




Reinforcement-learning-based control of turbulent channel flows at high Reynolds numbers

Zisong Zhou¹ , Mengqi Zhang²  and Xiaojue Zhu¹ 

¹Max Planck Institute for Solar System Research, Göttingen 37077, Germany

²Department of Mechanical Engineering, National University of Singapore, 9 Engineering Drive 1, Singapore 117575, Republic of Singapore

Corresponding author: Xiaojue Zhu, zhux@mps.mpg.de

(Received 21 August 2024; revised 22 November 2024; accepted 27 December 2024)

Deep reinforcement learning (DRL) is employed to develop control strategies for drag reduction in direct numerical simulations of turbulent channel flows at high Reynolds numbers. The DRL agent uses near-wall streamwise velocity fluctuations as input to modulate wall blowing and suction velocities. These DRL-based strategies achieve significant drag reduction, with maximum rates 35.6 % at $Re_\tau \approx 180$, 30.4 % at $Re_\tau \approx 550$, and 27.7 % at $Re_\tau \approx 1000$, outperforming traditional opposition control methods. An expanded range of wall actions further enhances drag reduction, although effectiveness decreases at higher Reynolds numbers. The DRL models elevate the virtual wall through blowing and suction, aiding in drag reduction. However, at higher Reynolds numbers, the amplitude modulation of large-scale structures significantly increases the residual Reynolds stress on the virtual wall, diminishing the drag reduction. Analysis of budget equations provides a systematic understanding of the underlying drag reduction dynamics. The DRL models reduce skin friction by inhibiting the redistribution of wall-normal turbulent kinetic energy. This further suppresses the wall-normal velocity fluctuations, reducing the production of Reynolds stress, thereby decreasing skin friction. This study showcases the successful application of DRL in turbulence control at high Reynolds numbers, and elucidates the nonlinear control mechanisms underlying the observed drag reduction.

Key words: drag reduction, turbulence control, turbulence simulation

1. Introduction

Turbulent flows lead to significantly greater energy losses compared to laminar flows, presenting a major challenge in various engineering applications (Brunton & Noack 2015). For instance, wall friction contributes to approximately 50 % of total resistance in aircraft, up to 90 % in submarines, and nearly all resistance in pipeline flows (Gad-el-Hak & Blackwelder 1989). These applications typically operate under high Reynolds number conditions, making the turbulent drag reduction at high Reynolds numbers both theoretically significant and practically valuable.

In wall-bounded turbulence, coherent structures are strongly associated with high skin friction (Kravchenko, Choi & Moin 1993; Choi, Moin & Kim 1994; Xu & Huang 2005), leading to the concept that real-time control of these structures could effectively potentially suppress turbulence and reduce skin friction. At low Reynolds numbers, the near-wall region is dominated by velocity streaks and quasi-streamwise vortices. These structures are cyclically generated through a self-sustaining process (Jiménez & Moin 1991; Hamilton, Kim & Waleffe 1995), which can persist even without turbulence in the outer region (Jiménez & Pinelli 1999). The work by Choi *et al.* (1994) pioneered an active control method targeting the streamwise vortices, known as the opposition control strategy. In this method, the wall-normal velocity fluctuations are monitored on a hypothetical detection plane in the near-wall region. Based on these detected signals, counteracting wall-normal blowing and suction velocities are applied at the wall to suppress the ejection and sweep events caused by streamwise vortices, thereby reducing Reynolds shear stress and achieving drag reduction. The effectiveness of opposition control was confirmed by direct numerical simulations (DNS) of turbulent channel flows, as demonstrated by Choi *et al.* (1994), which showed a maximum drag reduction rate approximately 25 % at friction Reynolds number $Re_\tau = 180$. Subsequent investigations by Hammond, Bewley & Moin (1998) and Chung & Talha (2011) further elucidated the mechanisms behind this drag reduction. They found that wall-normal blowing and suction significantly limited momentum transport toward the wall, effectively creating a ‘virtual wall’ that hindered high-speed fluid motions towards the wall induced by streamwise vortices, thus reducing local high friction drag. Building on the concept of opposition control, various other strategies have been developed. These include neural-network-based control schemes (Lee *et al.* 1997) and suboptimal control schemes (Lee, Kim & Choi 1998; Fukagata & Kasagi 2004; Hasegawa & Kasagi 2011), which utilize measurable wall quantities to achieve drag reduction.

As Reynolds numbers increase, the efficacy of drag reduction schemes, such as opposition control, markedly declines. For instance, in turbulent channel flows, the maximum drag reduction rate achieved by opposition control decreases from 25 % at $Re_\tau = 180$ to 18 % at $Re_\tau = 720$ (Chang, Collis & Ramakrishnan 2002; Iwamoto, Suzuki & Kasagi 2002; Pamiès *et al.* 2007). As Reynolds numbers rise, large-scale structures and very-large-scale structures emerge in the logarithmic and outer regions (Jiménez 1998; Kim & Adrian 1999; del Álamo & Jiménez 2003; del Álamo *et al.* 2004; Guala, Hommema & Adrian 2006; Balakumar & Adrian 2007; Hutchins & Marusic 2007a; Monty *et al.* 2009). Hwang (2013) suggested that these structures contribute to Reynolds shear stress, thereby diminishing drag reduction rates. Furthermore, Mathis, Hutchins & Marusic (2009) classified the influence of outer large-scale structures on near-wall turbulence into two effects: superposition and amplitude modulation. The superposition effect, a linear process, represents the footprint of large-scale structures on near-wall turbulence (Hoyas & Jiménez 2006; Hutchins & Marusic 2007b). These large-scale structures extend deeply into the near-wall region, contributing significantly to turbulent kinetic energy (Hoyas & Jiménez 2006; Mathis *et al.* 2009; Marusic, Mathis & Hutchins 2010a).

On the other hand, amplitude modulation is a nonlinear process that describes how small-scale turbulent fluctuations are intensified in large-scale high-speed regions, and suppressed in low-speed regions. Deng & Xu (2012) highlighted that the reduced drag reduction rate at high Reynolds numbers is primarily due to the decreased effectiveness of near-wall turbulence control, which is related to the amplitude modulation effect of large-scale structures.

In recent years, the extensive application of deep reinforcement learning (DRL) has been highlighted in various domains such as video classification, voice recognition and language processing. In fluid mechanics, DRL has also been applied extensively to flow control problems (Guéniat *et al.* 2016; Rabault *et al.* 2019; Han & Huang 2020; Paris, Beneddine & Dandois 2021; Zeng & Graham 2021; Li & Zhang 2022; Varela *et al.* 2022; Lee, Kim & Lee 2023; Guastoni *et al.* 2023; Sonoda *et al.* 2023; Suárez *et al.* 2024). For instance, Varela *et al.* (2022) demonstrated DRL's capability to extend control strategies across varying Reynolds numbers, adapting to different flow characteristics as the Reynolds number increases. Additionally, Suárez *et al.* (2024) leveraged multi-agent DRL to develop three-dimensional strategies as three-dimensional instabilities emerged in the cylinder flow, achieving greater drag reduction than traditional methods. These advancements, driven by artificial intelligence and data science, underscore DRL's robust capability to model complex interactions between inputs and outputs (Jordan & Mitchell 2015). Unlike traditional control methods that depend heavily on researchers' insights, neural-network-based DRL can partially automate this process, constructing highly nonlinear models between input signals and output controls. This makes DRL-based turbulence control particularly appealing, as it offers greater flexibility in selecting input signals, and potentially devises control strategies more attuned to the nonlinear mechanisms of turbulence, thereby enhancing drag reduction effects. The initial foray into using machine learning for drag reduction in channel flows can be traced back to Lee *et al.* (1997), who employed a linear neural network with multiple neurons to predict wall-normal blowing and suction velocities based on spanwise wall shear stress, proposing a straightforward control scheme. In more recent developments, Han & Huang (2020) and Lee *et al.* (2023) utilized reinforcement learning to predict wall-normal velocity fluctuations at the detection plane, effectively replicating opposition control based solely on wall measurements. Moreover, Guastoni *et al.* (2023) and Sonoda *et al.* (2023) have achieved better control models and higher drag reduction rates with reinforcement learning compared to traditional opposition control methods. Collectively, these studies demonstrate the significant potential of DRL in reducing drag in wall-bounded turbulence. While DRL-based control strategies have shown great promise, their practical implementation still presents challenges. Many current approaches rely on detailed flow-domain information, such as velocities at specific wall-normal locations, which may be difficult to measure in real-world settings. This underscores the importance of developing strategies that can bridge the gap between numerical simulations and practical applications.

Higher Reynolds number studies also represent an essential step towards conditions more representative of real-world scenarios. However, previous studies on DRL for turbulence control have been limited to low friction Reynolds numbers, with most Re_τ not exceeding 180. Consequently, research on DRL-based control strategies at higher Reynolds numbers remains scarce. Additionally, there is a significant gap in understanding the drag-reduction mechanisms underlying DRL models. This study aims to address these gaps by extending DRL-based control strategies to high Reynolds numbers. To the best of the authors' knowledge, this is the first study applying DRL control to turbulent channel flows with Re_τ larger than 500. Our main purpose is to evaluate the effectiveness of DRL models

in achieving drag reduction at high Reynolds numbers, and to explore the underlying drag reduction mechanisms from both kinematic and dynamic perspectives.

The paper is organized as follows. The numerical methodologies, including DNS and DRL methods, are detailed in § 2. Section 3 presents the DNS results and their subsequent discussions. The performance of the DRL-based control strategy is evaluated in § 3.1, while velocity statistics are elaborated upon in § 3.2. The analysis of the drag-reduction mechanism is approached from both a kinematic perspective, based on virtual wall theory, and a dynamic perspective, using budget equations, in §§ 3.3 and 3.4, respectively. Finally, the conclusions are summarized in § 4.

2. Numerical methodology

2.1. The DNS of the turbulent channel flows

We consider the turbulent channel flows established between two parallel plates separated by $2h$, driven by a pressure gradient. The governing equations of the turbulent flow are the Navier–Stokes equations of an incompressible Newtonian fluid, written as

$$\frac{\partial u_j}{\partial x_j} = 0, \tag{2.1}$$

$$\frac{\partial u_i}{\partial t} + u_j \frac{\partial u_i}{\partial x_j} = -\frac{1}{\rho} \frac{\partial p}{\partial x_i} + \nu \frac{\partial^2 u_i}{\partial x_j \partial x_j} + f_1 \delta_{i1}, \tag{2.2}$$

where $x_i (i = 1, 2, 3) = (x, y, z)$ represents the coordinates in the streamwise, wall-normal and spanwise directions, respectively, and $u_i (i = 1, 2, 3) = (u, v, w)$ denotes the corresponding velocity components. Here, t is the time, ρ is the density, p is the pressure, and ν is the kinematic viscosity. A body force f_1 is introduced to maintain constant momentum in the channel, ensuring the averaged bulk velocity U_m in the channel flow.

The flow is assumed to be periodic in the streamwise and spanwise directions, with periods L_x and L_z , respectively. The upper wall imposes no-slip and no-penetration conditions, setting the velocities $u = v = w = 0$. On the other hand, the lower wall adheres to the no-slip condition with $u = w = 0$, and implements turbulent control through blowing and suction.

The code AFiD (Verzicco & Orlandi 1996; van der Poel *et al.* 2015; Zhu *et al.* 2018) was utilized to carry out the DNS of turbulent channel flows. An energy-conserving second-order finite difference scheme is applied in the spatial discretization, with velocities on a staggered grid. Time marching is performed using a third-order Runge–Kutta scheme, combined with a Crank–Nicolson scheme for the implicit terms. The grids are uniformly distributed in both the streamwise and spanwise directions, with wall-normal grid refinement applied near the walls.

The computational parameters are listed in table 1 for the three Reynolds numbers $Re = U_m h / \nu$ considered in this study. The friction velocity $u_\tau = \sqrt{\tau_w / \rho}$ and the friction Reynolds number $Re_\tau = h^+ = u_\tau h / \nu$ define wall units in the following discussions, denoted by a $+$ superscript, where τ_w is the skin friction. Here, $y^+ = y / \delta_\nu$, where $\delta_\nu = \nu / u_\tau = h / Re_\tau$ is the friction length.

2.2. The DRL methodology

In order to control the turbulent channel flows, blowing and suction based on the DRL, predictions are applied to the lower wall. The flow chart of the control driven by reinforcement learning is shown in figure 1. Our current program mainly consists of two parts: the numerical simulation part and the reinforcement learning part. The numerical

Set of cases	Re	Re_τ	L_x	L_z	Δ_x^+	Δ_z^+	$(\Delta_y^+)_{min}$	$(\Delta_y^+)_{max}$
C180	2800	176.9	$2\pi h$	πh	8.7	4.3	0.10	4.3
C550	10 000	544.3	$2\pi h$	πh	13.4	6.7	0.19	5.8
C1000	20 000	983.4	$2\pi h$	πh	12.1	6.0	0.19	7.9

Table 1. Computational parameters. Here, Δ_x , Δ_y and Δ_z are the resolutions in the streamwise, wall-normal and spanwise directions, respectively.

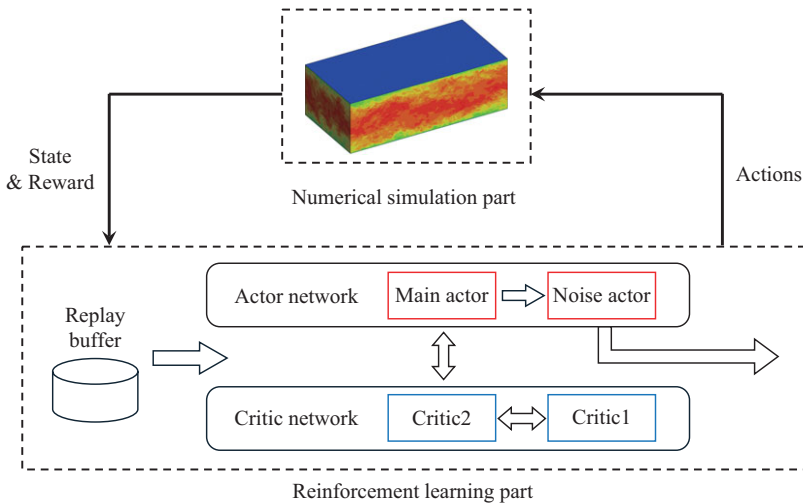


Figure 1. The flow chart of reinforcement-learning-driven control in turbulent channel flows.

simulation part, as discussed in § 2.1, acts as the environment and outputs the state s_t and reward r_t obtained in the flow field. The reinforcement learning part, acting as the agent, receives these variables, optimizes the decision-making policy $\pi(s_t)$ based on the reward, and outputs actions a_t based on the state. The numerical simulation part then uses these actions to control the flow, and advances the simulation in time. This creates a loop to achieve active control driven by reinforcement learning. Here, we select the wall blowing and suction velocities v'_w as the actions, and we choose the streamwise velocity fluctuations $u'(x, z)|_{y+=15}$ in the near-wall region as the states, similar to those adopted by Sonoda *et al.* (2023). Velocity fluctuations are defined based on the mean velocity profile of each case, where $u' = u(x, y, z) - U(y)$. The mean wall blowing and suction velocity is set to zero.

The agent that we adopted is based on the open-source code provided by Lee *et al.* (2023), which employs the twin-delayed deep deterministic policy gradient (TD3) model (Lillicrap *et al.* 2015), an actor–critic network structure. The TD3 model offers improved stability and performance in learning by addressing overestimation bias, incorporating delayed updates, and implementing target smoothing. It has been proven to be suitable for turbulence control optimization (Lee *et al.* 2023). In the TD3 model, the goal is to optimize the action value function $q_\pi(s_t, a_t)$ by satisfying the Bellman equation, where

$$q_\pi(s_t, a_t) = \mathbb{E}[r_t^d + \gamma^n q_\pi(s_{t+n}, \pi_\phi(s_{t+n}) + \epsilon)]. \quad (2.3)$$

Here, $r_t^d = \sum_{j=1}^n \gamma^{j-1} r_{t+j}$ is the n -step reward, γ is the discounted factor, $\pi_\phi(s_{t+n})$ is the delayed policy update, and ϵ is the clipped random noise. We adopt $n = 5$ and $\gamma = 0.95$

Set of cases	Δt^+	$\Delta T U_m/h$	N_{st}
C180	55.8	100.0	20
C550	59.3	40.0	20
C1000	48.4	20.0	20

Table 2. Parameters of state steps and episodes. Here, Δt and ΔT are the time lengths of each state step and episode, respectively, while N_{st} is the number of state steps in one episode.

in all cases, following Lee *et al.* (2023). The expected cumulative reward is predicted by the critic networks, and the objective function for updating the parameters of the critic networks is given by

$$J(\theta) = N^{-1} \sum [r_t^d + \gamma^n q_\pi(s_{t+n}, \pi_\phi(s_{t+n}) + \epsilon) - q(s_t, a_t)]^2, \tag{2.4}$$

where $N = 64$ is the minibatch size, and θ is the weight parameter of the critic networks. The actor network aims to find an optimal policy, guided by the policy objective function $J(\phi)$. Here, ϕ represents the weight parameters of the actor (or critic) networks. The objective function is updated by

$$\nabla_\psi J(\psi) = \mathbb{E}[\nabla_a q_\pi(s_t, a_t) |_{a=\pi(s_t)} \nabla_\psi \pi_\psi(s_t)], \tag{2.5}$$

where ψ is the weight parameter of the actor network. The actor network includes three convolutional layers, with the first two layers activated by the ReLU function. The numbers of filter kernels for these layers are set to 64, 32 and 1, respectively, with each filter kernel sized at 3×3 . In contrast, the critic network is structured with six convolutional layers followed by three fully connected layers, all activated by the ReLU function. Each convolutional layer contains 32 filter kernels of size 3×3 . Additionally, an average pooling layer is applied after every two convolutional layers. The fully connected layers each consist of 32 neurons, and the network ultimately outputs a q value to evaluate the control policy. Detailed hyperparameters can be referred to in Lee *et al.* (2023).

The choice of reward r_t is crucial for the effectiveness of the training outcomes. Inspired by the optimal control (Bewley, Moin & Temam 2001), we define the reward $r = 1 - e/e_0$ as the reduction of integrated turbulent kinetic energy (TKE) in the lower half-channel at the end of each state step. Here, e is the integrated TKE with control, defined as

$$e = \frac{1}{L_x L_z h U_m^2} \int_0^{L_z} \int_0^h \int_0^{L_x} \left(\frac{1}{2} u'_i u'_i \right) dx dy dz, \tag{2.6}$$

and e_0 is the integrated TKE without control. It is important to note that the reward is used only during the model training process, and is no longer needed once the model has converged.

Furthermore, the time lengths of each state step and episode for all cases are shown in table 2. First, Δt needs to be sufficiently long to allow the changes in the control strategy to fully develop. Therefore, $\Delta t^+ \approx 50$ is selected, which also meets the requirement for the prediction horizon $\Delta t^+ > 25$ in optimal control (Bewley *et al.* 2001). Additionally, given that the maximum Re_τ in our cases reaches 1000, it is crucial to train a control strategy that remains effective under large-scale structure evolution. Consequently, we ensure that ΔT is at least $20h/U_m$, which exceeds the time required for large-scale structures to advect streamwise across the entire channel, approximately $L_x/U_m \approx 6h/U_m$.

In summary, a comparison of our computational method with previous DRL-based studies on turbulent channel control is shown in table 3. Our work uses the TD3 algorithm,

Study	DRL algorithm	Input states	Rewards	Re_τ
Guastoni <i>et al.</i> (2023)	DDPG	(u', v')	Drag reduction (DR) rate	180
Sonoda <i>et al.</i> (2023)	DDPG	(u', v')	DR rate & control cost	150
Lee <i>et al.</i> (2023)	TD3	$(\frac{\partial u}{\partial y}, \frac{\partial w}{\partial y}) _{wall}$	DR rate	180, 360
Our work	TD3	u'	TKE reduction rate	180, 550, 1000

Table 3. Comparison of computational details in DRL-based turbulent channel control studies. Here, DDPG denotes the deep deterministic policy gradient algorithm. In all the studies, the output actions selected are the wall blowing and suction velocities v'_w .

Cases	Range of v'_w	Re_τ	DR (%)	P_S/P_I	ΔU_s^+	y_{vw}^+	$-\langle u'v' \rangle_{vw}^+$
C180-0	$v'_w = 0$	176.9	0	–	0	0	0
C180-opp	$v'_w = -v' _{y^+=15}$	154.2	24.0	27.3	2.92	7.1	0.029
C180-1	$v'_w \in [-u_\tau^0, u_\tau^0]$	153.9	24.9	12.4	3.30	9.5	0.144
C180-2	$v'_w \in [-2u_\tau^0, 2u_\tau^0]$	143.1	34.6	22.6	5.17	12.8	0.025
C180-3	$v'_w \in [-3u_\tau^0, 3u_\tau^0]$	142.0	35.6	26.3	5.18	12.7	0.023
C550-0	$v'_w = 0$	544.3	0	–	0	0	0
C550-opp	$v'_w = -v' _{y^+=15}$	480.7	22.0	11.4	3.00	7.3	0.054
C550-1	$v'_w \in [-u_\tau^0, u_\tau^0]$	484.3	20.8	12.1	3.21	7.4	0.125
C550-2	$v'_w \in [-2u_\tau^0, 2u_\tau^0]$	460.0	28.6	13.2	4.23	8.9	0.069
C550-3	$v'_w \in [-3u_\tau^0, 3u_\tau^0]$	454.0	30.4	13.5	4.28	8.7	0.074
C1000-0	$v'_w = 0$	983.4	0	–	0	0	0
C1000-opp	$v'_w = -v' _{y^+=15}$	887.8	18.5	9.2	2.64	6.7	0.044
C1000-1	$v'_w \in [-u_\tau^0, u_\tau^0]$	870.7	21.6	5.0	2.71	7.5	0.306
C1000-2	$v'_w \in [-2u_\tau^0, 2u_\tau^0]$	862.9	23.0	5.1	3.14	8.4	0.299
C1000-3	$v'_w \in [-3u_\tau^0, 3u_\tau^0]$	836.2	27.7	4.7	3.27	10.2	0.242

Table 4. The DNS cases and drag reduction results. Here, DR represents the drag reduction rate, P_S/P_I denotes the power saving ratio (where P_S is the power saving, and P_I is the power input), ΔU_s^+ denotes the shift of the mean velocity profile in the logarithmic region, y_{vw} indicates the height of the virtual wall, and $-\langle u'v' \rangle_{vw}$ is the averaged residual Reynolds stress on the virtual wall.

similar to that of Lee *et al.* (2023), with streamwise velocity fluctuations u' as input states, and the TKE reduction rate as the reward. We have extended DRL-based wall blowing and suction control to higher Reynolds numbers, reaching $Re_\tau = 1000$.

3. The DNS results and discussions

3.1. Performance of the DRL models

In this study, we focus on the DRL-optimized control models under different blowing and suction intensities, and their impact on the flow mechanism. The DNS cases that we utilized are detailed in table 4. Among these, cases with the suffices 0 and ‘opp’ did not use DRL models. The former denotes cases with no blowing or suction, whereas the latter represents cases with opposition control as suggested by Choi *et al.* (1994). Cases with suffices 1, 2 and 3 are based on DRL-optimized control models, where the magnitude of wall blowing and suction v'_w is limited to $[-u_\tau^0, u_\tau^0]$, $[-2u_\tau^0, 2u_\tau^0]$ and $[-3u_\tau^0, 3u_\tau^0]$, respectively. Here, the superscript 0 denotes variables before the application of turbulence control.

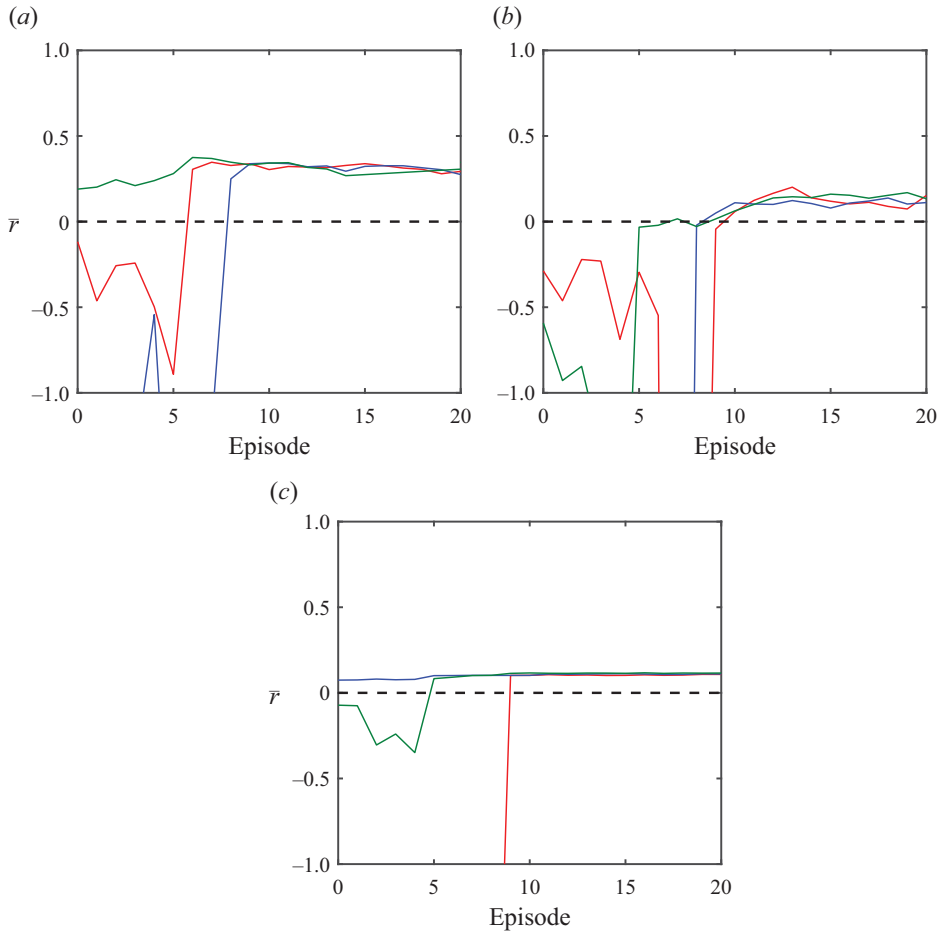


Figure 2. The evolution of the normalized reward over episodes during the training process: (a) C180, (b) C550, (c) C1000. — denotes cases with suffix 1; — denotes cases with suffix 2; — denotes cases with suffix 3.

Before further analysis, it is essential to confirm the training status of the current DRL models. The normalized reward \bar{r} , defined as $\bar{r} = \frac{\sum_{j=1}^n \gamma^{j-1} r_{t+j}}{\sum_{j=1}^n \gamma^{j-1}}$, serves as an indicator of learning performance during the training of the control strategy. The evolution of \bar{r} over episodes for different cases is illustrated in figure 2. Significant oscillations are observed primarily within the first 10 episodes, while the rewards for all analysed cases gradually converge after 10 episodes, indicating stabilization of the DRL models. Consequently, the model at 20 episodes will be utilized uniformly as the control strategy for subsequent analysis. Additionally, we observed that further training beyond 20 episodes does not significantly enhance learning performance, although this is not shown in figure 2. As suggested by Lee *et al.* (2023), prolonged training can lead to issues such as catastrophic forgetting or overfitting, potentially causing the training process to fail. Therefore, the models at 20 episodes are deemed appropriate for our purposes.

The drag reduction effects under different blowing and suction intensities are shown in table 4. As the range of v'_w is expanded, the drag reduction rates (DR) achieved using DRL models continuously improve. When $v'_w \in [-3u_\tau, 3u_\tau]$, the drag reduction effect significantly surpasses that of the traditional opposition control method, including

high-Reynolds-number cases. Specifically, the drag reduction rate is 35.6% in case C180-3, 30.4% in case C550-3, and 27.7% in case C1000-3. On the other hand, it is noted that the maximum drag reduction rate achieved using the DRL model decreases as the Reynolds number increases. This trend is similar to that observed with the opposition control method (Chang *et al.* 2002; Iwamoto *et al.* 2002; Pamiès *et al.* 2007; Toubert & Leschziner 2012; Hwang 2013). We also tested several cases with modified input states or rewards, along with the performance of trained models under different resolutions and Re_τ ; see Appendix A.

Furthermore, as proposed by Bewley *et al.* (2001), the energy efficiency of the active control policy can be quantified by the ratio of power saving to power input, P_S/P_I , as shown in table 4. The power input P_I is calculated as

$$P_I = \langle |p'_w v_w| \rangle + \rho \langle |v_w^3| \rangle / 2, \quad (3.1)$$

where p'_w represents pressure fluctuations on the wall, and ρ denotes the fluid density. The power saving P_S , in turn, is given by $P_S = (\tau_w^0 - \tau_w)U_m$, where τ_w^0 is the wall shear stress before control. For cases with $Re_\tau \leq 550$, the power saving ratio P_S/P_I shows an increasing trend as the range of v'_w expands. Specifically, at $Re_\tau = 180$, the power saving ratio achieved by DRL-based control is lower than that achieved by opposition control. On the other hand, at $Re_\tau = 550$, DRL-based control demonstrates more effective energy savings compared to opposition control. As the Reynolds number increases, the power efficiency ratio P_S/P_I gradually declines, eventually dropping to approximately 5 at $Re_\tau = 1000$. Notably, at $Re_\tau = 1000$, adjustments in the range of v'_w have only a minor impact on the energy savings achieved by DRL-based control.

3.2. Velocity statistics

Further investigation is needed to understand the impact of DRL-based control on flow field statistics and the underlying mechanisms affecting the flow. Therefore, this subsection will compare the velocity statistics across different cases.

Figure 3 shows the wall-normal distributions of the mean velocity profile, U^+ . In the viscous sublayer below $y^+ = 5$, the wall blowing and suction based on the DRL model result in a decrease in the mean velocity compared to the uncontrolled case. In the logarithmic region, the velocity profiles continue to follow the logarithmic law even in the presence of control, but with an upward shift relative to the uncontrolled case. This behaviour is analogous to what is observed with opposition control (Choi *et al.* 1994). The profile shift ΔU_s is detailed in table 4, with ΔU_s calculated as the averaged vertical shift between $y^+ = 50$ and $y/h = 0.5$. As the range of v'_w is extended progressively, the drag reduction rate increases consistently, resulting in a corresponding rise in the mean velocity profile in the logarithmic region, and a gradual increase in ΔU_s . Furthermore, ΔU_s at higher Reynolds numbers gradually decreases, corresponding to a decline in the drag reduction rate.

Wall-normal distributions of velocity fluctuations under different control strategies are illustrated in figure 4. After applying control, the streamwise velocity fluctuations u'_{rms} , indicated by solid lines, show a significant increase in the viscous sublayer below $y^+ = 5$. In contrast, at higher positions, particularly around $y^+ = 10$ in the near-wall region, the peak of u'_{rms} vanishes. The reduction in streamwise velocity fluctuations becomes more pronounced as the range of blowing and suction velocities is further expanded. This trend is observed consistently across different Reynolds numbers. However, it is noteworthy that for $Re_\tau^0 \approx 550$ and 1000, the impact of wall blowing and suction is trivial in the outer region, as depicted in figures 4(b) and 4(c). The current DRL-based control strategy hardly

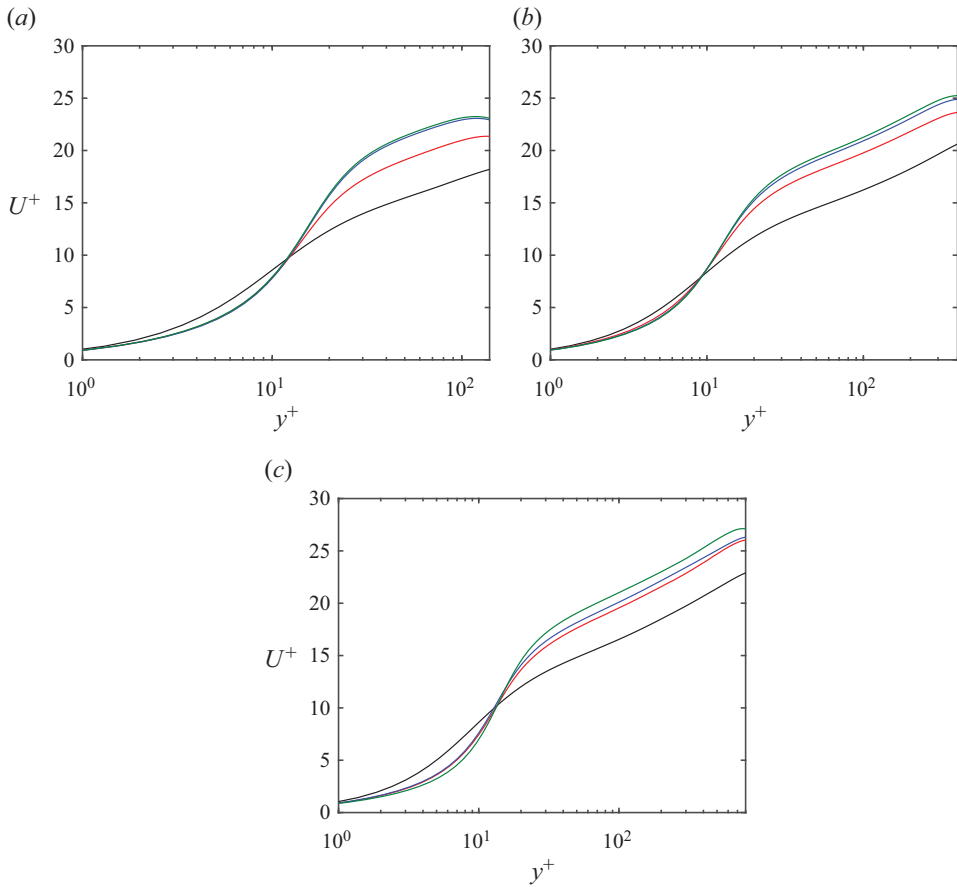


Figure 3. Mean velocity profile under different control strategies: (a) C180, (b) C550, (c) C1000. — denotes cases with suffix 0; — denotes cases with suffix 1; — denotes cases with suffix 2; — denotes cases with suffix 3.

affects the outer region, which is dominated by large-scale structures. On the other hand, the application of control significantly increases the wall-normal velocity fluctuations v'_{rms} within the viscous sublayer, due to the direct impact of blowing and suction. Conversely, the spanwise velocity fluctuations w'_{rms} below $y^+ = 5$ exhibit minimal changes. As the range of blowing and suction velocities is further extended, both v'_{rms} and w'_{rms} within $10 < y^+ < 30$ gradually decrease. This trend is particularly evident at the low Reynolds number $Re_\tau^0 \approx 180$, as shown in figure 4(a), but becomes less pronounced at higher Reynolds numbers. Additionally, in the controlled cases, v'_{rms} in the near-wall region initially decreases and then increases with height. The point of minimum v'_{rms} in the near-wall region can be defined as the position of the virtual wall y_{vw} (Hammond *et al.* 1998). The virtual wall and the residual fluctuations on it will be discussed further in § 3.3.

Figure 5 presents the wall-normal distributions of the averaged Reynolds shear stress, where $\langle \varphi \rangle$ denotes the variable $\varphi(x, y, z, t)$ averaged over the streamwise, spanwise and temporal directions. Although not shown in the figure, the Reynolds stress at the wall is always zero, confined by the boundary conditions. The Reynolds stress in the viscous sublayer is higher in the controlled case compared to the uncontrolled case, due to the application of blowing and suction at the wall. In the controlled case, the Reynolds stress slightly increases with height, reaching a peak near $y^+ = 5$, before rapidly decreasing

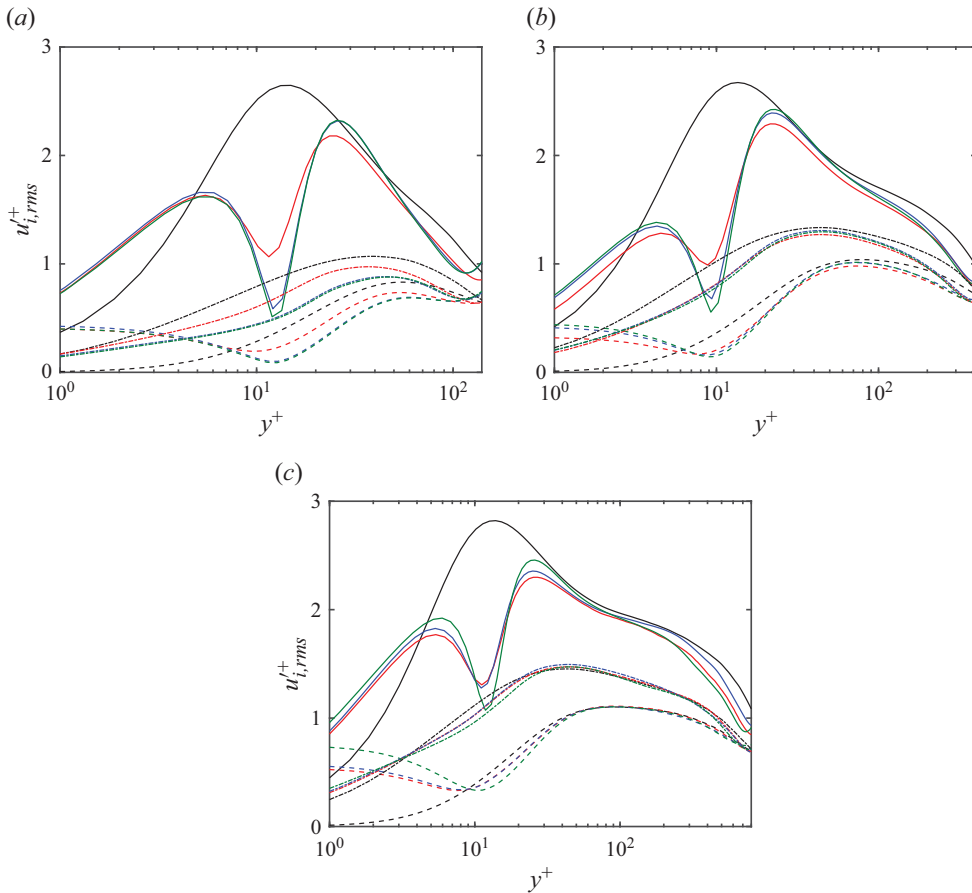


Figure 4. Wall-normal distributions of the velocity fluctuations under different control strategies. Solid, dashed and dash-dotted lines represent u'_{rms} , v'_{rms} and w'_{rms} , respectively: (a) C180, (b) C550, (c) C1000. — denotes cases with suffix 0; — denotes cases with suffix 1; — denotes cases with suffix 2; — denotes cases with suffix 3.

and forming a trough at approximately $y^+ = 10-12$. Compared to the uncontrolled case, the Reynolds stress with control significantly decreases in the region $10 < y^+ < 20$. This decreasing trend becomes more pronounced as the range of v'_w is further extended. According to the FIK identity proposed by Fukagata, Iwamoto & Kasagi (2002), this reduction in Reynolds stress also leads directly to a decrease in the skin friction. At a low Reynolds number $Re_\tau^0 \approx 180$, the Reynolds stress in the logarithmic region is lower with control, as shown in figure 5(a). However, this trend gradually disappears at higher Reynolds numbers, corresponding to a decrease in drag reduction rate.

The relationship between wall blowing and suction velocity and the velocity fluctuations at the detection plane is a crucial aspect of flow control. Traditional opposition control employs blowing and suction with equal magnitudes but opposite directions to the wall-normal velocity fluctuations at the detection plane. Consequently, the correlation R between the blowing and suction velocity v'_w and v' at $y^+ = 15$, defined as

$$R(v'_w, v') = \frac{\langle v'_w v' \rangle}{\sqrt{\langle v'_w v'_w \rangle \langle v' v' \rangle}}, \quad (3.2)$$

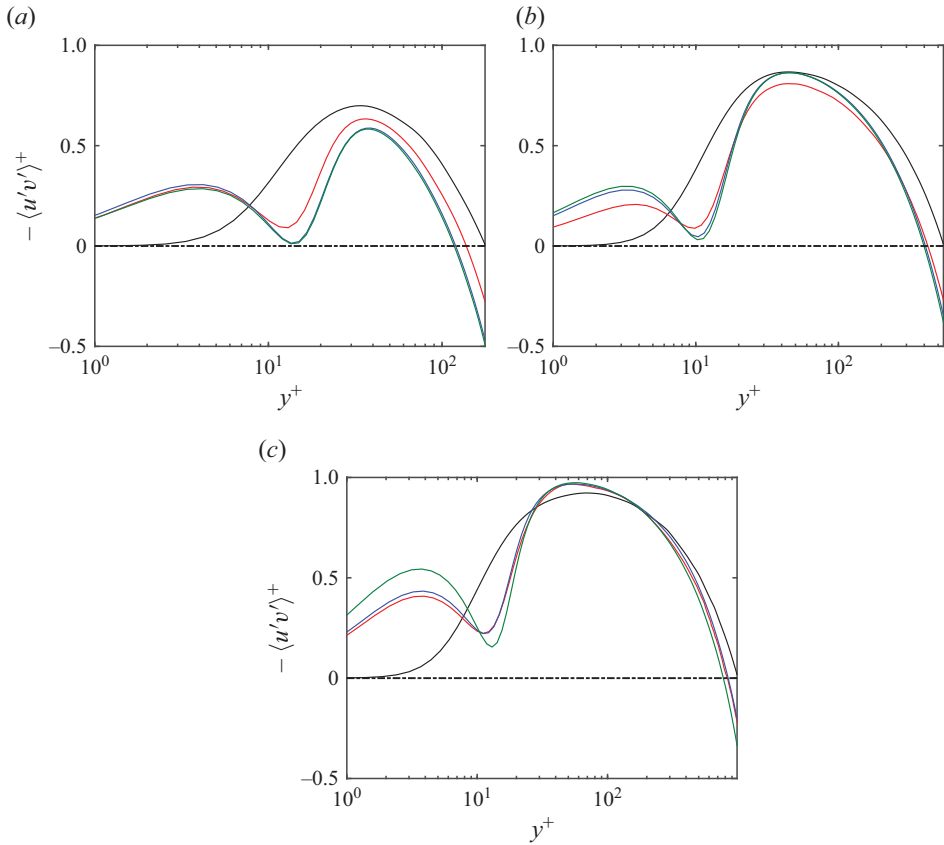


Figure 5. Wall-normal distributions of the averaged Reynolds shear stress under different control strategies: (a) C180, (b) C550, (c) C1000. — denotes cases with suffix 0; — denotes cases with suffix 1; — denotes cases with suffix 2; — denotes cases with suffix 3.

would be strictly -1 . Unlike opposition control, the current DRL-based control strategy is based on the streamwise velocity fluctuations u' at the detection plane $y^+ = 15$ as the input state. The joint probability density function (p.d.f.) of wall blowing and suction v'_w with velocity fluctuations at the near-wall detection plane is shown in figure 6, using the results from case C1000-3 as an example. Here, v'_w has a relatively weak correlation with v' at the detection plane $y^+ = 15$, as illustrated in figure 6(a), where $R(v'_w, v') = -0.10$. Conversely, the joint p.d.f. between v'_w and u' at $y^+ = 15$ in figure 6(b) is predominantly aligned with the first and third quadrants. This alignment indicates that wall blowing tends to occur beneath high-speed regions near the wall, while suction is more likely beneath low-speed regions. Furthermore, the correlation $R(v'_w, u') = 0.71$ is positive. This behaviour closely resembles the mechanism identified by Lee *et al.* (2023), where their DRL models, using wall streamwise shear stress as input states, exhibited a similar trend of applying blowing beneath high-speed regions and suction beneath low-speed streaks. They also observed a strong correlation between wall actuation and streamwise wall shear stress, similar to figure 6(b), suggesting that these DRL models effectively reduce drag through direct control of sweep and ejection events. Despite not being shown in the figure, this pattern appears consistently across all cases in our current work, implying a stronger connection between the DRL-based control strategies and the streamwise velocity fluctuations as the input state.

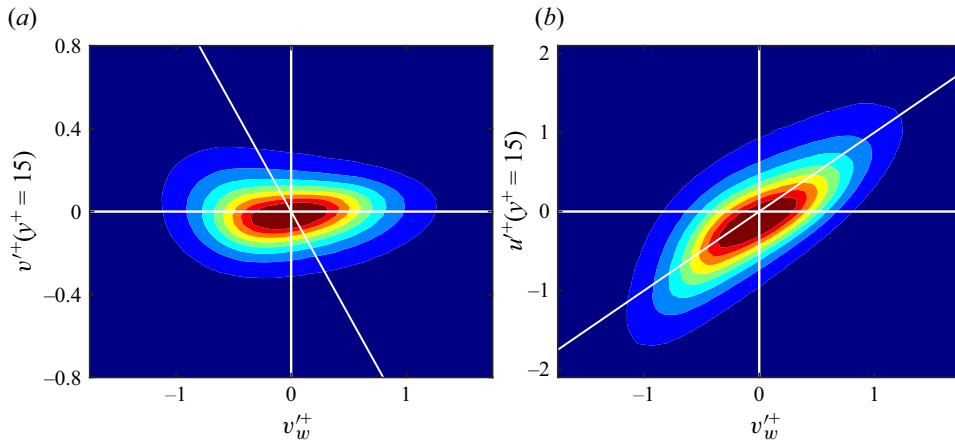


Figure 6. Joint p.d.f. of the wall blowing and suction v'_w with (a) v' and (b) u' at $y^+ = 15$ in case C1000-3. The white diagonals denote (a) $v' = -v'_w$ and (b) $u' = v'_w$, respectively. Contour levels are 0.1(0.1)0.8 of the maximum probability density.

3.3. Kinematic analysis of drag reduction based on virtual wall theory

The DRL-based control strategy could lead to larger drag reduction compared to the traditional opposition control; however, the underlying mechanism requires further investigation. This subsection utilizes the virtual wall theory proposed by Hammond *et al.* (1998) to analyse the drag reduction mechanism from a kinematic perspective.

According to the virtual wall theory by Hammond *et al.* (1998), wall blowing and suction create a virtual wall between the actual wall and the detection plane. This virtual wall hinders streamwise vortices from bringing high-speed fluid to the wall, which would otherwise create local high friction zones, thereby resulting in drag reduction. The drag reduction effect is influenced mainly by two factors: the height of the virtual wall y_{vw} , and the magnitude of the residual Reynolds stress on the virtual wall $-\langle u'v' \rangle_{vw}$. Specifically, the higher the virtual wall, the better the drag reduction effect. And the lower the residual Reynolds stress, the stronger the virtual wall's ability to impede wall-normal momentum transport, resulting in better drag reduction.

The height of the virtual wall and the residual Reynolds stress under different control strategies are detailed in table 4. At the low Reynolds number $Re_\tau^0 \approx 180$, as the range of blowing and suction velocities is further expanded, the height of the virtual wall gradually increases, and the residual Reynolds stress on the virtual wall gradually decreases. Both these changes correspond to an improvement in drag reduction effect. The values of y_{vw} and $-\langle u'v' \rangle_{vw}$ for C180-2 and C180-3 are similar, resulting in comparable drag reduction rates for both cases. Compared to the traditional opposition control method, the DRL-based control strategy in C180-3 does not show a significant reduction in residual stress on the virtual wall. However, its primary benefit is the ability to further elevate the virtual wall. As the Reynolds number increases, the height of the virtual wall under the DRL-based control strategy is significantly lower for $Re_\tau^0 \approx 550$ and 1000 compared to the results for $Re_\tau^0 \approx 180$. Additionally, the residual stress $-\langle u'v' \rangle_{vw}$ on the virtual wall rapidly increases with the rising Reynolds number. In case C550-3, $-\langle u'v' \rangle_{vw}$ is approximately three times that of case C180-3, and in case C1000-3, $-\langle u'v' \rangle_{vw}$ exceeds that of case C180-3 by more than ten times. These two factors together lead to a decrease in the drag reduction efficiency of the DRL-based control strategy.

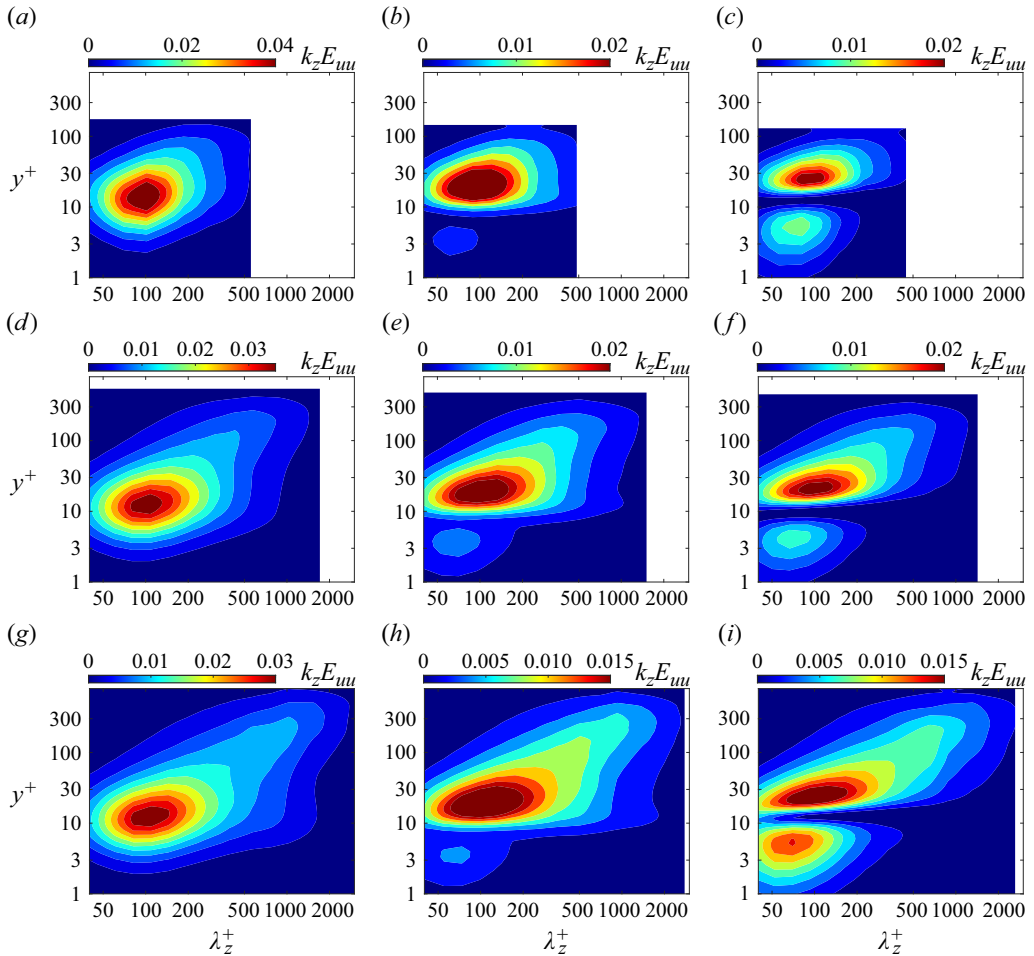


Figure 7. Premultiplied spanwise energy spectra $k_z E_{uu}$ of streamwise velocity fluctuations u' under different control strategies. For $Re_\tau^0 \approx 180$, (a) C180-0, (b) C180-opp, (c) C180-3. For $Re_\tau^0 \approx 550$, (d) C550-0, (e) C550-opp, (f) C550-3. For $Re_\tau^0 \approx 1000$, (g) C1000-0, (h) C1000-opp, (i) C1000-3.

Moreover, at higher Reynolds numbers, the control strategy optimized through DRL is less effective at impeding wall-normal momentum transport compared to traditional opposition control, as evidenced by the contrast of residual Reynolds stress. The main advantage of DRL optimization lies in its ability to effectively plan the wall blowing and suction in an expanded range, thereby elevating the virtual wall to a higher position and achieving better drag reduction efficiency.

To further quantify the impact of control strategies on the scales of the structures at different heights, especially the flow structures near the virtual wall, figure 7 presents the premultiplied energy spectra $k_z E_{uu}$ of u' . Here, k_z is the spanwise wavenumber, and $\lambda_z = 2\pi/k_z$ is the corresponding wavelength. In the near-wall region, the flow is dominated by streaks with spanwise scale $\lambda_z^+ \approx 100$ and wall-normal height concentrated around $y^+ = 15$, as depicted in figures 7(a), 7(d) and 7(g). After applying wall blowing and suction control, the peak velocity fluctuations in the near-wall region shift to a higher position, and a second spectral peak emerges in the viscous sublayer. These two peaks are separated by the virtual wall, as also suggested by Hammond *et al.* (1998). Notably, in

cases utilizing the DRL-based control strategy (C180-3, C550-3 and C1000-3), the virtual wall is significantly higher than in cases using traditional opposition control, corroborating the conclusions drawn from [table 4](#). The peak of velocity fluctuations corresponding to the near-wall streaks in the buffer layer also rises to a higher position. Furthermore, the intensity of velocity fluctuations in the viscous sublayer significantly increases after applying the DRL-based control strategy.

On the other hand, the characteristic scales of flow structures remain largely unchanged under different control strategies. The spanwise sizes of the near-wall streaks are consistently $\lambda_z^+ \approx 100$, while the peak of velocity fluctuations in the viscous sublayer due to wall blowing and suction stays within $\lambda_z^+ = 60\text{--}80$, i.e. slightly smaller than the spanwise sizes of the streaks. At high Reynolds numbers, wall blowing and suction have a trivial effect on the spanwise sizes of the outer large-scale structures, which remain at $\lambda_z \approx O(h)$. In the cases without control, the footprint of outer large-scale structures penetrates deeply into the near-wall region, as shown by the near-wall large-scale components in [figures 7\(d\)](#) and [7\(g\)](#). This phenomenon, known as the superposition effect (Hoyas & Jiménez 2006; Hutchins & Marusic, 2007b; Mathis *et al.* 2009; Marusic *et al.* 2010a), is noteworthy. After applying wall control, however, the footprint of outer large-scale structures cannot penetrate the virtual wall to reach the viscous sublayer or contribute to the residual velocity fluctuations on the virtual wall. This is particularly evident in cases C550-3 and C1000-3 using DRL models, as shown in [figures 7\(f\)](#) and [7\(i\)](#). Thus the superposition effect does not directly cause the increasing residual Reynolds stress on the virtual wall at rising Reynolds numbers. Its impact on the decreasing drag reduction rate of the DRL models at high Reynolds numbers is also trivial.

To identify the source of $-\langle u'v' \rangle_{vw}$ at high Reynolds numbers, [figure 8](#) illustrates the distributions of streamwise velocity fluctuations at $y^+ = y_{vw}^+$ and $y^+ = 150$. The DRL-based control strategy reveals strong fluctuations on the virtual wall, characterized by clustered small-scale fluctuations concentrated in specific areas. Although these fluctuations are mitigated when the range of blowing and suction velocities is expanded, they remain stronger than those observed after opposition control. It shall be noted that these fluctuations are much smaller in size compared to the outer large-scale structures. Therefore, they are unlikely to be induced by the linear superposition effect, but are more plausibly related to the nonlinear amplitude modulation mechanism of the large-scale structures. As indicated by the black rectangles in [figure 8](#), regions of strong fluctuations on the virtual wall often share similar spanwise locations with the outer large-scale high-speed regions, further supporting this point. In the streamwise direction, areas with clustered fluctuations are frequently situated upstream of the large-scale high-speed regions. This phenomenon can be attributed to the inclination angle of the large-scale coherent structures, as suggested by the near-wall fluctuation predictive models proposed by Marusic, Mathis & Hutchins (2010b) and Mathis, Hutchins & Marusic (2011).

Further statistical evidence is required to support the relationship between the amplitude modulation of outer large-scale structures and the residual Reynolds stress at the virtual wall. The streamwise velocity fluctuations u'_0 at the centre of the logarithmic region $y_0^+ \approx 3.9\sqrt{Re_\tau}$ can be utilized to characterize outer large-scale structures (Mathis *et al.* 2009, 2011). A positive u'_0 indicates a large-scale high-speed region, while a negative u'_0 denotes a low-speed region. On the other hand, the residual fluctuations at the virtual wall exhibit a clustered distribution, as illustrated in [figure 8](#). In areas with strong fluctuations, the streamwise and spanwise scales of the fluctuations are smaller, and the spatial alternation between positive and negative values is more pronounced. Considering the impact of spatial alternation, we select the envelope of the Reynolds stress at the virtual

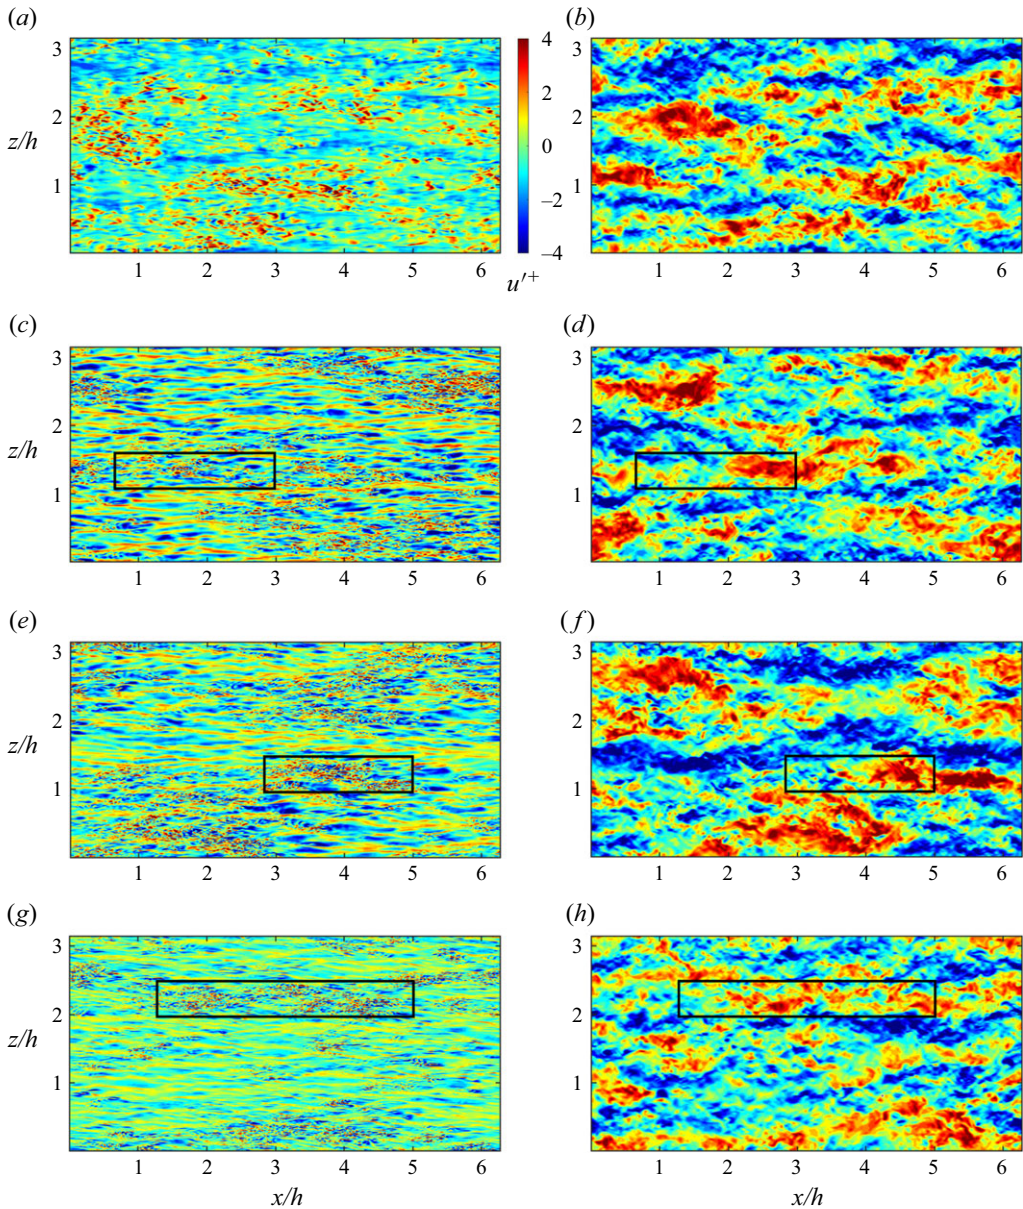


Figure 8. Instantaneous distributions of u' on the (x, z) plane at (a,c,e,g) $y^+ = y_{vw}^+$ and (b,d,f,h) $y^+ = 150$, for cases (a,b) C1000-opp, (c,d) C1000-1, (e,f) C1000-2, (g,h) C1000-3. The black rectangles represent some sample areas on the virtual wall where velocity fluctuations are stronger.

wall, denoted as $|\mathcal{H}(\langle u'v' \rangle_{vw})|$, to measure the strength of the residual stress fluctuations, where \mathcal{H} represents the operator of the two-dimensional Hilbert transform. Additionally, the inclination angle θ_L of the large-scale structures should be considered. The outer large-scale structures affecting the near-wall region are located downstream of this region. Hence we will examine primarily the relationship between the virtual wall fluctuations

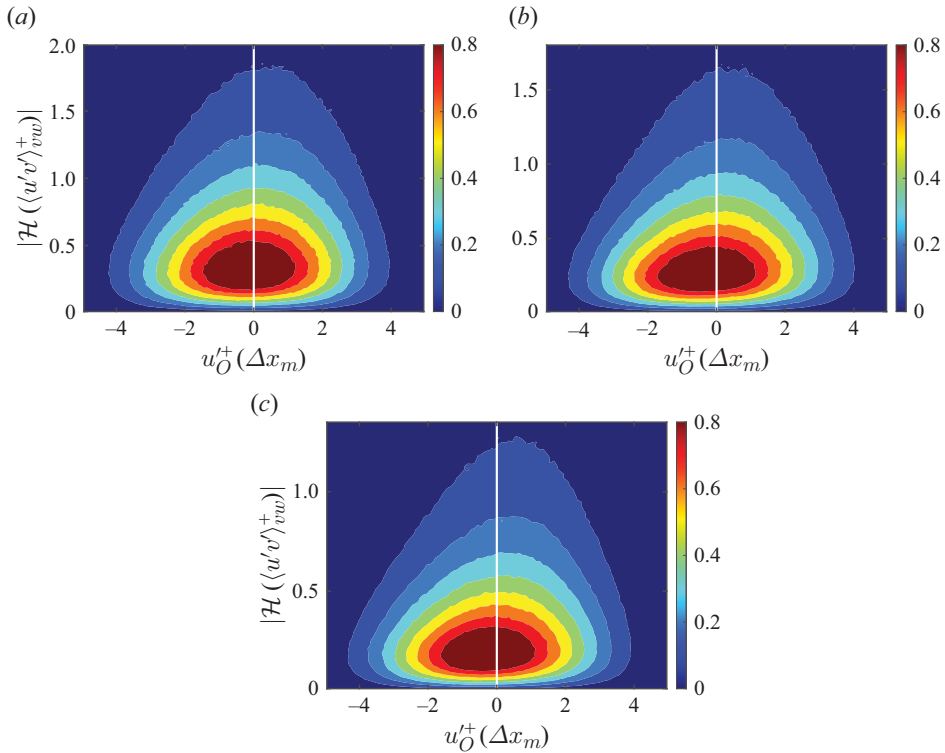


Figure 9. Joint p.d.f. of the streamwise velocity fluctuations $u'_O(\Delta x_m)$ at the centre of the logarithmic region and the envelope of Reynolds stress $|\mathcal{H}(\langle u'v' \rangle_{vw})|$ at the virtual wall, for cases (a) C1000-1, (b) C1000-2, (c) C1000-3. Contour levels are 0.1(0.1)0.8 of the maximum probability density.

and $u'_O(\Delta x_m)$ at a downstream displacement Δx_m . Here, $\Delta x_m = (y_O - y_{vw}) / \tan(\theta_L)$ and $\theta_L = 11^\circ - 15^\circ$ according to Mathis *et al.* (2011). We select $\theta_L = 13^\circ$ for the subsequent discussions, noting that the results are robust within the range $\theta_L = 11^\circ - 15^\circ$.

The joint p.d.f. between the outer $u'_O(\Delta x_m)$ and the envelope of Reynolds stress at the virtual wall is depicted in figure 9. At the position where $|\mathcal{H}(\langle u'v' \rangle_{vw})|$ approaches 0, the joint p.d.f. tilts to the left, indicating negative u'_O in low-speed large-scale motions. As the envelope of Reynolds stress gradually increases, the joint p.d.f. shifts, tilting to the right, which is particularly evident in the upper half of the distribution. This pattern suggests that locations with strong residual Reynolds stress fluctuations are typically situated below large-scale high-speed regions, while areas with weaker residual Reynolds stress generally correspond to large-scale low-speed regions. This observation is consistent with the findings illustrated in figure 8. As the range of blowing and suction velocities is extended, although the intensity of $-\langle u'v' \rangle_{vw}$ diminishes, the influence of outer large-scale structures on the distribution of Reynolds stress remains nearly unchanged. This further substantiates the relationship between the amplitude modulation of outer large-scale structures and the residual Reynolds stress at the virtual wall.

In summary, compared to the traditional opposition control method, the DRL-based control strategy demonstrates superior drag reduction capabilities by effectively elevating the virtual wall to a higher position. As the range of blowing and suction velocities

is expanded, the virtual wall ascends further and the residual Reynolds stress on the virtual wall decreases, both of which enhance the drag reduction rate of the DRL models. However, as the Reynolds number increases, large-scale structures emerge in the outer region. Their amplitude modulation effect significantly increases the residual Reynolds stress on the virtual wall, and disrupts the virtual wall's blockage in large-scale high-speed regions, thereby reducing the drag reduction rate of the DRL models.

3.4. Dynamic analysis of drag reduction using budget equations

In the previous subsection, the drag reduction mechanism was examined from a kinematic perspective using the virtual wall theory. This subsection will further discuss the dynamics mechanism behind drag reduction based on the analysis of budget equations.

According to the FIK identity proposed by Fukagata *et al.* (2002), the skin frictions in the current cases are primarily attributed to the Reynolds shear stress $\langle -u'v' \rangle$. Hence it is necessary to discuss how the DRL-based control strategies reduce the drag by altering $\langle -u'v' \rangle$. The transport equation of the Reynolds stress $\langle -u'v' \rangle$ is written as

$$\begin{aligned} \frac{\partial}{\partial t} \langle -u'v' \rangle = & \underbrace{\langle v'v' \rangle \frac{dU}{dy}}_{P_{12}} + \underbrace{\frac{d}{dy} \langle u'v'v' \rangle}_{D_{12,t}} + \underbrace{\left(-\mu \frac{d^2 \langle u'v' \rangle}{dy^2} \right)}_{D_{12,v}} \\ & + \underbrace{\frac{1}{\rho} \left\langle v' \frac{\partial p'}{\partial y} + u' \frac{\partial p'}{\partial x} \right\rangle}_{VP_{12}} + \underbrace{2\mu \left\langle \frac{\partial u'}{\partial x_j} \frac{\partial v'}{\partial x_j} \right\rangle}_{\varepsilon_{12}}, \end{aligned} \tag{3.3}$$

where P_{12} is the turbulent production, $D_{12,t}$ is the turbulent diffusion, $D_{12,v}$ is the viscous diffusion, VP_{12} is the velocity pressure-gradient term, and ε_{12} is the dissipation. Here, U is the mean streamwise velocity.

Figure 10 shows the wall-normal distributions of the budget terms on the right-hand side of (3.3). In the budget terms of the Reynolds shear stress, the production P_{12} and the velocity pressure-gradient term VP_{12} are significantly stronger and more dominant compared to the other terms, as suggested in figure 10(a,c,e). The viscous diffusion $D_{12,v}$, although large and negative in the viscous sublayer, decays rapidly above $y^+ = 5$ to become smaller than the dominant terms. Furthermore, the dissipation ε_{12} and the turbulent diffusion $D_{12,t}$ are much smaller than the other terms, and their contribution could be considered negligible. Among the two dominant terms, the velocity pressure-gradient term VP_{12} , which can be further divided into the pressure diffusion and the redistribution, mainly represents the transport of Reynolds stress at different heights, and the redistribution among different components caused by pressure. And it primarily acts as a negative term to offset the production P_{12} , which remains positive and determines the magnitude of the Reynolds shear stress.

In the uncontrolled cases, the turbulent production P_{12} increases with height, reaching a peak at approximately $y^+ = 15\text{--}20$, and then decreases continuously. After implementing wall blowing and suction, P_{12} at the wall is no longer zero, leading to a significant increase in P_{12} within the viscous sublayer. This results in a larger Reynolds stress in the viscous sublayer compared to that in the uncontrolled cases, as depicted in figure 5. However, this effect is confined to the narrow height range of the viscous sublayer, and has a limited impact on the overall skin friction. In the cases with control, P_{12} decreases rapidly with height, and reaches a trough at approximately $y^+ = 10$. In the range $10 < y^+ < 20$, P_{12} is significantly smaller than in the uncontrolled case, corresponding to a lower Reynolds

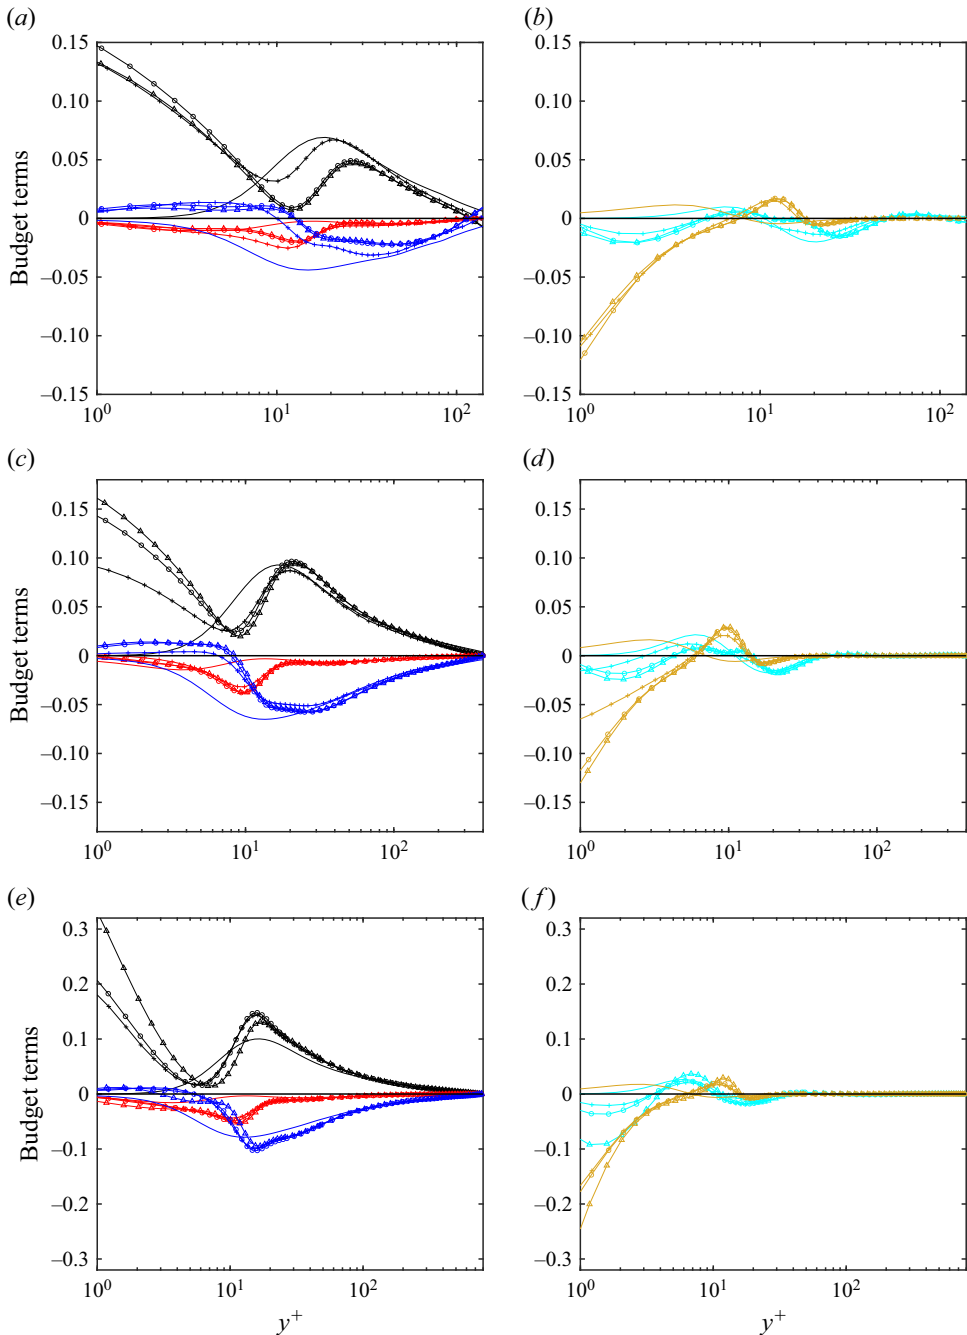


Figure 10. Wall-normal distributions of the budget terms of Reynolds shear stress $\langle -u'v' \rangle$ in (3.3): — P_{12} , — VP_{12} , — ε_{12} , — $D_{12,t}$, — $D_{12,v}$, for (a,b) C180, (c,d) C550, (e,f) C1000. Lines without markers indicate cases with suffix 0; plus signs indicate cases with suffix 1; circles indicate cases with suffix 2; triangles indicate cases with suffix 3.

shear stress in figure 5. As the range of blowing and suction velocities is expanded, the height corresponding to the trough gradually increases, and P_{12} at the trough further decreases, leading to a reduction in Reynolds shear stress. Moreover, the decrease in P_{12}

near the trough compared with the uncontrolled case is less pronounced at higher Reynolds numbers, as indicated in figure 10(e). This results in a reduced suppression effect on $\langle -u'v' \rangle$ at higher Reynolds numbers, shown in figure 5(c), further leading to a decreased drag reduction rate. As the height increases, P_{12} in the controlled case gradually rises above $y^+ = 15$, peaks, and then decreases continuously, eventually collapsing with the uncontrolled case in the outer region.

According to (3.3), the turbulent production P_{12} of Reynolds shear stress consists of two parts: $\langle v'v' \rangle$ and dU/dy . The latter could be viewed as the outcome associated with changes in Reynolds shear stress and skin friction. Therefore, the following discussion will focus primarily on the wall-normal kinetic energy $\langle v'v' \rangle$ to identify the source of changes in P_{12} . The wall-normal distributions of v_{rms} under different control strategies have already been shown and discussed in figure 4. After adopting the DRL models, the wall-normal velocity fluctuations in the buffer layer gradually decrease. As the range of blowing and suction velocities is extended, v_{rms} continues to decrease, but this decreasing trend slows down with increasing Reynolds number. This is similar to the evolution trend of the turbulent production P_{12} . The transport equation of the wall-normal kinetic energy $\langle v'v' \rangle$ is written as

$$\begin{aligned} \frac{\partial}{\partial t} \langle v'v' \rangle = & \underbrace{-\frac{d}{dy} \langle v'v'v' \rangle}_{D_{22,t}} + \underbrace{\mu \frac{d^2 \langle v'v' \rangle}{dy^2}}_{D_{22,v}} + \underbrace{\left(-\frac{2}{\rho} \frac{d}{dy} \langle p'v' \rangle \right)}_{D_{22,p}} \\ & + \underbrace{\frac{2}{\rho} \left\langle p' \frac{\partial v'}{\partial y} \right\rangle}_{\Phi_{22}} + \underbrace{\left(-2\mu \left\langle \frac{\partial v'}{\partial x_j} \frac{\partial v'}{\partial x_j} \right\rangle \right)}_{\varepsilon_{22}}, \end{aligned} \tag{3.4}$$

where $D_{22,t}$ is the turbulent diffusion, $D_{22,v}$ is the viscous diffusion, $D_{22,p}$ is the pressure diffusion, Φ_{22} is the redistribution, and ε_{22} is the dissipation.

The wall-normal distributions of the budget terms on the right-hand side of (3.4) are illustrated in figure 11. Among these budget terms, the wall-normal kinetic energy $\langle v'v' \rangle$ is influenced predominantly by redistribution Φ_{22} , pressure diffusion $D_{22,p}$, and dissipation ε_{22} . Conversely, the effects of turbulent diffusion $D_{22,t}$ and viscous diffusion $D_{22,v}$ are comparatively minor. In the viscous sublayer, the redistribution Φ_{22} is primarily negative, indicating that the wall-normal velocity fluctuations are being redistributed to other directions. This negative contribution is offset mainly by the positive pressure diffusion $D_{22,p}$. As the height increases, the redistribution Φ_{22} changes from negative to positive, indicating that the wall-normal velocity fluctuations are absorbing TKE from other components. Meanwhile, the pressure diffusion $D_{22,p}$ rapidly decreases and gradually approaches zero above $y^+ = 20$. On the other hand, the dissipation ε_{22} gradually increases, acting as a negative term to offset the positive contribution from the redistribution Φ_{22} . Based on the previous discussion, the drag reduction achieved by the DRL model stems mainly from the dynamic changes in the buffer layer, while significant changes in the viscous sublayer contribute very little to the overall drag reduction. Among the three dominant terms, the pressure diffusion $D_{22,p}$ primarily represents the TKE transport caused by pressure at different heights, with its intensity decreasing significantly in the buffer layer compared to the viscous sublayer. Therefore, in the following discussion, we will focus primarily on the redistribution Φ_{22} , which represents the exchange mechanism between wall-normal velocity fluctuations and other velocity components.

Compared to the uncontrolled cases, the DRL-based control strategy causes a significant decrease in Φ_{22} in the buffer layer, and raises the position where it changes from

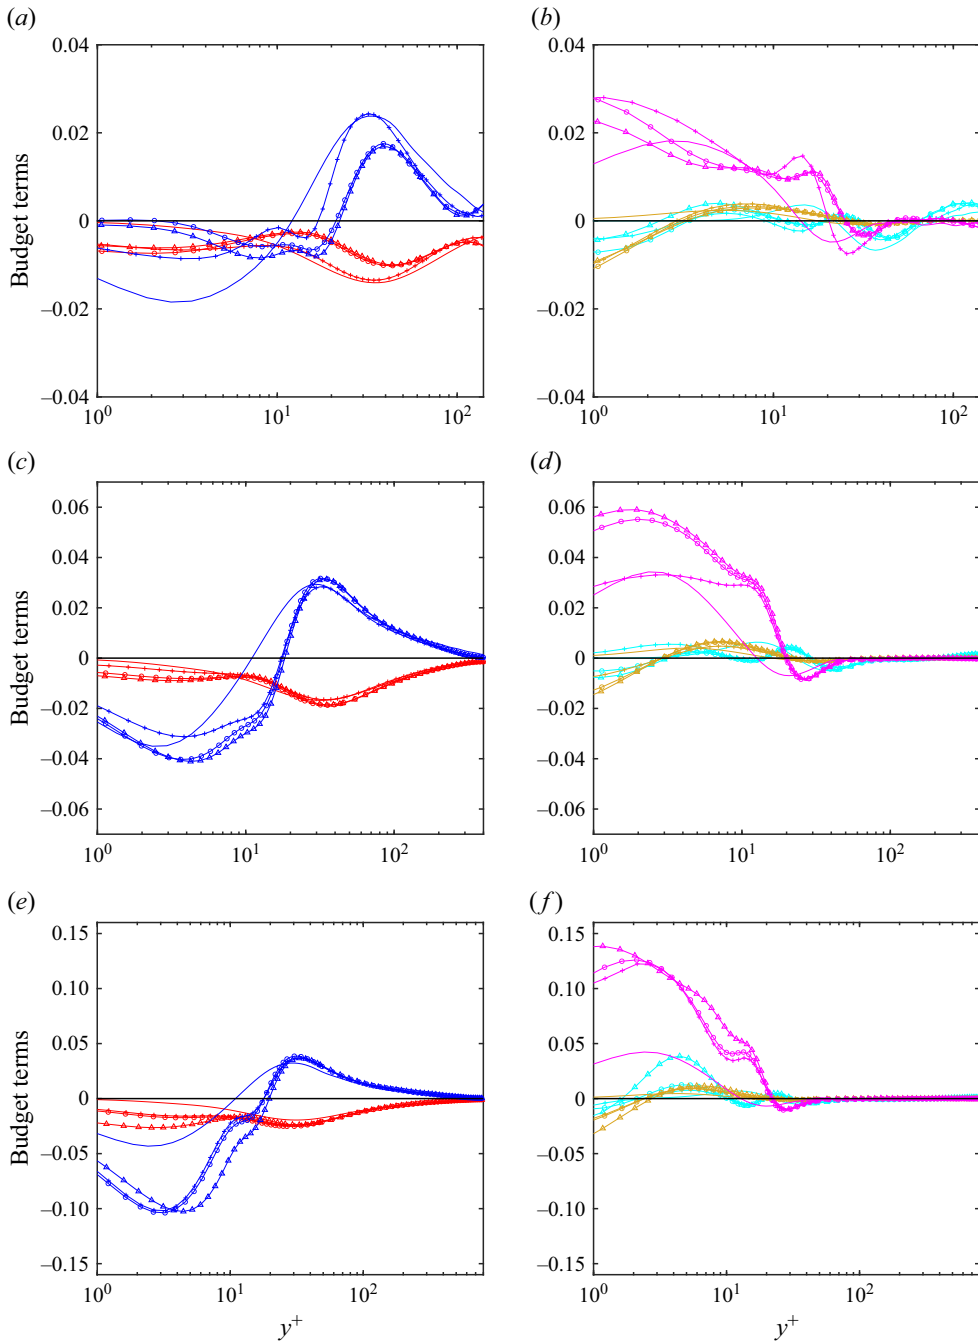


Figure 11. Wall-normal distributions of the budget terms of wall-normal kinetic energy $\langle v'v' \rangle$ in (3.4): — Φ_{22} , — ε_{22} , — $D_{22,p}$, — $D_{22,t}$, — $D_{22,v}$, for (a,b) C180, (c,d) C550, (e,f) C1000. Lines without markers indicate cases with suffix 0; plus signs indicate cases with suffix 1; circles indicate cases with suffix 2; triangles indicate cases with suffix 3.

negative to positive to approximately $y^+ = 20$. This leads to less kinetic energy being transferred to the wall-normal velocity fluctuations in the buffer layer, thereby suppressing the production of Reynolds stress. The suppressing effect of the DRL models on the

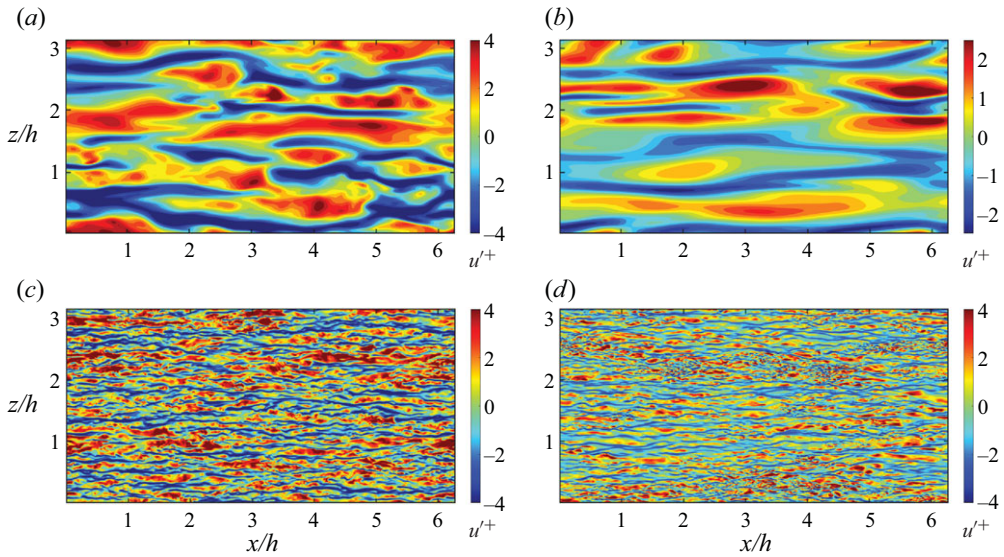


Figure 12. Instantaneous distributions of u' on the (x, z) plane at $y^+ = 20$, for cases (a) C180-0, (b) C180-3, (c) C1000-0, (d) C1000-3.

redistribution Φ_{22} increases further as the range of blowing and suction velocities is extended. At a low Reynolds number $Re_\tau^0 \approx 180$, this decreasing trend can extend to the logarithmic layer, as illustrated in figure 11(e). However, as the Reynolds number increases, the reduction in the redistribution Φ_{22} above $y^+ = 30$ nearly vanishes, and the suppression of Φ_{22} in the buffer layer by the DRL models becomes weaker. This change corresponds to the decrease in drag reduction rate at higher Reynolds numbers.

It is important to note that due to the incompressibility condition (where divergence equals zero), the sum of the redistribution terms for the three velocity components, namely $\Phi_{11} = (2/\rho)\langle p' \partial u'/\partial x \rangle$, $\Phi_{22} = (2/\rho)\langle p' \partial v'/\partial y \rangle$ and $\Phi_{33} = (2/\rho)\langle p' \partial w'/\partial z \rangle$, is 0. Above the viscous sublayer, the TKE is typically redistributed from the streamwise component to the wall-normal and spanwise components (Lee & Moser 2019), also indicated by the positive Φ_{22} observed in figure 11. This redistribution corresponds to the transient growth of the streamwise velocity streaks associated with u' , leading to the generation of quasi-streamwise vortices associated with v' and w' in the near-wall turbulent self-sustaining cycle. Thus the observed weakening of Φ_{22} due to the DRL-based control strategy can be interpreted as the suppression of the near-wall self-sustaining mechanism. Consequently, a larger proportion of TKE would remain in the streamwise component, resulting in smoother streak structures. Figure 12 illustrates the instantaneous distributions of u' in the near-wall region. The controlled cases exhibit significantly smoother near-wall streaks compared to the uncontrolled results, and this trend is consistent across different Reynolds numbers. This observation aligns with the suppressed near-wall self-sustaining mechanism, and also reflects the DRL model's influence on Φ_{22} from the perspective of flow structures.

Moreover, the streamwise velocity fluctuations merit further discussion due to their significant impact from the DRL-based control strategy, as illustrated in figure 4. The transport equation of the streamwise kinetic energy $\langle u'u' \rangle$ is expressed as

$$\frac{\partial}{\partial t} \langle u'u' \rangle = \underbrace{-2 \langle u'v' \rangle \frac{dU}{dy}}_{P_{11}} + \underbrace{\left(-\frac{d}{dy} \langle u'u'v' \rangle \right)}_{D_{11,t}} + \underbrace{\mu \frac{d^2 \langle u'u' \rangle}{dy^2}}_{D_{11,v}} \quad (3.5)$$

$$+ \underbrace{\frac{2}{\rho} \left\langle p' \frac{\partial u'}{\partial x} \right\rangle}_{\Phi_{11}} + \underbrace{\left(-2\mu \left\langle \frac{\partial u'}{\partial x_j} \frac{\partial u'}{\partial x_j} \right\rangle \right)}_{\varepsilon_{11}},$$

where P_{11} is the turbulent production, $D_{11,t}$ is the turbulent diffusion, $D_{11,v}$ is the viscous diffusion, Φ_{11} is the redistribution, and ε_{11} is the dissipation.

Figure 13 illustrates the wall-normal distributions of the budget terms on the right-hand side of (3.5). Among these budget terms, the production P_{11} plays a crucial role, consistently remaining positive across various heights. In contrast, the terms $D_{11,t}$ and $D_{11,v}$ are relatively smaller, and represent primarily the transport of TKE in the wall-normal direction. The dissipation term ε_{11} remains negative at different heights, counterbalancing the production P_{11} . Additionally, unlike Φ_{22} associated with wall-normal velocity fluctuations, the redistribution term Φ_{11} has a trivial impact on the streamwise TKE. In the uncontrolled cases, the production term P_{11} gradually increases with height, reaching a peak at approximately $y^+ = 10\text{--}15$. This height is similar to the u_{rms} peak in the near-wall region, as depicted in figure 4. After applying control, this peak disappears and transforms into a trough. The production term P_{11} in the viscous sublayer increases significantly compared with the uncontrolled case. In the buffer layer, P_{11} initially decreases and reaches the trough, then increases with height, with a second peak appearing at approximately $y^+ = 20$. It can be observed that after applying the DRL model, the trend of P_{11} changes in a manner highly consistent with the changes in u_{rms} . Notably, the turbulent production P_{11} of the streamwise TKE consists mainly of two parts: $\langle -u'v' \rangle$ and dU/dy . The latter can be considered as a result of changes in Reynolds stress. Therefore, we can infer that the significant changes in u_{rms} caused by the DRL-based control are primarily due to alterations in Reynolds shear stress.

In summary, figure 14 illustrates the dynamic mechanism through which DRL-based control strategies influence skin friction. The application of wall blowing and suction, directed by the DRL models, effectively suppresses the near-wall self-sustaining process, thereby leading to smoother velocity streaks. This suppression manifests as a decrease in the redistribution term of wall-normal TKE within the buffer layer, consequently reducing wall-normal velocity fluctuations. The reduction in $\langle v'v' \rangle$ further diminishes the production term of Reynolds stress, resulting in a decrease in $\langle -u'v' \rangle$. Ultimately, this decline in Reynolds stress results in a reduction of skin friction. Moreover, the weakening of streamwise velocity fluctuations can also be attributed to the decrease in Reynolds stress. When the range of blowing and suction velocities is expanded, the aforementioned effects are amplified, leading to an increase in the extent of drag reduction. Conversely, as the Reynolds number rises and drag reduction diminishes, these trends are reversed.

4. Summary and conclusions

This study employs deep reinforcement learning (DRL) to develop control strategies, aimed at reducing skin friction in DNS of turbulent channel flows at high Reynolds numbers. Utilizing the TD3 framework, DRL predictions regulated wall blowing and suction velocities, with streamwise velocity fluctuations at $y^+ = 15$ serving as the state input for the DRL agent.

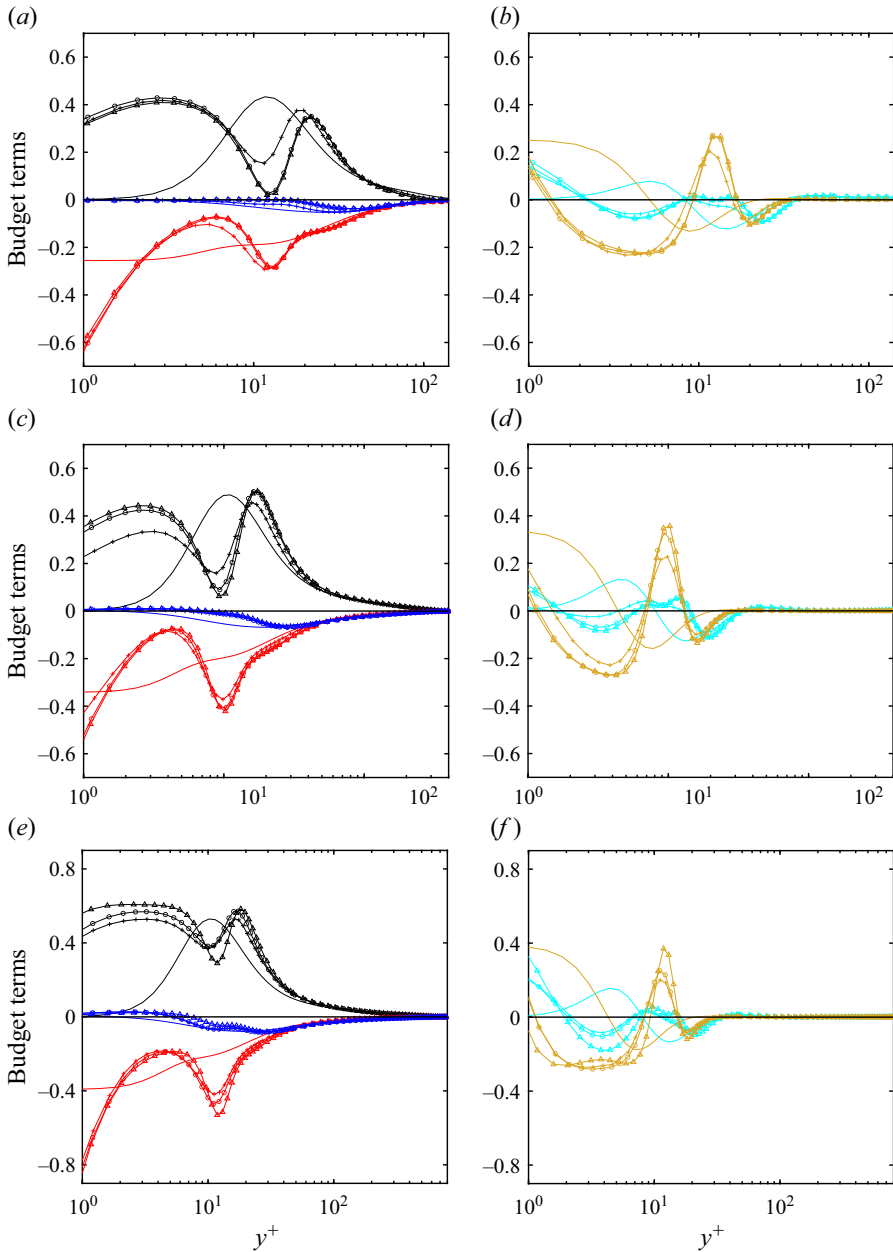


Figure 13. Wall-normal distributions of the budget terms of streamwise kinetic energy $\langle u'u' \rangle$ in (3.5): — P_{11} , — Φ_{11} , — ε_{11} , — $D_{11,t}$, — $D_{11,v}$, for (a,b) C180, (c,d) C550, (e,f) C1000. Lines without markers indicate cases with suffix 0; plus signs indicate cases with suffix 1; circles indicate cases with suffix 2; triangles indicate cases with suffix 3.

The DRL-based control strategies achieved significant drag reduction across various Reynolds numbers, with maximum reduction rates 35.6% at $Re_\tau \approx 180$, 30.4% at $Re_\tau \approx 550$, and 27.7% at $Re_\tau \approx 1000$. These results demonstrate superior drag reduction compared to traditional opposition control. As the range of blowing and suction velocities

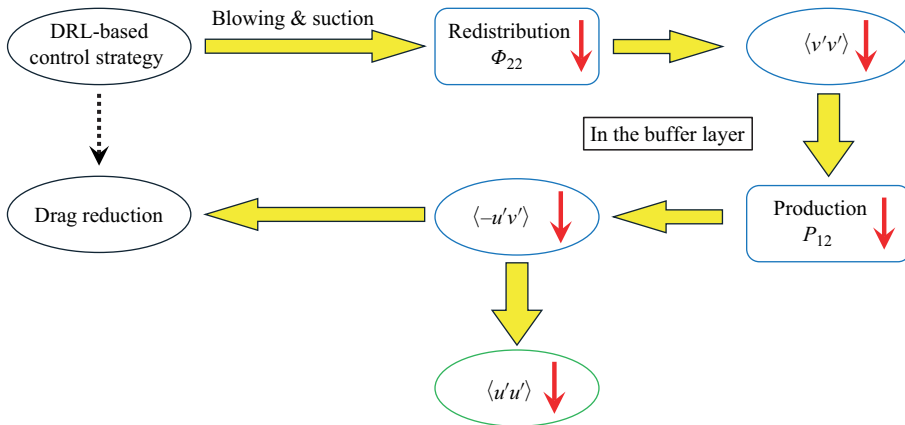


Figure 14. Schematic diagram of the dynamic mechanism through which DRL-based control strategies influence skin friction.

was extended, the drag reduction rates improved. Conversely, the effectiveness of DRL-based control decreased with higher Reynolds numbers, similar to opposition control methods. Further statistics indicate that the impact of DRL-based control on velocity fluctuations is limited to the near-wall region, with minimal effects on the outer region. Unlike opposition control, the wall blowing and suction velocities are more strongly correlated with the near-wall streamwise velocity fluctuations compared to the wall-normal component, owing to the u' input state of the DRL model.

According to the virtual wall theory (Hammond *et al.* 1998), the height of the virtual wall and the residual Reynolds stress on it are key indicators of drag reduction from a structural kinematic perspective. Compared to opposition control, the DRL model achieves higher drag reduction by elevating the virtual wall through blowing and suction. When the range of these actions is expanded, the virtual wall height increases, and residual Reynolds stress decreases, leading to further drag reduction. In contrast, an increase in Reynolds number significantly raises residual Reynolds stress, disrupting the virtual wall's effectiveness and resulting in decreased drag reduction rates. The contribution of residual Reynolds stress arises mainly from the amplitude modulation of large-scale structures, rather than the superposition effect. The footprint of outer large-scale structures is blocked above the virtual wall, while residual fluctuations on the virtual wall manifest as clusters of small-scale structures after DRL control. These small-scale fluctuations tend to be distributed beneath large-scale high-speed regions, indicating that the virtual wall's blockage is disrupted mainly in these areas.

On the other hand, analysing the budget equations elucidates the dynamic mechanisms through which DRL-based control strategies impact skin friction. Our observations indicate that the DRL models primarily reduce skin friction by inhibiting the redistribution term of wall-normal turbulent kinetic energy. This effect manifests as the suppression of the near-wall self-sustaining mechanism, resulting in smoother near-wall streaks. The reduction in the redistribution term leads to decreased wall-normal velocity fluctuations in the buffer layer, thereby diminishing the turbulent production of Reynolds stress. This chain of effects further weakens the Reynolds shear stress, ultimately reducing skin friction. Notably, when the range of blowing and suction velocities is extended, these effects are amplified, leading to even greater drag reduction. Conversely, an increase in the Reynolds number has the opposite effect, counteracting the benefits provided by the DRL-based strategies.

Cases	Input states	Rewards	Re_τ	DR (%)
C550-3	$u' _{y^+=15}$	TKE reduction rate	454.0	30.4
C550-drag	$u' _{y^+=15}$	Drag reduction rate	454.4	30.3
C550-v15	$v' _{y^+=15}$	TKE reduction rate	459.2	28.8
C550-u20	$u' _{y^+=20}$	TKE reduction rate	471.7	24.9

Table 5. Drag reduction results under different input states and rewards.

Funding. We gratefully acknowledge financial support from the Max Planck Society, the German Research Foundation (DFG) from grants 521319293, 540422505 and 550262949, and the Daimler and Benz foundation. We also thank the HPC systems of the Max Planck Computing and Data Facility (MPCDF) for the allocation of computational time. The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time on the GCS Supercomputers SuperMUC-NG at the Leibniz Supercomputing Centre (www.lrz.de) and JUWELS (Julich Supercomputing Centre 2021) at the Julich Supercomputing Centre (JSC). We gratefully acknowledge the grant WBS A-8001172-00-00 from the Ministry of Education, Singapore.

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

Appendix A

This appendix discusses the drag reduction performance of the DRL-driven control strategy based on varied input states and reward, as well as the performance of the trained control strategy across different resolutions and Reynolds numbers.

To evaluate the drag reduction performance of the DRL-driven control strategy under varied input states and rewards, we conduct three test cases, summarized in [table 5](#). These cases are all based on C550-3, with consistent parameters except for changes in input states and rewards. In C550-3-drag, the reward is modified to drag reduction rate, contrasting with C550-3; in C550-3-v15, the input state is changed to v' at $y^+ = 15$; and in C550-3-u20, the input state is changed to u' at a higher position $y^+ = 20$. We observe that each modified case converged within 10 episodes, and we select the models at 20 episodes, consistent with C550-3. The drag reduction results are presented in [table 5](#). The drag reduction rate in C550-3-drag nearly collapses with C550-3, suggesting that similar control performance can be achieved regardless of whether TKE or total drag is used as the optimization function. In C550-3-v15, where wall-normal velocity fluctuations are used as the input state, a slight decrease in drag reduction rate is observed. In C550-3-u20, with streamwise velocity fluctuations at a higher position, Re_τ increases, accompanied by a smaller drag reduction rate. These two cases indicate that despite variations in sensing parameters, the DRL strategy remains effective in developing flow control strategies based on the selected input variables.

Furthermore, we tested the performance of the trained control strategy across different resolutions and Reynolds numbers. In the first case, we applied the model trained in C550-3 to control a flow field around $Re_\tau \approx 550$ with the streamwise and spanwise grids refined by a factor of 2. We found that after grid refinement, the drag reduction rate collapsed with the result from case C550-3. This suggests that grid resolution has a trivial effect on the drag reduction performance of the DRL-derived control policy. In the second case, we applied the model trained in case C1000-3 to control a flow field around $Re_\tau \approx 550$. The resulting drag reduction rate was 28.9%, only slightly lower than the 30.4% achieved in C550-3. This indicates that the control policy trained in C1000-3 remains effective when the Reynolds number is reduced.

- BALAKUMAR, B.J. & ADRIAN, R.J. 2007 Large- and very-large-scale motions in channel and boundary-layer flows. *Phil. Trans. R. Soc. Lond. A: Math. Phys. Engng Sci.* **365** (1852), 665–681.
- BEWLEY, T.R., MOIN, P. & TEMAM, R. 2001 DNS-based predictive control of turbulence: an optimal benchmark for feedback algorithms. *J. Fluid Mech.* **447**, 179–225.
- BRUNTON, S.L. & NOACK, B.R. 2015 Closed-loop turbulence control: progress and challenges. *Appl. Mech. Rev.* **67** (5), 050801.
- CHANG, Y., COLLIS, S.S. & RAMAKRISHNAN, S. 2002 Viscous effects in control of near-wall turbulence. *Phys. Fluids* **14** (11), 4069–4080.
- CHOI, H., MOIN, P. & KIM, J. 1994 Active turbulence control for drag reduction in wall-bounded flows. *J. Fluid Mech.* **262**, 75–110.
- CHUNG, Y.M. & TALHA, T. 2011 Effectiveness of active flow control for turbulent skin friction drag reduction. *Phys. Fluids* **23** (2), 025102.
- DEL ÁLAMO, J.C. & JIMÉNEZ, J. 2003 Spectra of the very large anisotropic scales in turbulent channels. *Phys. Fluids* **15** (6), L41–L44.
- DEL ÁLAMO, J.C., JIMENEZ P., ZANDONADE, P. & MOSER, R.D. 2004 Scaling of the energy spectra of turbulent channels. *J. Fluid Mech.* **500**, 135–144.
- DENG, B.-Q. & XU, C.-X. 2012 Influence of active control on STG-based generation of streamwise vortices in near-wall turbulence. *J. Fluid Mech.* **710**, 234–259.
- FUKAGATA, K., IWAMOTO, K. & KASAGI, N. 2002 Contribution of Reynolds stress distribution to the skin friction in wall-bounded flows. *Phys. Fluids* **14** (11), L73–L76.
- FUKAGATA, K. & KASAGI, N. 2004 Suboptimal control for drag reduction via suppression of near-wall Reynolds shear stress. *Intl. J. Heat Fluid Flow* **25** (3), 341–350.
- GAD-EL-HAK, M. & BLACKWELDER, R.F. 1989 Selective suction for controlling bursting events in a boundary layer. *AIAA J.* **27** (3), 308–314.
- GUALA, M., HOMMEMA, S.E. & ADRIAN, R.J. 2006 Large-scale and very-large-scale motions in turbulent pipe flow. *J. Fluid Mech.* **554**, 521–542.
- GUASTONI, L., RABAULT, J., SCHLATTER, P., AZIZPOUR, H. & VINUESA, R. 2023 Deep reinforcement learning for turbulent drag reduction in channel flows. *Eur. Phys. J. E* **46** (4), 27.
- GUÉNIAT, F., MATHELIN, L. & HUSSAINI, M.Y. 2016 A statistical learning strategy for closed-loop control of fluid flows. *Theor. Comput. Fluid Dyn.* **30** (6), 497–510.
- HAMILTON, J., KIM, J. & WALEFFE, F. 1995 Regeneration mechanisms of near-wall turbulence structures. *J. Fluid Mech.* **287**, 317–348.
- HAMMOND, E.P., BEWLEY, T.R. & MOIN, P. 1998 Observed mechanisms for turbulence attenuation and enhancement in opposition-controlled wall-bounded flows. *Phys. Fluids* **10** (9), 2421–2423.
- HAN, B.-Z. & HUANG, W.-X. 2020 Active control for drag reduction of turbulent channel flow based on convolutional neural networks. *Phys. Fluids* **32** (9), 095108.
- HASEGAWA, Y. & KASAGI, N. 2011 Dissimilar control of momentum and heat transfer in a fully developed turbulent channel flow. *J. Fluid Mech.* **683**, 57–93.
- HOYAS, S. & JIMÉNEZ, J. 2006 Scaling of the velocity fluctuations in turbulent channels up to $Re_\tau = 2003$. *Phys. Fluids* **18**, 011702.
- HUTCHINS, N. & MARUSIC, I. 2007a Evidence of very long meandering features in the logarithmic region of turbulent boundary layers. *J. Fluid Mech.* **579**, 1–28.
- HUTCHINS, N. & MARUSIC, I. 2007b Large-scale influences in near-wall turbulence. *Phil. Trans. R. Soc. Lond. A: Math. Phys. Engng Sci.* **365** (1852), 647–664.
- HWANG, Y. 2013 Near-wall turbulent fluctuations in the absence of wide outer motions. *J. Fluid Mech.* **723**, 264–288.
- IWAMOTO, K., SUZUKI, Y. & KASAGI, N. 2002 Reynolds number effect on wall turbulence: toward effective feedback control. *Intl. J. Heat Fluid Flow* **23** (5), 678–689.
- JIMÉNEZ, J. 1998 The largest scales of turbulence. In *CTR Annual Research Briefs*, pp. 137–154. Stanford University.
- JIMÉNEZ, J. & MOIN, P. 1991 The minimal flow unit in near-wall turbulence. *J. Fluid Mech.* **225**, 213–240.
- JIMÉNEZ, J. & PINELLI, A. 1999 The autonomous cycle of near-wall turbulence. *J. Fluid Mech.* **389**, 335–359.
- JORDAN, M.I. & MITCHELL, T.M. 2015 Machine learning: trends, perspectives, and prospects. *Science* **349** (6245), 255–260.
- Jülich Supercomputing Centre. 2021 JUWELS cluster and booster: exascale pathfinder with modular supercomputing architecture at Juelich Supercomputing Centre, *J. Large-Scale Res. Facilities* **7**, A138.
- KIM, K.C. & ADRIAN, R.J. 1999 Very large-scale motion in the outer layer. *Phys. Fluids* **11** (2), 417–422.

- KRAVCHENKO, A.G., CHOI, H. & MOIN, P. 1993 On the relation of near-wall streamwise vortices to wall skin friction in turbulent boundary layers. *Phys. Fluids A* **5** (12), 3307–3309.
- LEE, C., KIM, J., BABCOCK, D. & GOODMAN, R. 1997 Application of neural networks to turbulence control for drag reduction. *Phys. Fluids* **9** (6), 1740–1747.
- LEE, C., KIM, J. & CHOI, H. 1998 Suboptimal control of turbulent channel flow for drag reduction. *J. Fluid Mech.* **358**, 245–258.
- LEE, M. & MOSER, R.D. 2019 Spectral analysis of the budget equation in turbulent channel flows at high Reynolds number. *J. Fluid Mech.* **860**, 886–938.
- LEE, T., KIM, J. & LEE, C. 2023 Turbulence control for drag reduction through deep reinforcement learning. *Phys. Rev. Fluids* **8** (2), 024604.
- LI, J. & ZHANG, M. 2022 Reinforcement-learning-based control of confined cylinder wakes with stability analyses. *J. Fluid Mech.* **932**, A44.
- LILLICRAP, T.P., HUNT, J.J., PRITZEL, A., HEESS, N., EREZ, T., TASSA, Y., SILVER, D. & WIERSTRA, D. 2015 Continuous control with deep reinforcement learning. [arXiv:1509.02971](https://arxiv.org/abs/1509.02971)
- MARUSIC, I., MATHIS, R. & HUTCHINS, N. 2010a High Reynolds number effects in wall turbulence. *Intl. J. Heat Fluid Flow* **31** (3), 418–428.
- MARUSIC, I., MATHIS, R. & HUTCHINS, N. 2010b Predictive model for wall-bounded turbulent flow. *Science* **329** (5988), 193–196.
- MATHIS, R., HUTCHINS, N. & MARUSIC, I. 2009 Large-scale amplitude modulation of the small-scale structures in turbulent boundary layers. *J. Fluid Mech.* **628**, 311–337.
- MATHIS, R., HUTCHINS, N. & MARUSIC, I. 2011 A predictive inner–outer model for streamwise turbulence statistics in wall-bounded flows. *J. Fluid Mech.* **681**, 537–566.
- MONTY, J.P., HUTCHINS, N., NG, H.C.H., MARUSIC, I. & CHONG, M.S. 2009 A comparison of turbulent pipe, channel and boundary layer flows. *J. Fluid Mech.* **632**, 431–442.
- PAMIÈS, M., GARNIER, E., MERLEN, A. & SAGAUT, P. 2007 Response of a spatially developing turbulent boundary layer to active control strategies in the framework of opposition control. *Phys. Fluids* **19** (10), 108102.
- PARIS, R., BENEDDINE, S. & DANDOIS, J. 2021 Robust flow control and optimal sensor placement using deep reinforcement learning. *J. Fluid Mech.* **913**, A25.
- RABAULT, J., KUCHTA, M., JENSEN, A., RÉGLADE, U. & CERARDI, N. 2019 Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control. *J. Fluid Mech.* **865**, 281–302.
- SONODA, T., LIU, Z., ITOH, T. & HASEGAWA, Y. 2023 Reinforcement learning of control strategies for reducing skin friction drag in a fully developed turbulent channel flow. *J. Fluid Mech.* **960**, A30.
- SUÁREZ, P., ÁLCANTARA-ÁVILA, A., FRANCISCO, V., JEAN, R., ARNAU, M., BERNAT, F., LEHMKUHL, O. & VINUESA, R. 2024 Flow control of three-dimensional cylinders transitioning to turbulence via multi-agent reinforcement learning. [arXiv:2405.17210](https://arxiv.org/abs/2405.17210)
- TOUBER, E. & LESCHZINER, M.A. 2012 Near-wall streak modification by spanwise oscillatory wall motion and drag-reduction mechanisms. *J. Fluid Mech.* **693**, 150–200.
- VAN DER POEL, E.P., ERWIN, P., OSTILLA-MÓNICO, R., DONNERS, J. & VERZICCO, R. 2015 A pencil distributed finite difference code for strongly turbulent wall-bounded flows. *Comput. Fluids* **116**, 10–16.
- VARELA, P., SUÁREZ, P., ÁLCANTARA-ÁVILA, F., FRANCISCO, M., ARNAU, R., JEAN, F., BERNAT, G.-C., MIGUEL, L., LEHMKUHL, O. & VINUESA, R. 2022 Deep reinforcement learning for flow control exploits different physics for increasing Reynolds number regimes. *Actuators* **11**, 359.
- VERZICCO, R. & ORLANDI, P. 1996 A finite-difference scheme for three-dimensional incompressible flows in cylindrical coordinates. *J. Comput. Phys.* **123** (2), 402–414.
- XU, C.-X. & HUANG, W.-X. 2005 Transient response of Reynolds stress transport to spanwise wall oscillation in a turbulent channel flow. *Phys. Fluids* **17** (1), 018101.
- ZENG, K. & GRAHAM, M.D. 2021 Symmetry reduction for deep reinforcement learning active control of chaotic spatiotemporal dynamics. *Phys. Rev. E* **104** (1), 014210.
- ZHU, X. *et al.* 2018 AFiD-GPU: a versatile Navier–Stokes solver for wall-bounded turbulent flows on GPU clusters. *Comput. Phys. Commun.* **229**, 199–210.