1st Spanish Fluid Mechanics Conference
Cádiz, June 19-22, 2022

Active drag reduction in minimal flow units via deep reinforcement learning

Alejandro Güemes,[1] Alejandro Martín,[1] Rodrigo Castellanos,[1, 2] Oscar Flores,[1] and Stefano Discetti[1, *]

[1] *Aerospace Engineering Research Group, Universidad Carlos III de Madrid, Spain*
[2] *Theoretical and Computational Aerodynamics Branch, Flight Physics Department, Spanish National Institute for Aerospace Technology (INTA), Spain*

This contribution explores the suitability of deep reinforcement learning (DRL) as an active control strategy to reduce the skin friction in wall-bounded flows. For that purpose, experiments are carried out on a direct numerical simulation of a channel flow at $Re_\tau \approx 200$ with a small computational domain. The actuation is simulated as a volumetric force applied in the center of one of the channel walls. The DRL agent is an artificial neural network fed by wall-shear stress and pressure sensors. The agent learns an active control strategy tuning the volumetric force according to the upstream flow conditions. Preliminary results show modest skin friction reduction, consistent with an opposition control strategy.

Keywords: wall turbulence, deep reinforcement learning, skin friction, flow control.

INTRODUCTION

The quest for efficient control strategies of wall-bounded turbulent flows has fueled abundant research over the last decades due to its relevance in countless industrial and aeronautical devices. Among others, opposition control has demonstrated being a simple, yet effective, closed-loop flow-control solution to tamper with near-wall structures [1, 2], as well as with the large-scale features present in the logarithmic layer [3].

Recent advances in machine learning (ML) are raising the question on whether such methods can be efficient for flow control. The challenge is dual: on the one hand we search for effective control laws, even if in black-box form; on the other, there is the chance that, interpreting what is inside the black box, we might discover new intriguing strategies to control the flow. Among ML control methods, deep reinforcement learning (DRL) is certainly one of the most promising. DRL has been developed to identify active control policies for complex dynamical systems [4]. The DRL agent, controlled by a deep neural network (DNN), performs actions on an environment according to its state. The actions are graded with a reward (based on the goal to be achieved and, often, on the cost of the action), that is used to progressively learn an optimized control strategy. In the area of fluid mechanics, DRL has been successfully applied in several applications, such as controlling the wake behind a cylinder [5].

The capability of DRL to reduce the wall-shear stress in a turbulent channel flow is explored in this work. The environment is a minimal flow unit [6], which ensures a good representation of the structures populating the near-wall region and their dynamics.

METHODOLOGY

A direct numerical simulation (DNS) of a turbulent channel flow based on the numerical implementation presented by Vela-Martín *et al.* [7], sets the environment to evaluate the performances of DRL in wall-bounded-flow control. The friction Reynolds number is $Re_\tau = u_\tau h/\nu \approx 200$, being u_τ the friction velocity, h the half channel height and ν the kinematic viscosity. The size of the computational box in the streamwise and spanwise directions is πh and $\pi h/2$ respectively, which, together with Re_τ, provides a good representation of the structures and the dynamics of the near-wall region [6]. The agent is trained using the proximal policy optimization (PPO) algorithm [4] to control a volumetric force of characteristic size L_f and duration T_f applied at the bottom wall, as defined by Pastor *et al.* [2]. The agent reads shear and pressure information from nine sensors placed on the bottom wall, as inputs, and provides the values for the forcing intensity f_0, as output. The reward function driving the learning process is defined as the temporal mean of the difference in instantaneous skin friction at the lower wall between the controlled and uncontrolled simulation. The skin-friction sensor for the reward computation covers a region over the entire streamwise direction at the bottom wall and spanning $1.2h$ in the spanwise direction.

The training procedure consists in running two parallel simulations with the same initial condition and allowing the agent to actuate only in one of them. Comparison of the wall shear in the two simulations allows a direct evaluation of the effect of the actuation. The control command consists of a single volumetric force action generated by the agent. After each actuation, the forced simulation is discarded and replaced by the non-forced case. In this way, the agent learns how to act on a well-developed turbulent flow without reminiscent effect of the previous control action due to the periodicity of the channel flow. The coefficients of the DNN driving the agent are adjusted every 8 episodes based on the effectiveness of the control in reducing skin-friction.

RESULTS

Preliminary results are presented for a training set of 2800 episodes of duration $T_{epi} \approx 0.55h/u_\tau$ with forcing duration $T_f = 0.25h/u_\tau$ and size $L_f = 0.2h$. Figure 1 shows the averaged initial states of wall-normal velocity pressure and shear conditioned to $\Delta\tau > 0.001\tau_0$, from which it can be deduced that the agent learns to operate with an opposition control strategy.

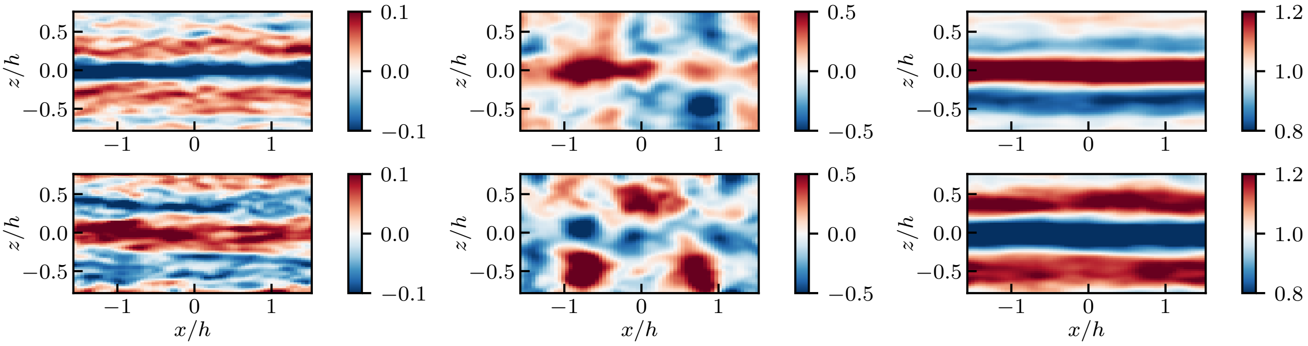

FIG. 1. Fluid state before actuation conditioned to $\Delta\tau > 0.01\tau_0$ and a specific sign of f_0. Top row, up-ward forcing, $f_0 > 0$. Bottom row, downward forcing, $f_0 < 0$. Left column, vertical velocity at $y^+ = 15$. Centre column, wall pressure. Right column, wall shear. Velocity is normalised with u_τ, while τ_0 is used for pressure and shear.

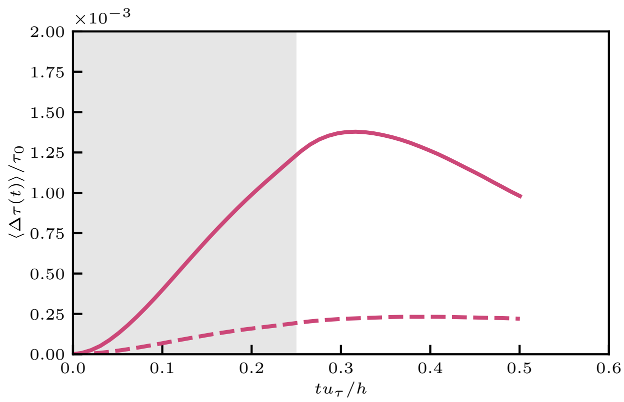

FIG. 2. Ensemble mean $\langle\rangle$ of $\Delta\tau(t)$ versus time, normalised with the mean skin friction of the base flow. Lines indicate $\Delta\tau$ computed over (solid) reward region and (dashed) entire bottom wall. Gray region indicates the period in which the agent acts.

When $f_0 > 0$ results in a skin-friction reduction, the initial state is characterised by a Q2 event, with the vertical velocity going towards the wall. Conversely, for $f_0 < 0$ the typical characteristics of Q4 events can be observed.

Figure 2 presents the ensemble mean of $\Delta\tau(t)$ as a function of time for both the reward region and the entire lower wall. Note that a single actuator covering 0.6% of the wall area reduces the skin friction by a factor up to 0.15% in the reward region (0.01% in the entire wall) with respect to the non-forced case. The actuation ratio between positive and negative rewards is 1.42.

CONCLUSIONS

The preliminary results hints to the suitability of DRL strategies for the active control of wall-bounded flows for skin friction reduction. Although a deeper analysis is required to account for the different parameters involved in the actuation, the reported results pave the way toward model-free active control of wall-bounded turbulence via DRL. By the time of the conference, it is expected that this analysis will be extended to account for both the effect of episode duration and the width of the reward sensor. If the agent maintains its control capacity, its range of action can be increased by controlling the duration of the applied volumetric force.

ACKNOWLEDGEMENT

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 949085) and from the BBVA Foundation under the "Becas Leonardo a Investigadores y Creadores Culturales" programme (Grant No. IN[20]_ING_ING_0163).

* sdiscett@ing.uc3m.es

[1] H. Choi, P. Moin, and J. Kim, Active turbulence control for drag reduction in wall-bounded flows, J. Fluid Mech. **262**, 75 (1994).
[2] R. Pastor, A. Vela-Martin, and O. Flores, Wall-bounded turbulence control: statistical characterisation of actions/states, in *J. Phys. Conf. Ser.*, Vol. 1522 (2020) p. 012014.
[3] A. Guseva and J. Jiménez, Linear instability and resonance effects in large-scale opposition flow control, J. Fluid Mech. **935** (2022).
[4] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, Proximal policy optimization algorithms, arXiv preprint arXiv:1707.06347 (2017).
[5] J. Rabault, M. Kuchta, A. Jensen, U. Réglade, and N. Cerardi, Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control, J. Fluid Mech. **865**, 281 (2019).
[6] J. Jiménez and P. Moin, The minimal flow unit in near-wall turbulence, J. Fluid Mech. **225**, 213 (1991).
[7] A. Vela-Martín, M. P. Encinar, A. García-Gutiérrez, and J. Jiménez, A low-storage method consistent with second-order statistics for time-resolved databases of turbulent channel flow up to $Re_\tau = 5300$, J. Comput. Sci. **56**, 101476 (2021).