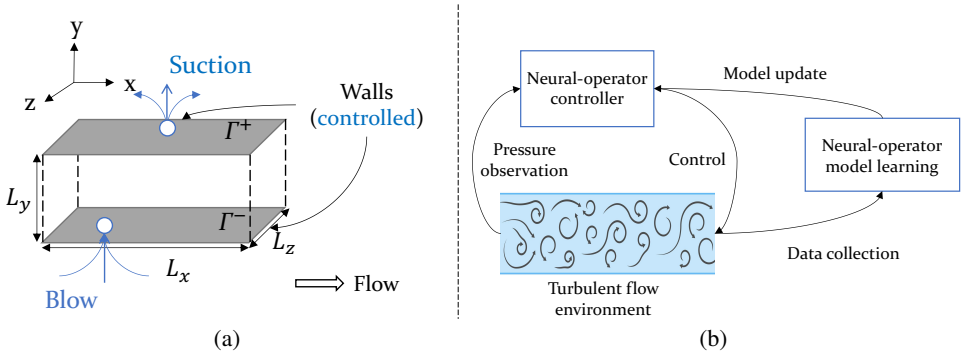

Fig. 1: **(a)**: Channel flow moving along the streamwise direction x, with the control applied at the wall via suction or blowing. **(b)**: Overall schematic of PINO-PC. The controller takes pressure observations as input, performs predictive control, and applies the control to the turbulent flow environment. Data is collected from the turbulent flow environment and is used to train the model.

A simple active control method in wall-bounded flows, known as "opposition control" [6], proposes applying blowing and suction at the wall, opposite to the normal velocity at a plane off the wall called a "detection plane". While such an approach achieves a reasonable drag reduction, it requires placement of sensors at the detection plane, which is often impractical [7, 9]. A machine-learning-based approach is introduced [9] to predict the velocity function on the detection plane based on boundary information (e.g., boundary pressure function). The opposition control approach is limited to computing the normal velocity and simply reversing it, which is considered myopic and is not an optimal control policy. As mentioned above, closed-loop feedback control is a potential approach to tackle this limitation, however, such a method requires the knowledge of a well-defined, often, state-space dynamical model of the system to synthesize controllers.

To address the above limitations in both learning and optimality, reinforcement learning (RL) methods are developed that control the drag reduction in turbulent flow [17]. In particular, deep deterministic policy gradient (DDPG) [18], are deployed to control flows, achieving superior drag reduction compared to opposition control [16]. It controls the flow by changing the mass flow rates of two jets on the sides of a cylinder. Rabault et al. [19] and Fan et al. [20] use RL methods to control the cylinder or bluff body flow. Tang et al. [21] proposes a smoothing technique to reduce the drag fluctuations while enabling the RL agent to generalize to unseen Reynolds numbers. Recently, [22] proposes to transfer discovered two-dimensional controls to three-dimensional cylinder flows via reinforcement learning. These RL approaches often have large variances during training [16] and have inferior performance when the flow is of a high turbulent level (at a high Reynolds number [23]).

The above-mentioned approaches are only applicable to a certain discretization of pressure, velocity, and other functional quantities of interest. These methods are often built on the assumption of full observability of the dynamics, do not incorporate physics information, and can cause unstable behavior in turbulence settings [16]. Besides, these approaches do not respect the functional nature of the problem, while in fluid dynamics, it is crucial to

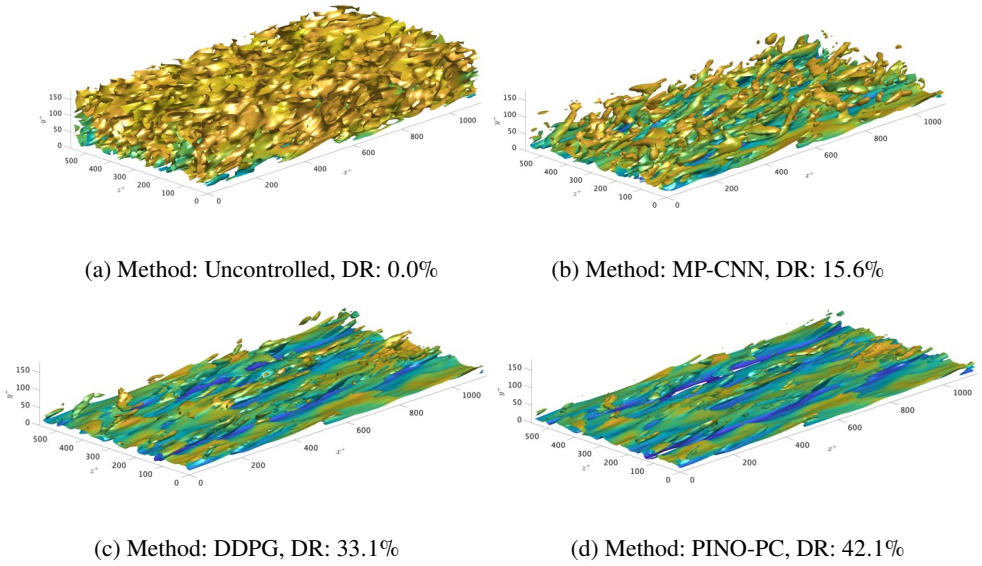

(a) Method: Uncontrolled, DR: 0.0%

(b) Method: MP-CNN, DR: 15.6%

(c) Method: DDPG, DR: 33.1%

(d) Method: PINO-PC, DR: 42.1%

Fig. 2: Visualizations of the isosurface of the Q-criterion (a vorticity indicator) [13–15] under the same threshold ($Q_\tau = 50$) after control with Reynolds number $\mathrm{Re}_u = 3k$. The isosurface is colored by the velocity magnitude. Four flows under four flow control methods are shown in this diagram: Uncontrolled, MP-CNN [9], DDPG [16], and our approach. Please refer to the main text for details.

capture fine-scale features and high-resolution details to accurately model and predict fluid behavior [8]. Without these, the models may fail to represent the complexity of fluid flows, especially in turbulent or complex flow regimes.

Our approach: We propose a deep RL predictive control framework for drag reduction on turbulent flows, termed as physics-informed neural operator predictive control (PINO-PC). This framework interactively collects data (pressure, velocity, and drag) from the turbulent flow environment and consists of two main components, the observer model and policy model, as illustrated in Figure 1b. It trains the observer model to predict the internal field velocity based on the applied control. The observer model is trained using the PINO framework, through physics-informed learning [24, 25]. Physics-informed learning only relies on the PDE but does not require data, which is crucial when data is limited such as flows of high Reynolds numbers. Simultaneously, the policy model is trained to minimize the kinetic energy and actuation norms on the trained observer predictive model and control cost function [7]. The policy model is deployed as a controller to predict the control based on the boundary pressure, and the control is applied to the turbulent flow environment.

The observer and policy models are built using neural operators [26], a generalization of standard neural networks, enabling model learning in function spaces. Utilizing the properties of neural operators, we formulate our RL approach on function spaces in which the control problem is defined. The observer model flexibly learns from data at various resolutions, and it outputs the next state in function spaces, enabling us to utilize the PINO framework for training

using PDE constraints. The policy neural operator model is resolution agnostic, making it flexible to operate on various observer models' resolutions and controller resolutions. The proposed framework borrows insights from conventional RL methods. PINO-PC utilizes the memory to store collected experiences, inspired by deep Q networks [27], a seminal deep RL method on discrete action spaces. It learns an observer model to predict action outcomes as a model-based RL component of our framework [28]. Nevertheless, unlike previous RL methods, our approach is based on neural operators, which enables policy learning in functional space. Furthermore, PINO-PC is a model-based approach because it learns to predict the outcome of the control, which is necessary to generalize to unseen Reynolds numbers.

Summary of empirical results: Our numerical experiments show that PINO-PC demonstrate a better drag reduction compared to alternative machine-learning approaches [26]: results from numerical simulations demonstrate the efficacy of our proposed control method, achieving a 43.5% reduction in drag for Reynolds numbers not included in the training data. Our approach outperforms both opposition control [6] by 24.9% and conventional optimal control techniques [7, 8] by 9.6% in challenging generalization scenarios. Also, it achieves a remarkable improvement of 26.5% in drag reduction when compared to MP-CNN, while it outperforms DDPG by 9.0% in drag reduction. Furthermore, PINO-PC leverages physics-informed learning, which boosts generalization to unseen flows. The experimental results show that the generalization performance can improve drag reduction performance up to 2.2% when using physics-informed learning compared to PINO-PC without physics-informed learning.

In this work, we addressed the challenge of turbulent flows in the wall-bounded scenario. We considered an active control setup, implementing control through blowing and suction at the wall. We proposed a novel machine-learning-based predictive control scheme, PINO-PC. This framework leveraged a policy model and an observer model. The policy model was used to predict the control (applied boundary velocity) based on the boundary pressure. The observer model predicted the control outcome (internal field velocity) based on the control.

Our approach demonstrated superior accuracy and drag reduction compared to alternative machine-learning methods. Notably, PINO-PC achieved a remarkable 43.5% drag reduction for Reynolds numbers not included in the training data, surpassing both opposition control and the optimal control baseline. The proposed iterative learning procedure, with extensive observer and policy learning, proved effective in achieving more robust turbulence control. This work provided a foundation for more efficient and practical turbulence control methodologies.

2 Results

We perform the direct numerical simulation (DNS) of a turbulent channel flow [6, 7, 9, 16]. The schematic of the channel flow is presented in Figure 1a. The simulations are based on discretizing the incompressible Navier-Stokes equations, while the equations are solved with an explicit third-order Runge Kutta (RK3) method for time advancement. The control is deployed by applying a normal velocity at the wall, while the control target is to minimize the drag. More details of the problem setting can be found in Section 3.1.

We experiment with flows of various Reynolds numbers to show the drag reduction results of different approaches. The first setup is ***Single***, which means the training and testing bulk-velocity Reynolds number is $\mathrm{Re}_b \approx 3k$. In this scenario, the friction Reynolds number is around $\mathrm{Re}_\tau \approx 178$, which is close to the default setup of previous studies [6, 9, 16, 29]. The

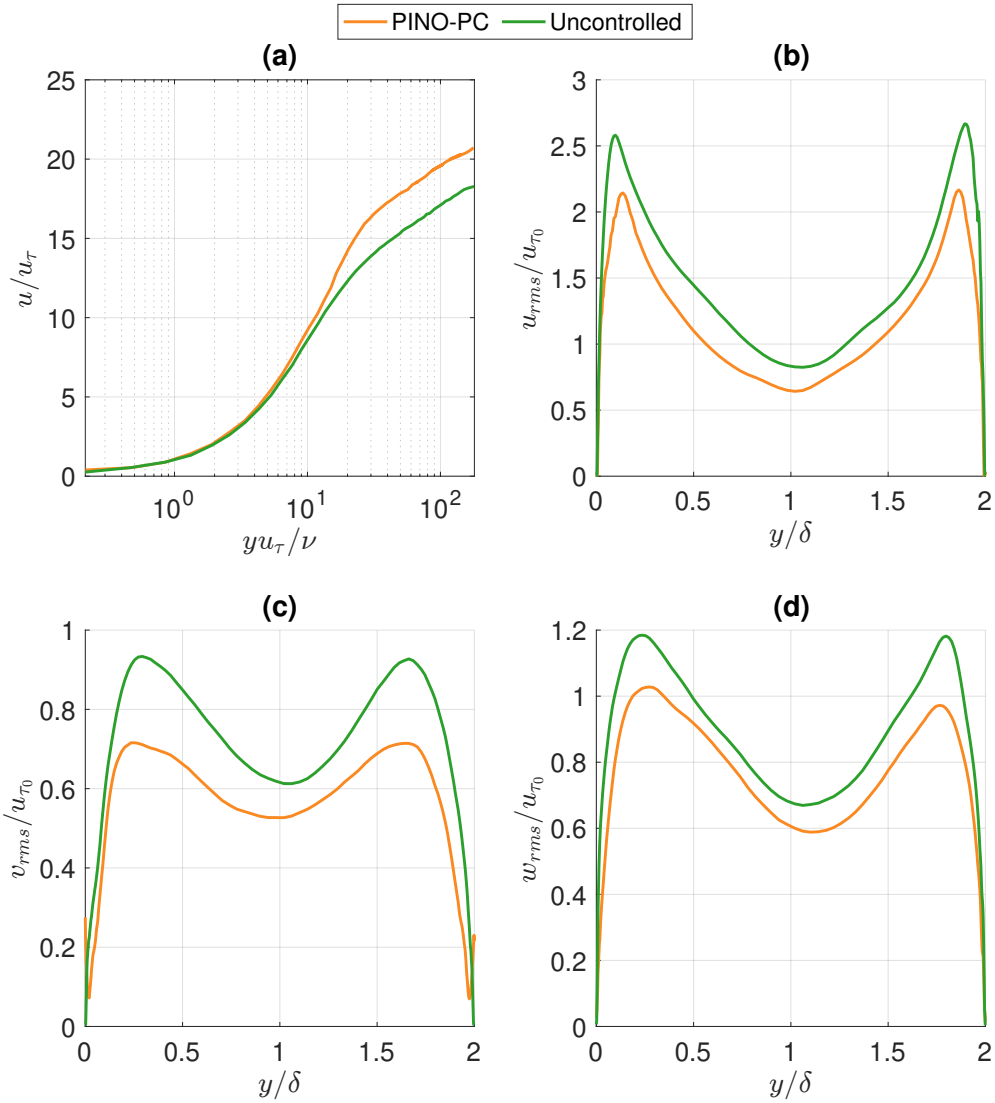

Fig. 3: Time averaged statistics of the uncontrolled flow and PINO-PC (after control) in the full channel flow case. (a) The mean stream-wise velocity profile. (b, c, d) Stream-wise, wall-normal, and span-wise r.m.s. velocity fluctuations in the wall-normal direction.

second setup ***Interpolation*** refers to cases where the testing Reynolds number is smaller than the largest training Reynolds number and larger than the smallest training Reynolds number. Third, ***Extrapolation-A*** occurs when the testing Reynolds number surpasses all Reynolds numbers in the training data. ***Extrapolation-B*** is another setup in which the model is solely trained on one low Reynolds number and is then tested on a wide range of higher Reynolds numbers. To reduce the effect of noises and randomness, we conduct each experiment with

Generalization setting	Single		Interpolation		Extrapolation-A		Extrapolation-B	
Phase	Training	Testing	Training	Testing	Training	Testing	Training	Testing
Re_b	3k	3k	3k, 6k, 9k, 15k	12k	3k, 6k, 9k, 12k	15k	3k	6k, 9k, 12k, 15k
Opposition control [6]	-	17.2 ±1.5	-	16.3 ±1.2	-	16.6 ±2.1	-	15.1 ±1.5
DNS-PC [8]	-	38.1 ±3.5	-	31.1 ±4.8	-	33.1 ±4.7	-	25.7 ±3.8
MP-CNN [9]	15.3 ±2.5	15.1 ±2.6	16.1 ±3.1	14.2 ±1.4	15.4 ±3.4	15.4 ±1.6	15.3 ±2.5	13.1 ±1.2
DDPG [16]	41.2 ±7.1	40.5 ±7.8	36.1 ±5.6	32.1 ±7.3	44.2 ±7.9	31.4 ±7.7	41.2 ±7.1	14.6 ±9.3
PINO-PC	**45.1** ±6.7	**45.5** ±5.4	**45.3** ±4.0	**41.2** ±4.6	**42.1** ±5.7	**38.5** ±6.4	**45.1** ±6.7	**33.5** ±5.5

Table 1: Performances in varied Reynolds numbers comparing several flow control methods in the **minimum channel flow** case. The metric is drag reduction rate (DR) in percentage. Each experiment is repeated three times while we report both mean and variance in this table. Opposition control [6] and DNS-PC [8] do not have training performance scores because they are not machine-learning-based methods. We experimented with different generalization settings, including ***Single***, ***Interpolation***, ***Extrapolation-A***, and ***Extrapolation-B***. Please refer to the main text to see descriptions of these settings.

Generalization setting	Single		Interpolation		Extrapolation-A		Extrapolation-B	
Phase	Training	Testing	Training	Testing	Training	Testing	Training	Testing
Re_b	3k	3k	3k, 6k, 9k, 15k	12k	3k, 6k, 9k, 12k	15k	3k	6k, 9k, 12k, 15k
Opposition control [6]	-	17.4 ±1.4	-	15.2 ±1.9	-	15.8 ±2.9	-	14.1 ±1.9
DNS-PC [8]	-	40.3 ±3.4	-	30.2 ±5.4	-	30.5 ±4.1	-	29.4 ±3.9
MP-CNN [9]	15.8 ±2.3	15.6 ±2.4	18.4 ±3.2	15.2 ±1.6	15.9 ±3.5	16.1 ±1.4	15.8 ±2.3	13.4 ±2.0
DDPG [18]	34.1 ±6.9	33.1 ±5.9	36.2 ±5.2	31.4 ±7.0	38.2 ±8.1	32.5 ±7.9	34.1 ±6.9	14.6 ±9.1
PINO-PC	**43.5** ±3.9	**42.1** ±4.9	**43.1** ±3.9	**40.3** ±3.2	**40.1** ±6.2	**35.1** ±5.9	**43.5** ±3.9	**39.0** ±4.0

Table 2: Performances in varied Reynolds numbers comparing several flow control methods in the **full channel flow** case. The metric is the drag reduction rate (DR). Other setups are the same as Table 1.

three different flow initializations and show mean and standard errors in tables and curve plots. We use different flow initializations in training and testing to assess the generalization performance of machine learning models.

Table 1 presents numerical control results in the minimal channel. Opposition control [6] reported a drag reduction of $\approx 14\%$. Our results indicate that opposition control reaches a drag reduction of $\approx 17.2\%$ in the single Reynolds setup, which is in agreement with their result. Opposition control [6] has similar results in other Reynolds numbers under other settings. The machine-learning-based opposition control method called MP-CNN [9] has lower drag reduction rates in all settings than the traditional opposition control [6] because it does not perfectly imitate the opposition control policy. DNS-PC [8] has a much better drag reduction than opposition control because it can access interior information and optimizes control for a period of time. Our reproduction shows that they can achieve a drag reduction of $\approx 38.1\%$ in the minimum channel flow case, which is close to their report ($\approx 40\%$). DNS-PC [8] has lower performances in higher Reynolds numbers. DDPG [18] is a strong baseline and performs highly in the single Reynolds number setup. However, when scaled to higher Reynolds numbers, it suffers from overfitting and cannot perform well in the test

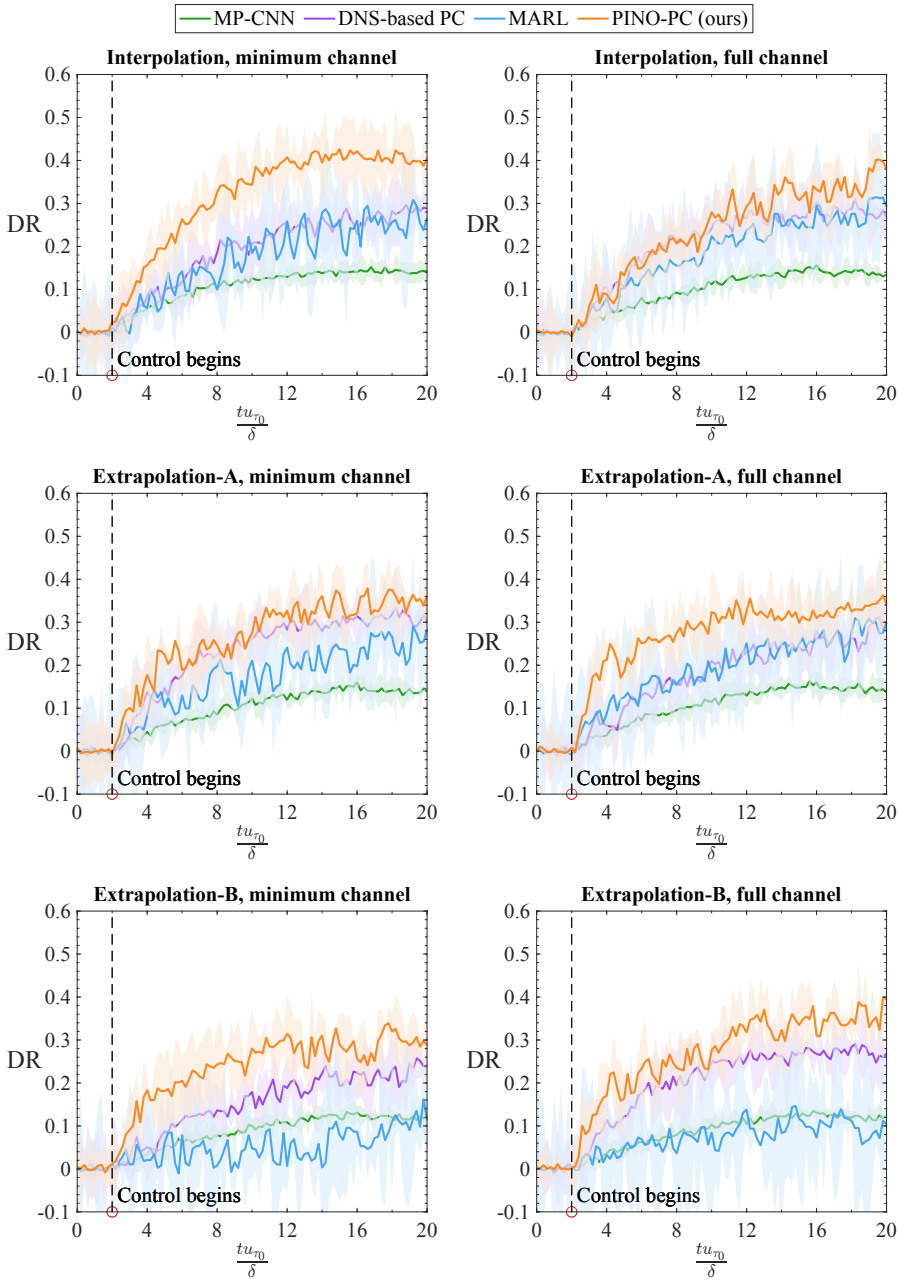

Fig. 4: Drag reduction curves comparing several flow control methods: MP-CNN [9], Local suboptimal [7], DDPG [16] and PINO-PC (ours). The x-axis denotes the non-dimensional timestep, and the y-axis denotes the drag reduction rate (DR) related to the uncontrolled case. The control beginning time is indicated via a red circle.

Generalization setting	Single		Interpolation		Extrapolation-A		Extrapolation-B	
Phase	Training	Testing	Training	Testing	Training	Testing	Training	Testing
Re_b	3k	3k	3k, 6k, 9k, 15k	12k	3k, 6k, 9k, 12k	15k	3k	6k, 9k, 12k, 15k
Ours (w/o PC)	34.3	33.5	36.0	31.1	38.0	32.3	34.3	15.3
Ours (RNO [30] → CNN [9])	39.8	36.5	40.7	37.9	34.9	31.6	41.8	36.9
Ours (RNO [30] → FNO [31])	43.4	42.0	42.9	39.9	39.4	33.8	43.4	40.8
Ours (w/o MF [32])	43.1	42.0	43.0	39.8	35.8	32.2	43.1	37.1
Ours (w/o L_{pde})	42.0	39.2	41.8	38.4	36.0	32.9	42.0	38.8
Ours	**43.5**	**42.1**	**43.1**	**40.3**	**40.1**	**35.1**	**43.5**	**39.0**

Table 3: Ablation studies of PINO-PC. The metric is the drag reduction rate in the full channel flow of varied Reynolds numbers, and please refer to Section 4.3 for more descriptions of mentioned ablations.

splits of Extrapolation-A and Extrapolation-B settings. Also, it has a larger variance than MP-CNN [9]. We find PINO-PC consistently outperforms other methods across different Reynolds numbers and generalization settings, achieving the highest mean drag reduction rates. In the test splits of Extrapolation-A and Extrapolation-B setups, PINO-PC performs much better than DDPG [18], which indicates the effectiveness of the proposed physics-informed learning scheme in narrowing the generalization gap. We also observe that PINO-PC has larger variances than MP-CNN [9] because it also needs to interact with the flows during training, which increases the training variance. Further, PINO-PC has a smaller variance than DDPG [18], which reveals the effect of using a physics-informed observer model to lower training variance.

We also provide the experimental results of the full channel flow in Table 2. We observe that opposition control [6] and MP-CNN [9] achieve a similar result in full channel flow compared to the minimum channel case. This is because opposition-control-based methods are not affected much by scales [9], and they usually have a smaller actuation intensity [16]. DNS-PC [8] demonstrates competitive performances in the full channel flow. It does not behave much differently in the minimum channel flow case because it resolves the optimization problem and calculates the control policy under each scenario. DDPG [18]'s performance downgrades significantly in the full channel flow compared to the minimum one. One possible explanation is that their policy model based on fully connected networks (FCN) is sensitive to the flow scale. Nonetheless, DDPG performs better than MP-CNN [9] and opposition control [6]. Furthermore, we observe that PINO-PC also has strong performance in the full channel flow case, which is primarily due to the fact neural operators [26] can scale to other input dimensions because they learn solutions in the function space.

Figure 4 shows comparison curves under several setups. We use shadowed regions to denote variance in this plot, and the control beginning time is marked by a red circle in the plot. We observe that opposition control [6] and MP-CNN [9] have smaller training variances with poor control outcomes. DNS-PC [8] performs strongly in some scenarios, especially in high Reynolds numbers. DDPG [16] has large training variances because deep RL algorithms often require extensive trial-and-error, and unsuccessful explorations fail to bring drag reduction. By contrast, PINO-PC has a smaller training variance than DDPG [16] and performs better, especially in high Reynolds numbers, which suggests the benefits of adopting the physics-informed neural-operator-based observer. Our approach further uncovers the power of machine-learning-based methods in reducing drag.

Figure 3 provides the flow statistics of adopting PINO-PC after control in the full channel flow case under a Reynolds number of $\mathrm{Re}_b = 3k$ and $\mathrm{Re}_\tau = 178$. We observe that the velocity fluctuations in three dimensions decrease after control. The turbulent kinetic energy (TKE) also decreases after control via PINO-PC, which verifies the effectiveness of the control algorithm from another perspective. These trends align with previous flow control works [6–8].

Finally, we conduct an isosurface visualization [13, 15, 33] to compare the performances of several control methods, and the result is presented in Figure 2. The setup is also a full-channel flow under a Reynolds number of $\mathrm{Re}_b = 3k$ and $\mathrm{Re}_\tau = 178$. The isosurface is formed by computing the Q-criterion [15], which is associated with vorticity in the flow. The Q-criterion can be computed as:

$$Q \equiv \frac{1}{2}\left(\|\boldsymbol{\Omega}\|^2 - \|\boldsymbol{S}\|^2\right), \tag{1}$$

where $\|\boldsymbol{S}\| = \left[\mathrm{tr}\left(\boldsymbol{S}\boldsymbol{S}^T\right)\right]^{1/2}$, $\|\boldsymbol{\Omega}\| = \left[\mathrm{tr}\left(\boldsymbol{\Omega}\boldsymbol{\Omega}^T\right)\right]^{1/2}$. Here, $\boldsymbol{S}$ and $\boldsymbol{\Omega}$ are the symmetric and anti-symmetric components of the velocity gradient tensor ∇u. Thus, $\boldsymbol{S}$ is the rate of strain tensor, while Ω is the vorticity tensor. Therefore, Q represents the difference between vorticity magnitude and shear strain rate. The more complex the isosurface of the Q-criterion is, the more vorticity is involved. We can observe from Figure 2 that PINO-PC has the most simple isosurface when compared to baselines (Uncontrolled, MP-CNN [9], DDPG [16]). In this plot, we color the isosurface via the velocity magnitude. A large velocity corresponds to a yellow-colored surface, while a smaller velocity is associated with a blue-colored surface. We can observe from this plot that the middle region of the channel flow is associated with large velocities, while Q vanishes at the wall [15].

3 Methods

In this section, we first introduce the problem setup in Section 3.1. We then introduce our proposed method, called physics-informed neural-operator-based predictive control (PINO-PC), in several upcoming subsections. In Section 3.2, we introduce the algorithm outline and overview of our proposed predictive control scheme. Subsequently, we propose details of two machine learning models adopted in our framework in Section 3.4 and Section 3.3.

3.1 Problem setting

In this work, we perform a direct numerical simulation (DNS) of a turbulent channel flow, which has been studied in previous drag reduction works [6, 7, 9, 16]. The flow domain is designed such that the x direction indicates the streamwise direction while y and z direction denote the wall-normal and spanwise directions, respectively.

3.1.1 The governing equation

The governing incompressible Navier-Stokes equations can be formulated as

$$\begin{cases} \dfrac{\partial u_j}{\partial x_j} = 0, \\ \dfrac{\partial u_i}{\partial t} + u_j \dfrac{\partial u_i}{\partial x_j} = -\dfrac{1}{\rho}\dfrac{\mathrm{d}P}{\mathrm{d}x_1}\delta_{1i} - \dfrac{1}{\rho}\dfrac{\partial p}{\partial x_i} + \nu \dfrac{\partial^2 u_i}{\partial x_j \partial x_j}, \end{cases} \tag{2}$$

where δ_{ij} is the Kronecker delta, ρ is the density, and ν is the kinematic viscosity. Here $(x_1, x_2, x_3) = (x, y, z)$ is the position vector, $(u_1, u_2, u_3) = (u, v, w)$ is the corresponding velocity, and $-\mathrm{d}P/\mathrm{d}x_1$ is the applied mean pressure gradient to drive the flow. We use u_{τ_0} to denote the wall shear velocity of the uncontrolled flow, and we use u_τ to denote the wall shear velocity during the control. Note that here, we use a different term, p, written in lowercase to represent the pressure fluctuation. The friction Reynolds number of the flow is defined as $\mathrm{Re}_\tau = u_{\tau_0}\delta/\nu$, where δ is the channel half height.

3.1.2 The solver

The simulations are performed by discretizing the incompressible Navier–Stokes equations (Equation (2)) with a staggered, second-order-accurate, central finite-difference method in space [34], and an explicit third-order-accurate Runge–Kutta method for time advancement [35]. The system of equations is solved via an operator splitting approach [36]. This code has been validated by prior studies on turbulent channel flow [29, 37, 38]. The computation domain of the simulation is denoted by Ω, and we use Γ^+ and Γ^- to denote two walls located at $y = 0$ and $y = 2\delta$, respectively. Periodic boundary conditions are applied to both the spanwise and streamwise directions, and the no-slip boundary condition in the streamwise and spanwise directions is applied at the walls. A no-penetration boundary condition is used at the walls for the uncontrolled case, whereas blowing and suction boundary conditions are used for controlled cases. The implementation of this solver comes from previous studies [29]. From an RL perspective, the solver is considered as an environment.

3.1.3 Control setups

The active control is achieved by applying a wall-normal velocity at the wall, which can either be blow or suction at walls. We call such a velocity a control or an action, denoted by ϕ. Controls are applied at both walls (as shown on the left of Figure 1), while we focus on one wall in methodology formulation for simplicity. We use variables with subscript w to denote physical variables associated with the wall. For example, we use p_w to denote pressure at the wall. We interchangeably use the terms "control", "action", or "actuation" to denote ϕ in this paper. During the control, the channel flow is driven by a uniform mass fluxr [6, 8, 9] which fixes the bulk Reynolds number $\mathrm{Re}_b = u_b\delta/\nu$, where u_b is the mass flow rate. This is achieved by adapting the mean pressure gradient $\mathrm{d}P/\mathrm{d}x_1$ at each time step [39]. The control target is to reduce the mean pressure gradient to achieve a drag reduction (DR):

$$DR = \frac{(-\frac{\mathrm{d}P}{\mathrm{d}x_1})_{t=0} - (-\frac{\mathrm{d}P}{\mathrm{d}x_1})_{t=T}}{(-\frac{\mathrm{d}P}{\mathrm{d}x_1})_{t=0}}, \tag{3}$$

where $(-\mathrm{d}P/\mathrm{d}x_1)_{t=0}$ denotes the pressure gradient of the uncontrolled flow and $(-\mathrm{d}P/\mathrm{d}x_1)_{t=T}$ denotes the pressure gradient after the control (at the termination timestep $t = T$).

3.1.4 Computation domains

We discretize the computation domain via a staggered central finite-difference method. We consider two different simulation domains [16], where the first one is called a *minimal channel*

	Minimum channel			Full channel		
Re_b	Nx	Ny	Nz	Nx	Ny	Nz
$3k$	32	130	32	128	130	128
$6k$	64	260	64	256	260	256
$9k$	96	390	96	384	390	384
$12k$	160	520	160	512	520	512
$15k$	192	650	192	640	650	640

Table 4: A list of bulk Reynolds numbers Re_b along with corresponding grid resolutions used in our study.

with size $\Omega = L_x \times L_y \times L_z = 1.77\delta \times 2\delta \times 0.89\delta$. This channel flow is large enough to reflect relevant near-wall turbulent statistics and is used to reduce the computational burden. We also adopt another larger channel domain called a *full channel*, which is of size $L_x = 2\pi\delta$, $L_y = 2\delta$, and $L_z = \pi\delta$. For the case with $\mathrm{Re}_\tau \approx 180$, we discretize the streamwise and spanwise directions uniformly using $N_x \times N_z = 32 \times 32$ for the minimum channel and $N_x \times N_z = 128 \times 128$ for the full channel, which result in streamwise and spanwise grid spacing of $\Delta x^+ \approx 10$ and $\Delta z^+ \approx 5$. Here, the superscript + denotes the wall units defined by ν and u_{τ_0}. In the wall-normal direction, we use a hyperbolic-tangent stretching function of size $N_y = 130$, which results in a wall-normal spacing of $\min(\Delta y^+) \approx 0.17$ at the wall and $\max(\Delta y^+) \approx 7.6$ at the center of the channel domain.

3.1.5 The Reynolds numbers

For the remainder of the paper, we use Re to denote Re_b. We experiment on different Reynolds numbers in the numerical study to test the generalization ability of the concerned methods. When changing the Reynolds numbers, we also proportionally scale the grid resolution N_x, N_y, and N_z to preserve the grid resolutions in wall units. We give specific configurations of those parameters in Table 4.

3.2 Algorithm outline of PINO-PC

In this paper, we propose a machine-learning-based framework for flow control inspired by previous predictive control (PC) studies [7, 8, 40]. The overall schematic is shown on the right of Figure 1. The proposed framework adopts a policy model denoted as M_p, and an observer model denoted as M_O. The policy model M_p predicts the control ϕ based on the boundary pressure p_w. The observer model M_O predicts the interior velocity field u based on the control ϕ. The velocity field u is used to predict the outcome (the reward) of the control ϕ to revise the control accordingly. For convenience, we focus on the top wall in the whole section, while the control also happens in the bottom wall in practice.

We propose a machine learning algorithm to jointly optimize the policy model and the observer model in Algorithm 1. We use a memory (a "replay buffer" [27]) to store collected data, which is used to optimize those models. Initially, the policy model and the observer model are initialized with zero weights, and the memory is initialized as an empty collection. The main loop of the algorithm iterates through episodes and conducts data collection, control, and learning during each episode: At the start of one episode (Line 4), PINO-PC collects the

wall pressure p_w, the field velocity u and the drag $-\frac{dP}{dx_1}$ from the solver, and stores them to the memory. Next, the policy model predicts the control ϕ based on the boundary pressure (Line 5). Then, the control ϕ is applied to the wall, and the solver is updated to the next timestep (Line 6). Subsequently, the algorithm optimizes the observer model and the policy model based on sampled data (Lines 8 – 12). For each training epoch, a data tuple is sampled from memory (Line 9). Then, the observer model is learned through optimizing two loss functions (introduced later): supervised loss L_{data} (Equation (13)) and physics loss L_{pde} (Equation (14)). After that, the policy model is trained to optimize a policy loss L_{policy} (Line 11). Here, the observer model is used to predict the outcome of the policy model, so when training the policy model, the observer model is fixed, which means the weights of the observer model are not changed during the optimization of the policy model.

We build the policy model and the observer model based on neural operators [26]. Neural operators are neural networks that learn the map between infinite-dimensional function spaces, such as the Fourier neural operator (FNO [31]). Prior study [26] shows that neural operators demonstrate stronger performance than existing machine-learning-based methodologies. We again verify that the neural-operator-based model outperforms prior state-of-the-art machine-learning-based approaches in our experiments.

Algorithm 1 Physics-informed neural-operator-based predictive control (PINO-PC)

1: Initialize the policy model, the observer model, and the memory;
2: **for** each episode **do**
3: **for** each timestep in episode **do** ▷ Data collection: roll out and collect trajectory
4: Collect p_w, $\boldsymbol{u}$, and $-\frac{dP}{dx_1}$ to the memory;
5: Predict the control ϕ with the policy model M_p;
6: Apply the control ϕ to the wall, and step the solver to the next timestep;
7: **end for**
8: **for** each epoch **do**
9: Sample a data tuple $(p_w, u, -\frac{dP}{dx_1})$ from the memory;
10: Update the observer model based on L_{data} and L_{pde};
11: Fix the observer model, update the policy model to optimize the target L_{policy};
12: **end for**
13: **end for**

3.3 The policy model M_p

We introduce a policy model M_p to predict the control ϕ based on the boundary pressure p_w. The structure of the policy model is shown on the left of Figure 6.

3.3.1 The FNO encoder of the policy model

First, the policy model leverages an FNO [31], which is a variant of neural operators [26]. Given the boundary pressure $p_w : \Gamma^+ \to \mathbb{R}$, we use an FNO encoder to encode it to the latent function h_p:

$$h_p = \text{FNO}(p_w). \tag{4}$$

3.3.2 Conditioning on the Reynolds number

To help the model adapt to unseen Reynolds numbers, we also condition the decoder model on the Reynolds number Re. We propose to use the multiplicative filter (MFN) [32] to encode the Reynolds number Re to get latent representations. The multiplicative filter takes the hidden feature and the Reynolds number Re as input and outputs a new feature h_{pm}:

$$h_{pm} = \mathrm{MFN}(h_p), \tag{5}$$

where the MFN is based on the sinusoidal filter g to encode the Reynolds number Re:

$$g\left(\mathrm{Re}; \theta^{(i)}\right) = \sin\left(\omega^{(i)} \frac{\mathrm{Re}}{\mathrm{Re_m}} + \tau^{(i)}\right). \tag{6}$$

Here $\theta^{(i)} = \{\omega^{(i)}, \tau^{(i)}\}$ are parameters of the sinusoidal filter, and a constant Reynolds number $\mathrm{Re_m} = 100,000$ is used to normalize the input Reynolds number Re. The MFN then performs the following recursion with L layers:

$$\begin{aligned} z^{(1)} &= h_p, \\ z^{(i+1)} &= \left(W^{(i)} z^{(i)} + b^{(i)}\right) \circ g\left(\mathrm{Re}; \theta^{(i+1)}\right), i = 1, \ldots, L-1, \\ h_{pm} &= W^{(L)} z^{(L)} + b^{(L)}, \end{aligned} \tag{7}$$

where $W^{(i)}$ and $b^{(i)}$ are the learnable linear transform and bias.

3.3.3 The FNO decoder of the policy model

Finally, the policy model uses an FNO decoder to get the control:

$$\phi = M_p(p_w) = \mathrm{FNO}(h_{pm}) - \mathrm{mean}(\mathrm{FNO}(h_{pm})), \tag{8}$$

where we use a normalization function to ensure that we don't add mass to the system.

3.3.4 The policy loss

To optimize the policy model, we adopt a policy loss (corresponding to Line 11 in Algorithm 1). The policy loss is written in two terms: the turbulent kinetic energy (TKE) and the norm of the control:

$$L_{policy}(\phi) = E_t\left(\int_\Omega |\boldsymbol{u}(t+\Delta t)|^2 \, \mathrm{d}\boldsymbol{x} + \frac{\lambda_n}{\Delta t} \int_{\Gamma_2^+} \int_t^{t+\Delta t} \phi^2 \, \mathrm{d}\tau \, \mathrm{d}S\right). \tag{9}$$

In this equation, the expectation is taken over the concerned episode. The TKE is computed based on the field velocity, where the field velocity is the predicted outcome of the control with the fixed observer model (this prediction procedure will be introduced in the next subsection). We use $\boldsymbol{u}(t+\Delta t)$ to denote the velocity at the time after a period $t+\Delta t$, which helps obtain

long-term gain. Here $\lambda_n = 0.5$ is a balancing term of the regularization term, and Δt is the concerned time window.

3.4 The observer model

In this subsection, we present a PDE observer model called M_O, which predicts internal velocity field $\boldsymbol{u}$ given the control ϕ. We use ϕ_t to denote the control at discrete timestep t. The observer model is also conditioned on the Reynolds number Re to boost generalization to different Reynolds numbers. The observer model is shown on the right of Figure 6.

3.4.1 The FNO encoder of the observer model

The boundary velocity $\{\phi\}$ is normalized and passed through an FNO encoder. The FNO encoder takes the boundary velocity as input and outputs a hidden feature of controls $h_c : \Gamma^+ \to \mathbb{R}^{d_1}$ where d_1 is the hidden feature dimension.

3.4.2 The field decoder of the observer model

We then use a field decoder to transform the hidden feature h_c to the field velocity $\boldsymbol{u} : \Omega \to \mathbb{R}$, $\boldsymbol{v} : \Omega \to \mathbb{R}$, and $\boldsymbol{w} : \Omega \to \mathbb{R}$. To achieve this, we first generate latent representations for each of the field velocities, which is performed by an inflating hidden function $h_{in} : \Omega \to \mathbb{R}^{d_2}$, ($d_2$ is the dimension of the inflated feature) turning the 2D hidden feature h_c to the 3D space:

$$h_{in}(x, y, z) = h_c(x, z) \oplus \text{PosEmb}(y), \tag{10}$$

where $\oplus$ denotes the concatenation operator, and PosEmb is a positional embedding function which turns y into a hidden feature:

$$\gamma(y, j) = \begin{cases} \sin(2^{\lfloor j/2 \rfloor} \pi y), & \text{if } j \bmod 2 = 0, \\ \cos(2^{\lfloor j/2 \rfloor} \pi y), & \text{else.} \end{cases} \tag{11a}$$

$$\text{PosEmb}(y) = \gamma(y, 1) \oplus \gamma(y, 2) \oplus \ldots \oplus \gamma(y, n_p), \tag{11b}$$

where n_p is the number of the trigonometric functions adopted to form the positional embeddings. After that, we decode the hidden functions into u, v, w with 3D FNO modules [31]:

$$u = \text{FNO3D}(\text{h}_{\text{in}}), v = \text{FNO3D}(\text{h}_{\text{in}}), w = \text{FNO3D}(\text{h}_{\text{in}}). \tag{12}$$

3.4.3 The data loss

We introduce the data loss to train the observer model, which penalizes the $L2$ distance between the predicted and the ground-true field velocity, denoted by u and u_{gt} correspondingly:

$$L_{data} = \mathbb{E}_{x,y,z} \left(\frac{u_{gt}(x, y, z) - u(x, y, z)}{\bar{u}(x, y, z)} \right)^2, \tag{13}$$

where $\bar{u}$ is the root-mean-square velocity at each point (x, y, z), and the expectation is taken in the full domain Ω.

3.4.4 Physics-informed learning and the PDE loss

Despite the supervised loss introduced in Equation (13), we further introduce a PDE loss L_{pde} to optimize the model by leveraging the governing PDE (Equation (2)). This approach belongs to physics-informed learning [24, 41], where through optimizing the PDE loss, the observer can be optimized without ground-true data acquired from precise simulation. This technique is helpful when the training data is scarce, such as in the high Reynolds number region. Our experiments show that physics-informed learning can help the observer model generalize to unseen Reynolds numbers.

We implement the PDE loss based on the difference between the temporal gradient of the predicted velocity and the right-hand-side terms. We denote the predicted velocity via $\frac{du_\theta}{dt}$, $\frac{dv_\theta}{dt}$, and $\frac{dw_\theta}{dt}$, and we denote rest terms via a function R, then the PDE loss is given as:

$$L_{pde} = \left|\frac{du_\theta}{dt} - R(u_\theta)\right| + \left|\frac{dv_\theta}{dt} - R(v_\theta)\right| + \left|\frac{dw_\theta}{dt} - R(w_\theta)\right|, \tag{14}$$

where the velocity gradients are estimated via temporal difference, and we compute $R(u_i)$ as:

$$R(u_i) = -u_j \frac{\partial u_i}{\partial x_j} - \frac{\mathrm{d}P}{\mathrm{d}x_1}\delta_{1i} - \frac{\partial p}{\partial x_i} + \frac{1}{\mathrm{Re}_\tau}\frac{\partial^2 u_i}{\partial x_j \partial x_j}. \tag{15}$$

Here, u_i denotes any of the predicted velocity u_θ, v_θ, and w_θ.

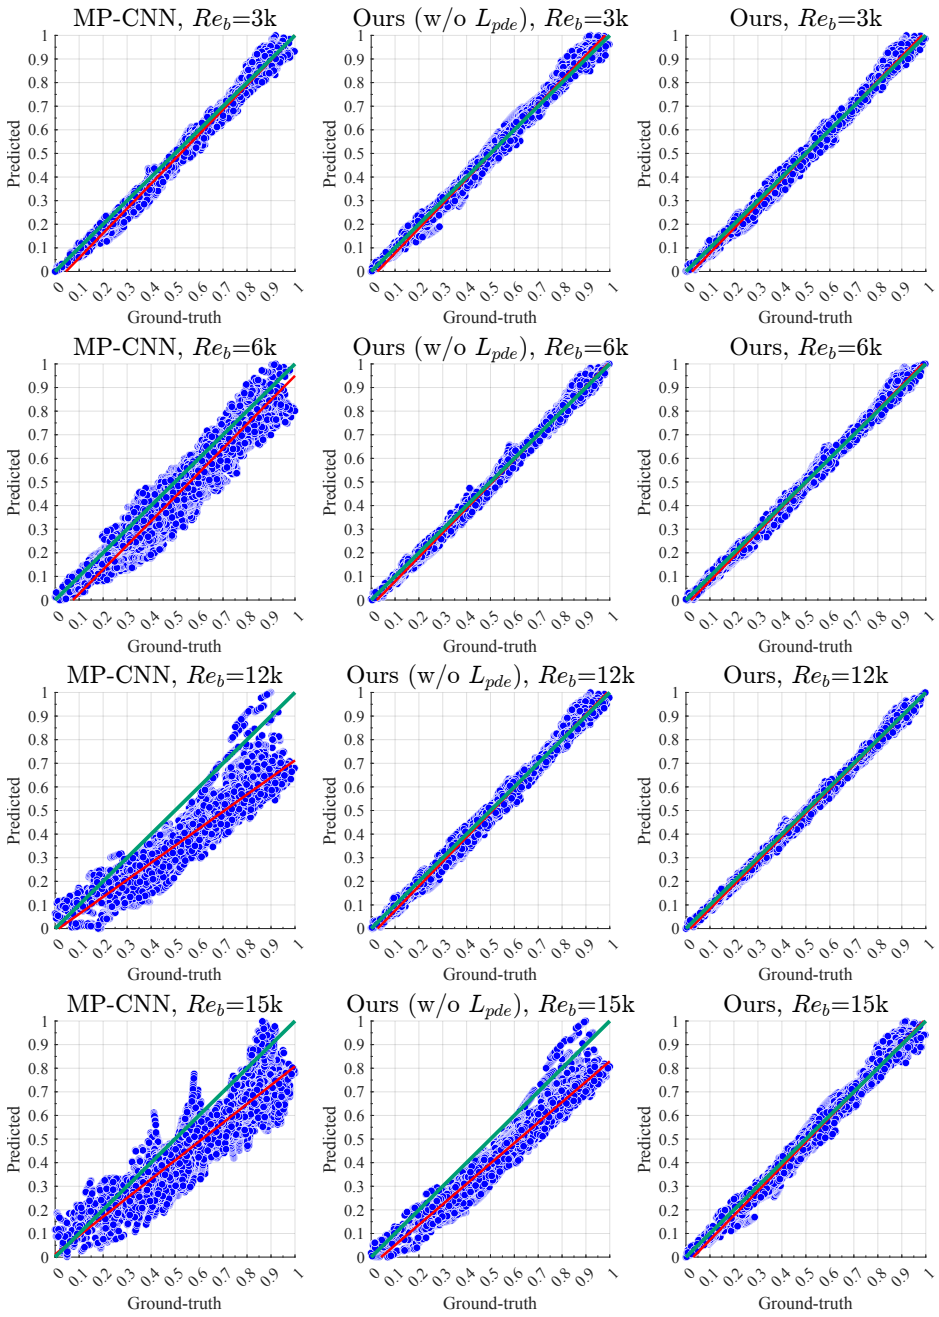

Fig. 5: Scatter plot of predicted velocities (y-axis) and the gt velocities (x-axis) under different Reynolds numbers. We normalize the velocities into the [0, 1] range before plotting. We plot the $y = x$ line in green and the fitting line in red. The more accurate a model is, the closer the plotted points should surround the $y = x$ line.

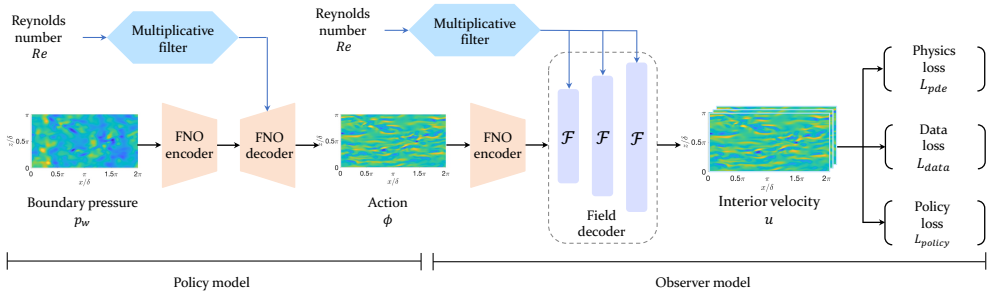

Fig. 6: The policy model and the observer model in PINO-PC. Both the policy model and the observer model are conditioned on the Reynolds number Re, encoded by multiplicative filters [32]. The function of the policy model is to give the boundary velocity ϕ based on the boundary pressure p_w. It is instantiated with an FNO encoder and an FNO decoder. The output boundary velocity ϕ is sent to the observer model, which leverages the FNO encoder to encode the boundary velocity. Then, the observer model uses a field decoder to output interior velocity u. We use a data loss L_{data} (Equation (13)) and a physics loss L_{pde} (Equation (14)) to optimize the observer. Meanwhile, we adopt a policy loss L_{policy} (Equation (9)) to optimize the policy model.

Method name	Opposition	DNS-PC	Local suboptimal	MP-CNN	DDPG	PINO-PC
Reference	[6]	[8]	[7]	[9]	[16]	Ours
1. Machine learning model backbone	N/A	N/A	N/A	CNN	FCN	Neural operators
2. Based on predictive control	✗	✔	✔	✗	✗	✔
3. Only need boundary observation	✗	✗	✔	✔	✔	✔
4. Experiment with varied Reynolds numbers	✗	✗	✗	✔	✗	✔
5. Use a PDE observer	N/A	N/A	N/A	✔	✗	✔
6. Use physics-informed learning	N/A	N/A	N/A	✗	✗	✔

Table 5: In the context of the flow control problem, we present comparisons between previous and our approaches. Each column corresponds to a specific method, and each row denotes a particular property. The first property indicates the machine learning model used in those approaches, with the first three methods not employing machine learning techniques. The second property pertains to whether a method is grounded in predictive control. PINO-PC is rooted in predictive control as it performs control based on the predicted impact of the boundary velocity. The subsequent property addresses whether the method can function solely with boundary information, excluding the need for internal field data. All three machine learning models listed can achieve this, except for Local suboptimal [7]. The fourth property outlines whether a method has been verified in flows of varied Reynolds numbers. The final two properties are techniques used in machine learning methods. Both MP-CNN [9] and our approach leverage a PDE observer. Furthermore, our model is the only approach that employs physics-informed learning in the control procedure.

4 Appendix

4.1 Comparisons with other methods

A comparative analysis of our approach and related methodologies is provided in Table 5, affording a comprehensive understanding of PINO-PC's features.

4.2 Theorectical results of DDPG

DDPG models the control problem as a Markov decision process (MDP) in function space. The MDP consists of several components: a state space $\mathcal{X}$, an action space $\mathcal{A}$. A state $x \in \mathcal{X}$ and an action $a \in \mathcal{A}$ are functions. We use $p_1(x)$ to denote the initial probabilistic measure over states, and $p(x_{t+1}|x_t)$ to denote the probabilistic measure describing the transition dynamics distribution, while the MDP satisfies the Markov property $p(x_{t+1}|x_1,...,x_t,a_t) = p(x_{t+1}|x_t,a_t)$, for any trajectory $x_1,a_1,x_2,a_2,...,x_T,a_T$ in state-action space, and a reward function $r : \mathcal{X} \times \mathcal{A} \to \mathbb{R}$. We define a deterministic policy $\mu_\theta : \mathcal{X} \to \mathcal{A}$ parameterized by θ. The discount factor $\gamma \in [0,1)$ is given to calculate the total discounted reward (i.e., the return) r_t^γ. We denote the Q function to be $Q^\mu(x,a) = \mathbb{E}[r_1^\gamma | X_1 = x, A_1 = a; \mu]$, and we denote the V function as $V^\mu(x) = \mathbb{E}[r_1^\gamma | X_1 = a; \mu]$. For simplicity, we superscript value functions by μ instead of μ_θ. We denote the Fréchet derivative of the state-action value function with respect to the action to be $D_a(Q^\mu(x,a))$, and we denote the derivative of the transition probability to be $D_a(p(x_{t+1}|x_t,a))$. We further denote the discounted state occupancy measure by $\rho^\mu(x') = \int_{\mathcal{X}} \sum_{t=1}^{\infty} \gamma^{t-1} p_1(x) p(x \to x', t, \mu) \, \mathrm{d}x$, and we consider the following performance

objective:

$$J(\mu_\theta) = \mathbb{E}_{x\sim\rho^\mu}[r(x,\mu_\theta(x)] = \int_{\mathcal{X}} V^\mu(x)\mathrm{d}p_1(x). \tag{16}$$

Then, we introduce the deterministic policy gradient theorem for neural operators as follows:

Theorem 1 (Deterministic policy gradient theorem for neural operators)**.** *Suppose that the MDP satisfies the following regularization conditions*

1. *$p(x'|x,a)$, $D_a(Q^\mu(x,a))$, $\mu_\theta(x)$, $r(x,a)$, $D_a r(x,a)$, and $p_1(x)$ are continuous in all parameters and variables x, a, x' and θ,*
2. *there exists b and L such that $\sup_s p_1(x) < b$, $\sup_{a,x,x'} p(x'|x,a) < b$, $\sup_{a,x} \|r(x,a)\| < b$, $\sup_{a,x,x'} \|D_a p(x'|x,a)\| < L$, and $\sup_{a,x} \|D_a r(x,a)\| < L$,*

then $\nabla_\theta \mu_\theta(x)$ and $D_a(Q^\mu(x,a))$ exist and the deterministic policy gradient is given as:

$$\begin{aligned}\nabla_\theta J(\mu_\theta) &= \int_{\mathcal{X}} \rho^\mu(x)\nabla_\theta\mu_\theta(x) D_a(Q^\mu(x,a))\Big|_{a=\mu_\theta(x)}\mathrm{d}x \\ &= \mathbb{E}_{x\sim\rho^\mu}\left[\nabla_\theta\mu_\theta(x) D_a(Q^\mu(x,a))|_{a=\mu_\theta(x)}\right].\end{aligned}$$

Proof. The proof follows the deterministic policy gradient algorithms [18]. We first derive the gradient of the value function as:

$$\nabla_\theta V^\mu(x) = \nabla_\theta Q^\mu(x,\mu_\theta(x)) \tag{17}$$

$$= \nabla_\theta\left(r(x,\mu_\theta(x)) + \int_{\mathcal{X}} \gamma p(x' \mid x,\mu_\theta(x)) V^\mu(x')\mathrm{d}x'\right), \tag{18}$$

$$\begin{aligned}&= \nabla_\theta\mu_\theta(x) D_a r(x,a)|_{a=\mu_\theta(x)} + \\ &+ \int_{\mathcal{X}} \gamma\left(p(x' \mid x,\mu_\theta(x))\nabla_\theta V^\mu(x') + \nabla_\theta\mu_\theta(x) D_a p(x' \mid x,a)|_{a=\mu_\theta(x)} V^\mu(x')\right)\mathrm{d}x'\end{aligned} \tag{19}$$

$$\begin{aligned}&= \nabla_\theta\mu_\theta(x) D_a\left(r(x,a) + \int_{\mathcal{X}} \gamma p(x' \mid x,a) V^\mu(x')\mathrm{d}x'\right)\Big|_{a=\mu_\theta(x)} \\ &+ \int_{\mathcal{X}} \gamma p(x' \mid x,\mu_\theta(x))\nabla_\theta V^\mu(x')\mathrm{d}x'\end{aligned} \tag{20}$$

$$= \nabla_\theta\mu_\theta(x) D_a Q^\mu(x,a)|_{a=\mu_\theta(x)} + \int_{\mathcal{X}} \gamma p(x\to x',1,\mu_\theta)\nabla_\theta V^\mu(x')\mathrm{d}x' \tag{21}$$

$$= \int_{\mathcal{X}} \sum_{t=0}^{\infty} \gamma^t p(x\to x',t,\mu_\theta)\nabla_\theta\mu_\theta(x') D_a Q^\mu(x',a)\Big|_{a=\mu_\theta(x')}\mathrm{d}x', \tag{22}$$

where we apply the definition of the value function in Equation (17) and Equation (18). We apply the chain rule for Fréchet derivative in Equation (19). We use the Leibniz integral rule to exchange the order of derivative and integration based on the regularization conditions in Equation (21) and Equation (22). Furthermore, Equation (22) is derived by iterating Equation (21) based on the formula of the value derivative $\nabla_\theta V^\mu$.

Now consider the definition of the target in Equation (16), we can derive that:

$$\nabla_\theta J(\mu_\theta) = \nabla_\theta \int_{\mathcal{X}} p_1(x) V^\mu(x) dx \tag{23}$$

$$= \int_{\mathcal{X}} p_1(x) \nabla_\theta V^\mu(x) dx \tag{24}$$

$$= \int_{\mathcal{X}} \int_{\mathcal{X}} \sum_{t=0}^{\infty} \gamma^t p_1(x) p(x \to x', t, \mu_\theta) \nabla_\theta \mu_\theta(x') D_a Q^\mu(x', a) \Big|_{a=\mu_\theta(x')} dx' dx \tag{25}$$

$$= \int_{\mathcal{X}} \rho^\mu(x) \nabla_\theta \mu_\theta(x) D_a Q^\mu(x, a) \Big|_{a=\mu_\theta(x)} dx, \tag{26}$$

where we leverage again exchange the order of integration and derivative in Equation (25) and Equation (26), and we consider the definition of ρ^μ in the last line. □

4.3 Ablation studies of PINO-PC

In this subsection, we compare various machine-learning-based models and, therefore, provide rationals and insights behind our model choice. First, we validate whether the proposed policy learning framework is superior to existing model-free reinforcement learning methods. Furthermore, we explore whether neural operators [26] can bring benefits over convolutional neural networks (CNNs). Finally, we investigate the effectiveness of the chosen machine learning modules such as multiplicative filters. The result of this ablation study is presented in Table 3.

4.3.1 Analysis via the supervised representation learning

In this subsection, we collect several datasets to train and evaluate observer models, while the target is to measure the fitting performance of each model. All collected datasets are using the full channel flow. We change the Reynolds number by altering the kinematic viscosity, as stated in Section 3.1. Those datasets are obtained from the DNS of a turbulent channel flow, with a bulk velocity Reynolds number of $3k$, $6k$, $9k$, $12k$, and $15k$ correspondingly. We experiment with four setups: Single, Interpolation, Extrapolation-A, and Extrapolation-B. Descriptions of those datasets are presented in the main text. Different from other experiments that do not use the normal velocity $v(y^+ = 10)$ as a supervision signal, in this section, all models are trained and compared under a supervised learning setup where the normal velocity is the ground truth. Each setup has flow data of three splits: `train`, `validation`, and `test`. The `training` split is used to optimize machine-learning models, the `validation` split is used to tune hyper-parameters of models, and the `test` split is for evaluating machine-learning models. We use 700 instantaneous fields for the `training` split, 100 fields for the `validation` dataset, and 100 different fields for the `test` split. If a split contains more than one Reynolds number, then the number of flow data is equal for each Reynolds number. The training dataset size is approximately the same as Park and Choi [9] and is enough to train neural models. The collected wall pressure and normal velocity are normalized by their root-mean-square values before being fed into neural models.

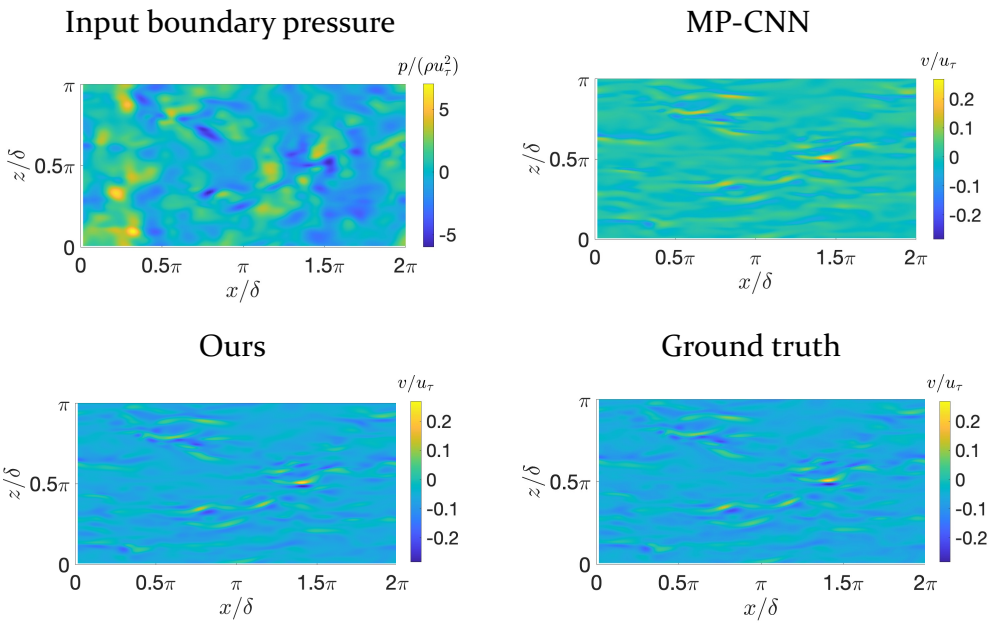

Fig. 7: Visualization of input pressure and predicted interior velocity $v_{y^+=10}$ in the full channel scenario.

We first compare the representation power of PINO-PC's observer model against another observer of MP-CNN [9] under a supervised learning setting, where we call the observer model used in PINO-PC as M_O. In this supervised learning setup, all models take the wall pressure p_w as input, predicting the normal velocity $v(y^+ = 10)$ at the detection layer $y^+ = 10$. Trained models under this supervised learning framework can be further applied to opposition control [9], where details will be given later in this subsection.

We compare to MP-CNN [9], which contains twenty hidden layers, an averaging pooling layer, and a linear layer with residual connections. Zero paddings are used to adjust the sizes of convolutional filters. Zero paddings are applied when the height or width in the input is an odd number. The input and output grid points are 32×32 and 16×16, corresponding to the x-axis and z-axis. The input and output are aligned in their centers. The weights in the model are initialized by the Xavier method [42]. We do not use the GAN loss [43] to train the MP-CNN [9] because computing the GAN loss needs another CNN as a discriminator. We experiment with two loss setups to test the performance of MP-CNN [9]. The first setup is to only train the model with L_{data} (w/o physics-informed learning), and the other setup is to train the model with L_{data} and L_{pde} (w/ physics-informed learning). The model parameters are optimized with Adam [44] with an initial learning rate of $1e-3$. No learning rate scheduler is adopted.

In this case, we let the observer model only output the interior velocity in the detection plane $v(y^+ = 10)$. We use the same size of inputs and outputs as MP-CNN [9]. The number of parameters in our observer model is smaller than that of MP-CNN. The optimizer and learning rate remains the same as that of MP-CNN.

Figure 5 shows scatter plots of the prediction and ground truth data, where these diagrams are produced under the Extrapolation-B setup. Under this setup, the training dataset uses a Reynolds number of $3k$ (corresponding to the first row of this figure), while testing datasets use varied Reynolds numbers of $6k$ to $15k$. We experiment with two different models by changing training losses. The first setting is to train M_O only with the data loss L_{data} (denoted by "Ours w/o L_{pde}"), while the physics-informed learning is not adopted in this case. The second setting is to train our model with the data loss L_{data} and the PDE loss L_{pde}. This full model is denoted by "Ours". In this plot, the x-axis denotes the ground truth, and the y-axis denotes the predicted values, where the velocity tensor is flattened into $1D$ before plotting this diagram. Therefore, the distance of each scatter point to the $y = x$ line (which means ground truth equals the prediction) reflects the prediction error. We observe that MP-CNN [9] performs well in a bulk-velocity Reynolds number $\mathrm{Re}_b = 3k$, but it deteriorates significantly in the high Reynolds number scenarios. M_O demonstrates superior performance than MP-CNN [9], even without physics-informed learning (the PDE loss L_{pde}). The PDE loss can enhance predictions of M_O, especially when the flow is highly turbulence. This demonstrates that M_O learns better neural features than MP-CNN [9].

Figure 7 provides 2D visualizations of input pressure and output velocity comparing ours against the MP-CNN [9] baseline. All models are trained only with L_{data} under a Reynolds number of $3k$ and are tested in an unseen Reynolds number of $6k$. In this figure, predicted velocities are on the test split. Our methods can produce closer predictions to the ground truth, while MP-CNN fails to predict the target velocity in many regions.

References

[1] Dutton, J.A., Panofsky, H.A.: Clear air turbulence: A mystery may be unfolding: High altitude turbulence poses serious problems for aviation and atmospheric science. Science **167**(3920), 937–944 (1970)

[2] McWilliams, J.C., Sullivan, P.P., Moeng, C.-H.: Langmuir turbulence in the ocean. Journal of Fluid Mechanics **334**, 1–30 (1997) https://doi.org/10.1017/S0022112096004375

[3] Ghalichi, F., Deng, X., De Champlain, A., Douville, Y., King, M., Guidoin, R.: Low Reynolds number turbulence modeling of blood flow in arterial stenoses. Biorheology **35**(4-5), 281–294 (1998)

[4] Stein, P.D., Sabbah, H.N.: Measured turbulence and its effect on thrombus formation. Circulation Research **35**(4), 608–614 (1974)

[5] Brunton, S.L., Noack, B.R.: Closed-loop turbulence control: Progress and challenges. Applied Mechanics Reviews **67**(5), 050801 (2015) https://doi.org/10.1115/1.4031175

[6] Choi, H., Moin, P., Kim, J.: Active turbulence control for drag reduction in wall-bounded flows. Journal of Fluid Mechanics **262**, 75–110 (1994)

[7] Lee, C., Kim, J., Choi, H.: Suboptimal control of turbulent channel flow for drag reduction. Journal of Fluid Mechanics **358**, 245–258 (1998)

[8] Bewley, T.R., Moin, P., Temam, R.: DNS-based predictive control of turbulence: an optimal benchmark for feedback algorithms. Journal of Fluid Mechanics **447**, 179–225 (2001)

[9] Park, J., Choi, H.: Machine-learning-based feedback control for drag reduction in a turbulent channel flow. Journal of Fluid Mechanics **904**, 24 (2020)

[10] Guastoni, L., Güemes, A., Ianiro, A., Discetti, S., Schlatter, P., Azizpour, H., Vinuesa, R.: Convolutional-network models to predict wall-bounded turbulence from wall quantities. Journal of Fluid Mechanics **928**, 27 (2021)

[11] Ho, C.-M., Huang, L.-S.: Subharmonics and vortex merging in mixing layers. Journal of Fluid Mechanics **119**, 443–473 (1982)

[12] Mettot, C., Sipp, D., Bézard, H.: Quasi-laminar stability and sensitivity analyses for turbulent flows: prediction of low-frequency unsteadiness and passive control. Physics of Fluids **26**(4) (2014)

[13] Blackburn, H.M., Mansour, N.N., Cantwell, B.J.: Topology of fine-scale motions in turbulent channel flow. Journal of Fluid Mechanics **310**, 269–292 (1996)

[14] Hammond, E.P., Bewley, T.R., Moin, P.: Observed mechanisms for turbulence attenuation and enhancement in opposition-controlled wall-bounded flows. Physics of Fluids **10**(9), 2421–2423 (1998)

[15] Jeong, J., Hussain, F.: On the identification of a vortex. Journal of Fluid Mechanics **285**, 69–94 (1995) https://doi.org/10.1017/S0022112095000462

[16] Guastoni, L., Rabault, J., Schlatter, P., Azizpour, H., Vinuesa, R.: Deep reinforcement learning for turbulent drag reduction in channel flows. The European Physical Journal E **46**(4), 27 (2023)

[17] Farahmand, A.-m., Nabi, S., Nikovski, D.N.: Deep reinforcement learning for partial differential equation control. In: 2017 American Control Conference (ACC), pp. 3120–3127 (2017). https://doi.org/10.23919/ACC.2017.7963427

[18] Lillicrap, T.P., Hunt, J.J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., Wierstra, D.: Continuous control with deep reinforcement learning. arXiv preprint arXiv:1509.02971 (2015)

[19] Rabault, J., Kuchta, M., Jensen, A., Réglade, U., Cerardi, N.: Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control. Journal of Fluid Mechanics **865**, 281–302 (2019)

[20] Fan, D., Yang, L., Wang, Z., Triantafyllou, M.S., Karniadakis, G.E.: Reinforcement learning for bluff body active flow control in experiments and simulations. Proceedings of the National Academy of Sciences **117**(42), 26091–26098 (2020)

[21] Tang, H., Rabault, J., Kuhnle, A., Wang, Y., Wang, T.: 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) (2020)

[22] Chatzimanolakis, M., Weber, P., Koumoutsakos, P.: Learning in two dimensions and controlling in three: Generalizable drag reduction strategies for flows past circular cylinders through deep reinforcement learning. Phys. Rev. Fluids **9**, 043902 (2024) https://doi.org/10.1103/PhysRevFluids.9.043902

[23] Lale, S., Azizzadenesheli, K., Hassibi, B., Anandkumar, A.: Model learning predictive control in nonlinear dynamical systems. In: 2021 60th IEEE Conference on Decision and Control (CDC), pp. 757–762 (2021). IEEE

[24] Li, Z., Zheng, H., Kovachki, N., Jin, D., Chen, H., Liu, B., Azizzadenesheli, K., Anandkumar, A.: Physics-informed neural operator for learning partial differential equations. arXiv preprint arXiv:2111.03794 (2021)

[25] Wu, J.-L., Xiao, H., Paterson, E.: Physics-informed machine learning approach for augmenting turbulence models: A comprehensive framework. Phys. Rev. Fluids **3**, 074602 (2018) https://doi.org/10.1103/PhysRevFluids.3.074602

[26] Kovachki, N.B., Li, Z., Liu, B., Azizzadenesheli, K., Bhattacharya, K., Stuart, A.M., Anandkumar, A.: Neural operator: Learning maps between function spaces. CoRR **abs/2108.08481** (2021) 2108.08481

[27] Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., *et al.*: Human-level control through deep reinforcement learning. Nature **518**(7540), 529–533 (2015)

[28] Kaiser, L., Babaeizadeh, M., Milos, P., Osinski, B., Campbell, R.H., Czechowski, K., Erhan, D., Finn, C., Kozakowski, P., Levine, S., et al.: Model-based reinforcement learning for Atari. arXiv preprint arXiv:1903.00374 (2019)

[29] Bae, H.J., Lozano-Duran, A., McKeon, B.J.: Nonlinear mechanism of the self-sustaining process in the buffer and logarithmic layer of wall-bounded flows. Journal of Fluid Mechanics **914**, 3 (2021)

[30] Liu-Schiaffini, M., Singer, C.E., Kovachki, N., Schneider, T., Azizzadenesheli, K., Anandkumar, A.: Tipping Point Forecasting in Non-Stationary Dynamics on Function Spaces (2023)

[31] Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., Anandkumar, A.: Fourier neural operator for parametric partial differential equations. arXiv preprint arXiv:2010.08895 (2020)

[32] Fathony, R., Sahu, A.K., Willmott, D., Kolter, J.Z.: Multiplicative filter networks. In: International Conference on Learning Representations (2020)

[33] Hunt, J.C., Wray, A.A., Moin, P.: Eddies, streams, and convergence zones in turbulent flows. Center for Turbulence Research, Proceedings of the Summer Program (1988)

[34] Orlandi, P.: Fluid Flow Phenomena: a Numerical Toolkit vol. 55. Springer, ??? (2000)

[35] Wray, A.A.: Minimal storage time advancement schemes for spectral methods. NASA Ames Research Center, California, Report No. MS **202** (1990)

[36] Chorin, A.J.: Numerical solution of the Navier–Stokes equations. Mathematics of Computation **22**(104), 745–762 (1968)

[37] Bae, H.J., Lozano-Durán, A., Bose, S.T., Moin, P.: Turbulence intensities in large-eddy simulation of wall-bounded flows. Phys. Rev. Fluids **3**, 014610 (2018) https://doi.org/10.1103/PhysRevFluids.3.014610

[38] Bae, H.J., Lozano-Durán, A., Bose, S.T., Moin, P.: Dynamic slip wall model for large-eddy simulation. Journal of Fluid Mechanics **859**, 400–432 (2019) https://doi.org/10.1017/jfm.2018.838

[39] Gokarn, A., Battaglia, F., Fox, R., Hill, J., Reveillon, J.: Large eddy simulations of incompressible turbulent flows using parallel computing techniques. International Journal for Numerical Methods in Fluids **56**(10), 1819–1843 (2008)

[40] Maciejowski, J.M., Huzmezan, M.: Predictive control. In: Robust Flight Control: A Design Challenge, pp. 125–134. Springer, ??? (2007)

[41] Raissi, M., Perdikaris, P., Karniadakis, G.E.: Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics **378**, 686–707 (2019)

[42] Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, pp. 249–256 (2010). JMLR Workshop and Conference Proceedings

[43] Güemes, A., Discetti, S., Ianiro, A., Sirmacek, B., Azizpour, H., Vinuesa, R.: From coarse wall measurements to turbulent velocity fields through deep learning. Physics of Fluids **33**(7) (2021)

[44] Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)