

Review

Flow control in wings and discovery of novel approaches via deep reinforcement learning

R. Vinuesa^{1,*}, O. Lehmkuhl², A. Lozano-Durán³ and J. Rabault^{4,5}¹ FLOW, Engineering Mechanics, KTH Royal Institute of Technology, Stockholm, Sweden.² Barcelona Supercomputing Center, Barcelona, Spain.³ Dept. of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA.⁴ IT Department, Norwegian Meteorological Institute, Oslo, Norway.⁵ Independent researcher.

* Email for correspondence: rvinuesa@mech.kth.se

Abstract: In this review we summarize existing trends of flow control used to improve the aerodynamic efficiency of wings. We first discuss active methods to control turbulence, starting with flat-plate geometries and building towards the more complicated flow around wings. Then, we discuss active approaches to control separation, a crucial aspect towards achieving high aerodynamic efficiency. Furthermore, we highlight methods relying on turbulence simulation, and discuss various levels of modelling. Finally, we thoroughly revise data-driven methods, their application to flow control, and focus on deep reinforcement learning (DRL). We conclude that this methodology has the potential to discover novel control strategies in complex turbulent flows of aerodynamic relevance.

Keywords: turbulence; flow control; simulation; aerodynamics; machine learning; deep reinforcement learning

1. Introduction

Over the past decades, aviation has become an essential component of today's globalized world: before the current pandemic of coronavirus disease 2019 (COVID-19), over 100,000 flights took off everyday allowing the transportation of people and goods, and the establishment of global commercial relations. Despite the significant impact of the pandemic on the aviation sector, a number of studies indicate that after the pandemic its relevance in the transportation mix will be similar to that before COVID-19 [1,2]. Aviation alone is responsible for 12% of the carbon dioxide emissions from the whole transportation sector, and for 3% of the total CO₂ emissions in the world [3]. On the other hand, fuel represents around 40% of the costs in regular airlines, corresponding to hundreds of millions of dollars spent yearly [4]. This has important implications in the context of the Sustainable Development Goals (SDGs) 9 (on sustainable infrastructure) and 11 (on sustainable cities), as well as an important impact on SDGs 3 (on health) and 13 (on climate change) from the United Nations (UN) 2030 Agenda [5–7]. Due to the major environmental and economical impacts associated with aviation, it is desirable to improve the aerodynamic performance of airplane wings, with the aim of reducing the fuel consumption and emissions associated with air travel.

In order to develop more efficient wings it is necessary to reduce the losses associated with their movement within the surrounding fluid. This implies reducing the force parallel to the incoming flow (the drag), and one of the strategies to achieve such a reduction is to perform flow control. A wide range of methods aimed at controlling the flow to reduce the drag have been reported and some have documented net-energy savings, *i.e.* taking into account the energy spent on the control, as documented by Fahland *et al.* [8]. These strategies include *passive methods*, such as riblets [9], which are drag-reducing surfaces proven successful in passenger aircraft [10], and *active techniques*, in which the drag-reduction effect is achieved through an action which requires additional energy to be transferred to the flow [11]. When the control action is determined based on the instantaneous flow

state, sensing is required and the control strategy is denoted as *reactive* [12], whereas in *predetermined* strategies the control is decided a-priori [13]. Predetermined strategies are appealing due to their simple implementation; however, reactive methods are the ones which have the highest potential to reduce drag. Note that currently-available reactive controls, such as opposition control [14,15] through blowing and suction at the wall, are simply based on an ad-hoc gain to relate the sensed velocity fluctuation and the actuation (see §1.1 for more details). It is important to highlight that the total drag in subsonic transport aircraft can be decomposed essentially into four components, namely: friction drag, lift-induced drag (where the lift is the aerodynamic force perpendicular to the incoming stream), wave drag and parasite drag. Whereas the two latter only represent up to 10% of the total drag, the first and second are responsible for around 50% and 40% respectively in cruise conditions [16,17]. Although the continuous design optimization carried out over the past decades has led to reductions of the drag, new innovative approaches are required in order to obtain significant improvements in aerodynamic performance [18]. In the following, we will discuss in more detail different types of active control.

1.1. Active control of turbulent flows

A widely-used method of predetermined active flow control is uniform blowing/suction. The first wind-tunnel experiments using the micro-blowing technique (MBT) [19] reported that it is possible to achieve a significant drag reduction with relatively moderate blowing, as well as to have net-energy savings in full-scale applications. More recent studies have confirmed this possibility, investigating the effects of MBT on more complex geometries and on adverse-pressure-gradient (APG) turbulent boundary layers (TBLs). A detailed description of the MBT technique is provided by Hwang [20], and Kornilov [21] discusses more recent developments, particularly regarding experimental results. On the other hand, high-fidelity numerical simulations have been used to better characterize the interaction between control and wall-bounded turbulence. One of the first numerical studies investigating TBLs with blowing and suction is that of Park and Choi [22], who employed direct numerical simulations (DNSs, in which all turbulent scales are resolved) and considered a Reynolds number based on displacement thickness δ^* and freestream velocity U_∞ of $Re_{\delta^*} = 500$. Kametani and Fukagata [23] performed DNS of a zero-pressure-gradient (ZPG) TBL with blowing and suction at Reynolds numbers based on momentum thickness Re_θ between 300 and 700, with intensities up to 1% of U_∞ . They also analyzed the energy input associated with uniform blowing to estimate the upper bound of control efficiency, and they confirmed that it is theoretically possible to achieve net-energy savings. As expected, uniform suction has opposite effects. Later, other numerical simulations (more relevant to the full-scale conditions) have been performed at higher Reynolds numbers. For instance, Kametani *et al.* [24] carried out high-resolution large-eddy simulations (LESs, where only the smallest turbulent scales are modelled) of a ZPG TBL at $Re_\theta = 2,500$, considering blowing and suction with an intensity of 0.1% of the freestream velocity. These authors achieved more than 10% drag reduction despite the relatively low blowing intensity.

The numerical studies discussed above are focused on the description of the effect of blowing on a spatially-developing ZPG TBL, which is an idealized study case. Firstly, more realistic scenarios exhibit more complex turbulent flows [25], including as pressure-gradient TBLs [26] and finite aerodynamic bodies [27,28]. In these cases it is not trivial to generalize the control techniques. Secondly, the skin-friction reduction is beneficial in engineering applications only if it corresponds to a reduction of total drag (which includes additional components as discussed above), and/or to an improvement of the aerodynamic efficiency (defined as the lift-to-drag ratio L/D). For these reasons, the two following experimental studies on the effects of blowing and suction in airfoils are of particular relevance. We first discuss the work by Eto *et al.* [29], who applied a blowing intensity of 0.14% of U_∞ to the suction side of a Clark-Y airfoil at $Re_c = 1,600,000$ (where Re_c is the Reynolds number based on U_∞ and wing chord c). They observed a local reduction of the skin friction between 20 and 40%, but they also reported an increase of the total drag. On the other hand, Kornilov *et al.* [30] carried out experiments on a NACA0012 airfoil at $Re_c = 700,000$, applying blowing and suction over both sides of the airfoil. They




Figure 1. Instantaneous visualization of coherent vortical structures [34] around a NACA4412 wing section at $Re_c = 400,000$, colored by streamwise velocity ranging from (dark blue) -0.2 to (dark red) 1.7 . The yellow line indicates the extent of the control region (with uniform blowing on the suction side), whereas the red one denotes the tripping location [35,36]. Figure extracted from Ref. [32], with permission of the publisher (Springer Nature).

confirmed that blowing over the suction side does not reduce the total drag, but they also observed that blowing over the pressure side and suction over suction side have a beneficial effect, achieving a reduction of total drag of around 10%. This highlights the additional difficulty of performing control in wings, where the various contributions to the total drag are tightly coupled. One of the first high-fidelity simulations of turbulent wings with control was conducted by Vinuesa and Schlatter [31] in 2017. In that work, as well as in more recent studies [32], high-resolution LES was used to study the turbulent flow around a NACA4412 wing section, up to $Re_c = 400,000$, where different combinations of blowing and suction rates over the suction and pressure sides were applied. Using predetermined active flow control, they achieved a maximum increase of the aerodynamic efficiency of 11% [32]. In Figure 1 we show the effect of applying uniform blowing on the suction side of the wing section, which leads to an increase of boundary-layer thickness and turbulence activity far from the wall. This produces a reduction of the wall-shear stress and an increase of the pressure drag, leading to a higher total drag; note that the opposite holds for uniform suction [32]. Another interesting numerical work is that of Albers *et al.* [33], who performed LES on an airfoil at $Re_c = 400,000$ to assess the effect of transversal surface waves, reporting a drag reduction of 7.5%. As discussed below, it is possible to obtain more sophisticated and efficient control strategies by sensing the flow and exploiting all the available information on its state, *i.e.*, by performing reactive flow control.

A widely-used method of reactive control is the so-called opposition control, where suction and blowing are introduced at the wall with the aim of suppressing the sweep and ejection events in the near-wall region, so as to reduce the skin friction [14]. Essentially, the velocity imposed at the wall v_w should be opposite to the wall-normal velocity v at a certain sensing plane y_s , according to the equation:

$$v_w(x, 0, z, t) = -\alpha [v(x, y_s, z, t) - V(y_s, t)]. \quad (1)$$

Note that here x , y and z denote the streamwise, wall-normal and spanwise coordinates, t is the time, α is a positive constant and V is the instantaneous wall-normal velocity averaged over the control area. Subtracting this term ensures a zero-net-mass-flux condition at the wall. It is important to note that, through this equation, the control aims at opposing the fluctuations at a certain wall-normal location, which is typically around the near-wall fluctuation peak, *i.e.* $y_s^+ = 15$. The superscript '+' denotes inner scaling, in terms of the viscous length $\ell^* = \nu / u_\tau$, where ν is the fluid kinematic viscosity and $u_\tau = \sqrt{\tau_w / \rho}$ is the friction velocity (with τ_w being the wall-shear stress and ρ the fluid density). The

constant α is set empirically, which means that the resulting control law is relatively simple. Despite this simplicity, Stroh *et al.* [15] reported drag-reduction rates of around 20% in turbulent channels and boundary layers up to $Re_\tau \simeq 660$ (which is the friction Reynolds number, based on the 99% boundary-layer thickness δ_{99} and the friction velocity u_τ). However, it may be possible to obtain more sophisticated control laws by formulating an optimization problem, as discussed in §3.

1.2. Active control of separation

Several studies exploring the capabilities of active-flow-control (AFC) actuators in high-lift devices with massive separation (i.e. suction and blowing, sweeping jets, fluidic oscillators, plasma actuators, synthetic jets) have been conducted in the literature. A brief review on the state of the art of active flow control techniques for civil aircrafts can be found Batikh *et al.* [37]. Khodadoust and Washburn [38] conducted wind-tunnel measurements on a high-lift device fitted with AFC actuators. They observed that the application of small amount of suction and blowing increased the lift performance. Khün *et al.* [39] simulated a 3D high-lift wing with constant blowing using RANS. The results showed that blowing can be beneficial to suppress massive separation in the flap. Radespiel *et al.* [40] reviewed different techniques for AFC using constant blowing showing that tangential blowing can be promising at increasing the lift at high angles of attack. Fricke *et al.* [41] simulated the AFC by means of pulsed blowing to control flow separations in the wing engine junction with RANS. Later, Schloesser *et al.* [42] conducted experimental investigations in the same configuration. Their results showed that AFC successfully suppressed the flow separation with a lift increase and that results are independent on the Reynolds and Mach numbers. Hue *et al.* [43] using RANS simulated constant and pulsed blowing devices and observed gains up to 3% in the lift and the retard in separation due to the nacelle. Fluidic actuators placed at a tail of an aircraft were simulated using unsteady RANS and validated by means experimental results by Shmilovich *et al.* [44]. More recently, Andino *et al.* [45] tested fluidic actuators in a generic tail at low speeds and demonstrated that a modest increase of the momentum coefficient can result in important increments of the side force. Whalen *et al.* [46] presented wind-tunnel test results of the AFC of the vertical tail of a Boeing 757 equipped with sweeping jet actuators; significant increase in the side force at maximum rudder deflection of 30° was observed. Other works involving the use of sweeping jets can be found in Refs. [47,48]. The effectiveness of microjets at drag reduction was experimentally studied by Aley *et al.* [49] in a simplified 2D wing; a significant wake velocity deficit reduction and thus, drag was observed when using the microjets actuation.

To finalise, we focus on the application of synthetic jets with zero net mass flux as a promising technique for AFC of wings. In these devices the fluid necessary to alter the boundary layer is intermittently injected through an orifice driven by the motion of a diaphragm located on a sealed cavity below the surface [50]. Indeed, synthetic jets have been shown to succeed at reducing the fuel burnt during the operations of take-off and landing [51]. In the context of synthetic jets for AFC, there have been significant advances in the past years in airfoils (see for instance Refs. [52–55]). However, whether they can be implement on a full aircraft is still subject of investigation. Recently, Jabbal *et al.* [56] analysed different system architectures for AFC for real-size civil aircrafts in terms of efficiency, power requirements and integration issues. They concluded that synthetic jets might be useful to control separation in short-duration operations. Shmilovich and Yadlin [57] studied different AFC strategies of a high-lift profile in the conditions of take-off and landing using RANS. Bauer *et al.* [58] conducted experiments on a two element wing with unsteady AFC near the leading edge and showed that stall can be delayed. Lin *et al.* [59] addressed different strategies in the flap of a high-lift profile comprising steady suction and blowing, and periodic excitation of the boundary layer. Several of these AFC strategies are planned to be tested experimentally by NASA for increasing lift-to-drag ratios (L/D) in take-off configurations [60]. Although most of the numerical studies conducted so far have been performed using RANS, Jansen *et al.* [61] proved, by comparing experimental and numerical simulations, that delayed detached eddy simulations are useful to analyse the effect of a synthetic jet on the flow field of a tail at $Re_c = 350,000$. Finally, Lehmkuhl *et al.* [62] have studied the aerodynamic




Figure 2. JAXA Standard Model high-lift configuration at $Re_c = 1.93 \times 10^6$ and $AoA = 21.57^\circ$ with AFC. Two-dimensional streamlines at different spanwise locations: baseline case vs different actuator-jet angles Φ . (Top left) baseline; (top right) $\Phi = 0^\circ$; (bottom left) $\Phi = 60^\circ$; and (bottom right) $\Phi = 45^\circ$. Results extracted from Ref. [62], with permission of the publisher (IOP Publishing).

performance of active flow control on wings using synthetic jets with zero net-mass flow by means of wall-modeled large-eddy simulations, see Figure 2. The performance of synthetic jets was evaluated for the high-lift configuration of the JAXA Standard Model at realistic Reynolds numbers for landing $Re_c = 1.96 \times 10^6$. The results show that, at high angles of attack, the control successfully eliminates the laminar/turbulent recirculations located downstream of the actuator, thus increasing the aerodynamic performance.

2. Turbulence-simulation approaches

The use of computational fluid dynamics (CFD) for external-aerodynamic applications has been a key tool for aircraft design in the modern aerospace industry [63–65]. CFD methodologies with increasing functionality and performance have greatly improved our understanding and predictive capabilities of complex flows. These improvements suggest that design of novel and highly reliable control strategies via CFD may soon be a reality. The fully-virtual design of flow-control strategies is expected to limit the number of required wind-tunnel tests, reducing both the turnover time and cost of the design cycle [66,67]. However, flow predictions from the state-of-the-art CFD solvers are still unable to comply with the stringent accuracy requirements and computational efficiency demanded by the industry [68]. These limitations are imposed, largely, by the ubiquity of turbulence [69]. To tackle current challenges and encourage further advances in CFD, simulation of an aircraft configuration across the full flight envelope has been posed as one of the Grand Challenge Problems in the recent NASA CFD Vision 2030 [68].

From the early days of industrial CFD to present times, the treatment of turbulence has been mostly based on closure models for the Reynolds-averaged Navier–Stokes (RANS) equations. The approach appears in different flavors: from pure RANS solutions to Hybrid methods such as Detached Eddy Simulation and its variants [70,71]. In the latter, RANS is utilized close to the wall, whereas the outer layer is modeled via eddy-resolving methodologies. Many RANS models (and their variants) have been devised to overcome the limitations of their predecessors, usually by expanding and calibrating its coefficients to account for missing physics. Despite the reliance of RANS-based approaches on tunable parameters and empirical correlations, they have dominated the CFD industry for external aerodynamic applications, including commercial aviation [72].

The sophistication of RANS closure models has increased over time [71]. Yet, no practical model has emerged as a competent approach across the broad range of flow regimes of interest to

industry. The latter encompass separated flows, afterbodies, mean-flow three-dimensionality, shock waves, aerodynamic noise, fine-scale mixing, laminar-to-turbulent transition, etc. In these scenarios, RANS predictions tend to be inconsistent and unreliable, especially for geometries and conditions representative of the flight envelope of commercial airplanes. An example of such deficiencies is the prediction of the onset and extent of three-dimensional separated flow in wing-fuselage junctures, in which RANS-based approaches have shown poor performance [72,73]. RANS accuracy is also known to decline in aeroacoustic noise and vibration prediction for transonic airfoils [72]. Additional CFD experience in aircrafts at high angles of attack has revealed that RANS-based solvers have difficulty predicting maximum lift and the corresponding angle of attack along with the physical mechanisms for stall. This was highlighted in the 3rd AIAA CFD High Lift Prediction Workshop [74], where RANS solutions exhibited a significant scatter in lift, drag, and pitching moment near stall.

Recently, large-eddy simulation (LES) has gained momentum as a tool for both research and industrial applications. In LES, the large eddies containing most of the energy are directly resolved, while the dissipative effect of the small scales is accounted for by a subgrid-scale (SGS) model. Additionally, if the near-wall flow is also modeled (i.e., wall modeling) such that only the large-scale motions in the outer region of the boundary layer are resolved, the grid-point requirements for this wall-modeled LES (WMLES) scale at most linearly with increasing Reynolds number [75]. The cost-efficiency of WMLES and its demonstrated predictive capabilities over the last decade, make this approach a realistic contender to overcome the deficiencies of RANS-based methodologies.

Several strategies for modeling the near-wall region in LES have appeared in the literature, and comprehensive reviews can be found in Piomelli and Balaras [76], Cabot and Moin [77], Larsson *et al.* [78], and the most recent review by Bose and Park [79]. Most wall models utilize as input the LES solution at a given location in the LES domain, and return the wall heat and momentum fluxes needed by the LES solver. Among the most widespread approaches are those computing the wall stress using either the law of the wall [80–82] or simplified RANS equations [83–90], while recent advances in wall modeling are rooted on mathematical and physical principles completely free of RANS empiricism [91–93].

Advances in machine learning and data science have also incited new efforts to complement the existing turbulence-modeling approaches in the WMLES community. One of the first attempts at using supervised machine learning for WMLES can be found in Yang *et al.* [94], who proposed a physics-informed neural-network (PINN) model to predict the wall stress in turbulent channel flows. Recently, Lozano-Durán and Bae [95] formulated a wall model using building-block units (such as turbulent channel flows, ducts, and separation bubbles), which provides a classification of the flow and confidence in the prediction. The model was validated in a realistic aircraft with trailing-edge separation. Radhakrishnan *et al.* [96] formulated a wall model using gradient-boosted decision trees and predicted the wall-shear stress in a turbulent channel flow and a wall-mounted bump. Along these lines, Eivazi *et al.* [97] have recently reported efforts towards RANS modelling by means of PINNs.

According to NASA Vision 2030 report [68], hybrid RANS/LES and WMLES are identified as the most viable approaches for predicting realistic flows at high Reynolds numbers in external aerodynamics. As such, both hybrid RANS/LES and WMLES will be instrumental in the development of control strategies for realistic external-aerodynamic applications, as shown in Figure 3.

3. Data-driven methods for control and deep reinforcement learning

As discussed above, the flow around wings is very complex and it is difficult to devise efficient control strategies to optimize the aerodynamic efficiency, even having access to flow information in real time as in the case of opposition control. One approach to obtain more efficient control strategies is to formulate an optimization problem aimed at *e.g.* minimizing the drag or maximizing the aerodynamic efficiency. There have been several data-driven approaches to do this for flow control, for instance in the context of genetic programming [98]. Genetic programming (GP) is based on automatically choosing the terms in a symbolic equation through evolution and selection of the best candidates, a



Figure 3. Prediction of the wall-shear stress in the High-Lift Common Research Model at an angle of attack of 21 degrees computed using wall-modeled large-eddy simulation with the Alya solver on a grid with 40 million control elements.

fact that ensures the interpretability of this method [99,100], although the formula obtained can be deeply nested and complex. The GP approach has been successfully employed for control of external flows by Li *et al.* [101] and Minelli *et al.* [102]. Another interesting data-driven approach is Bayesian regression based on Gaussian processes [103], which was employed by Morita *et al.* [104] in CFD optimization, and by Mahfoze *et al.* [105] to identify the best combination of control-region length and blowing amplitude to maximize the energy savings, also including intermittent control regions. Note that these authors also took into account the data by Kornilov and Boiko [106] to formulate a more realistic estimate of the power consumption by blowing, and they reported a net-energy saving of around 5%. It is interesting to note that other data-driven methods may help to model the near-wall region and consequently may provide novel venues for improved flow control [107–111].

One very promising data-driven approach to flow control is deep reinforcement learning (DRL), which we will focus on in the following. In DRL, an agent (usually built based on a neural network, NN) interacts with an environment (the flow) in a closed loop. At each time t , the agent receives a partial observation of the environment o_t used to choose an action a_t , which will influence the evolution of the environment. The agent periodically receives a reward r_t , which indicates the quality of those actions under a certain norm. The goal of DRL is to find an optimal decision policy π from which the action is derived, i.e. $a_t = \pi(o_t)$, such that the cumulative reward is maximized. This process is summarized in Figure 4, and the goal of the DRL algorithm is to learn by interacting with the environment by gathering experience [112]. A pioneering work in this direction in the context of instability control in fluid mechanics was conducted by Rabault *et al.* [113], who used DRL to optimize the actuation from two jets on a two-dimensional cylinder flow. This resulted in significant drag-reduction rates using synthetic control jets blowing with a very low mass-flow rate intensity (typically a fraction of percent of the incoming mass flow rate intersecting the cylinder). We want to highlight that DRL is a very promising approach to discover novel and potentially more efficient control strategies going beyond classical control, since i) it does not make any assumptions on the properties of the system, except for the ability to establish a close-loop control and to extract a reward signal, ii) it takes advantage of the




Figure 4. Schematic representation of the deep-reinforcement-learning process employed for flow control. The objective is to find the optimal decision policy π such that the cumulative reward is maximized.

efficiency of NNs at representing complex, nonlinear functions following their universal-approximator property [114].

The work by Rabault *et al.* [113] employed the proximal-policy-optimization (PPO) algorithm [115], which is an Actor-Critic policy gradient algorithm. The PPO algorithm is simpler and faster than other similar techniques, such as the trust region policy optimization (TRPO) methods [115], and it requires relatively little hyper-parameter tuning. Furthermore, it is more suitable for continuous control than the deep-Q-network (DQN) learning [116], as well as its variations [117]. Offering a detailed overview of the PPO algorithm is beyond the scope of this review, and for a detailed discussion about the PPO algorithm, the reader is referred to either the initial PPO paper [115], or to a fluid-mechanics focused review of the DRL and PPO method [118,119]. However, the main lines of the PPO algorithm are as follows. The PPO method is episode-based, *i.e.*, it learns from performing active control for a limited amount of time before analyzing the obtained results and continuing with the learning process in a new episode. The learning problem is aimed at iteratively training (*i.e.*, finding the weights of) the policy network. Denoting the set of weights of the policy NN by Θ , the aim is, therefore, to maximize the long-term discounted reward function $R(t) = \sum_t \gamma^t r_t$, where γ is a discount factor (usually in the range [0.95 – 0.99]), formulated as finding:

$$R_{\max} = \max_{\Theta} \mathbb{E} \left[\sum_{t=0}^H (\gamma^t r_t \mid \pi_{\Theta}) \right], \quad (2)$$

where π_{Θ} is the policy function described by the neural network with weights Θ , and s_t is the (hidden) state of the system. In the present context, s_t would correspond to the complete flow information, whereas the limited observations o_t would be obtained from sensors. This maximization problem is solved by means of gradient descent performed on the weights Θ of the network, following experimental sampling of the system through interaction with the environment.

Following the initial work controlling the vortex shedding in a cylinder wake, a number of further refinements and applications have proven the potential of the method for the control of flow instabilities. Bucci *et al.* [120] successfully applied DRL to control chaotic systems, such as the Kuramoto–Sivashinsky equation. Paris *et al.* [121] investigated how sensors providing an overview of the state of the system to the DRL agent can be placed optimally. Beintema *et al.* [122] demonstrated efficient control of the Rayleigh Benard instability in a 2D channel. Tang *et al.* [123] proved through numerical simulations that DRL can be able to perform robust control over a range of inflow conditions. Xu *et al.* [124] investigated in a simulation how small counter-rotating cylinders can be used to reduce the drag behind a cylinder, while Fan *et al.* [125] provided an experimental demonstration of the technique. Finally, Ren *et al.* [126] pushed the value of the Reynolds number to a weakly turbulent regime and demonstrated that DRL can control fluid motion in the turbulent regime.

While research articles have mostly focused on relatively simple flow configurations so far, as these are the easiest to tackle computationally both for the CFD and the DRL agent, there are a number of possible refinements in the use of the PPO algorithm that make it a promising method also for controlling more complex, 3D cases. Firstly, the PPO algorithm is able to sample data

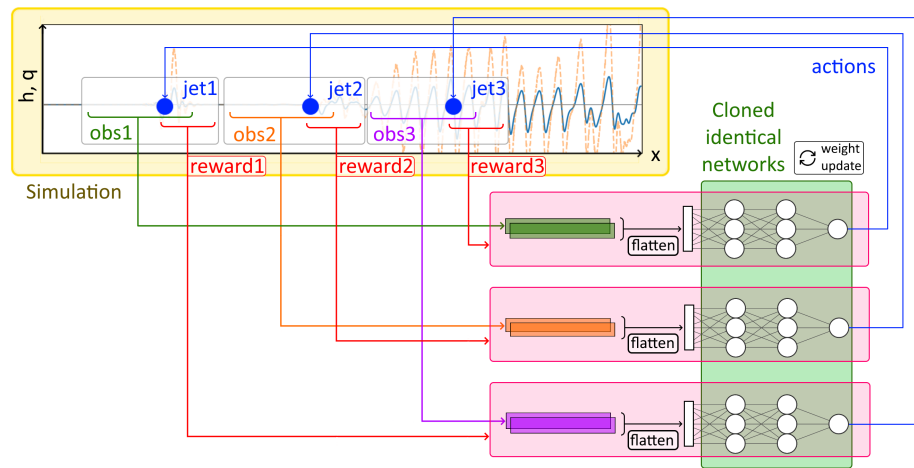


Figure 5. Illustration of the self-collaborative approach demonstrated by Belus *et al.* [128] for controlling systems with several similar control outputs. In this method, clones of the DRL agent are used on each of the output and share their experience samples and training. This allows to learn control laws in constant time, independently of the number of outputs used. This figure is reproduced with minor changes from Belus *et al.* [128].

from several independent environments when performing learning. Therefore, one can effectively parallelize the DRL training by using many CFD simulations running in parallel, as was presented by Rabault and Kuhnle [127]. This allows to drastically accelerate the training ([127] reports speed ups of up to 20, though the more complex the system to control, the higher the speed up attainable using this technique). This allows to perform reasonably fast PPO training, even in cases when the underlying environment is difficult to speed up. For example, in the case of CFD-based environment, this allows to scale the training to a number of compute cores $N \times M$, where N is the number of simulations run in parallel, and M is the maximum optimal parallelization of the CFD simulation itself. Secondly, it is possible to formulate the learning problem in such a way as to take advantage of invariants in the physical system that is undergoing control, as was illustrated by Belus *et al.* [128]. In their work, Belus *et al.* [128] formulated the learning problem as a self-collaborative interaction between several clones of the same environment, as is visible in Figure 5. This, in turn, allows the DRL agent to combine the information obtained at many locations that follow the same physical rules into a single policy. Belus *et al.* [128] argue based on theoretical considerations and prove empirically that, without such a technique, performing learning on a system with N_o outputs has a cost that scales as C^{N_o} , when C is the cost of training for a single output. This is prohibitively expensive as the number of outputs is increased. Belus *et al.* [128] then demonstrated empirically that, by contrast, using the approach presented in Figure 5 allows to perform training at a constant cost, independently of the number of outputs, as long as the control law to learn is similar at all outputs. This is a critical enabling factor for the application of DRL to realistic configurations, where many similar control outputs will be distributed across the physical system to control.

4. Conclusions and outlook

Machine Learning-based control methods are an exciting set of techniques that are receiving considerable attention recently for performing active flow control. This spike in interest follows both increases in computational power and the development of effective algorithms that can learn effectively through direct interaction with black-box, complex systems. These ML methods follow a completely different approach compared with how flow control strategies have usually been designed. Instead of performing a local analysis of the flow properties by considering the flow equations and using advanced mathematical and analytical tools to find optimal perturbations, ML techniques discover control strategy through a trial-and-error approach. There are a number of promising methods that

belong to the ML family of control algorithms, including for example Genetic Programming (GP) and Deep Reinforcement Learning (DRL). In this review, we focused on DRL methods in particular, and we discussed how recent works indicate that it can be efficiently used for controlling large, complex, non-linear systems arising from control tasks in fluid mechanics.

The efficiency of DRL has been demonstrated in a number of active flow control situations so far, and the fluid mechanics community is progressively tackling more and more complex flow configurations. In particular, recent works are pushing the use of DRL for flow control into intermediate Re values, leading to more non-linear and more complex flows, so far successfully. The next steps will be to demonstrate DRL control of complex 3D CFD simulations, and to further increase Re to reach fully turbulent conditions. While this will pose new challenges to the DRL method due to the inherent increase in complexity compared with the configurations studied so far, a number of preliminary works indicate that DRL is well adapted to controlling complex systems with a large number of control locations, and that the inherent parallelism present in the DRL experience sampling process can offer large speed ups on complex dynamical systems.

This push to more complex systems represents not only a scientific, but also a technical endeavour. Indeed, applying DRL to 3D flow control at moderate to high Re will pose a number of technical challenges regarding the amount of CFD computational power required, the ability to handle large amounts of data, and the coupling of CFD and DRL codes that were designed independently of each other at a time when the ability to couple them was not yet foreseen. All these aspects put tough requirements on the level of both expert knowledge (few people are expert in both DRL and large scale CFD) and general technical expertise (combining several different complex software stacks into a single system, and deploying this in HPC environments). In our opinion, these technical challenges, rather than fundamental issues, are presently the main limiting factor for applying ML control to active flow control. A possible way out of this challenge is to follow the example set by the ML community and adopt a resolute open source release policy of codes, scripts, tutorials, and trained networks, to reduce the barrier to entry for new groups joining in this research direction.

Author Contributions: All authors contributed towards the ideation, writing of the original draft and editing of the manuscript. All authors have read and agreed to the published version of the manuscript.

Acknowledgments: RV acknowledges the financial support from the Swedish Research Council (VR). OL is partially financed by a Ramón y Cajal postdoctoral contract by the Ministerio de Economía y Competitividad, Secretaría de Estado de Investigación, Desarrollo e Innovación, Spain (RYC2018-025949-I).

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Bouwer, J. Will airline hubs recover from COVID-19? Available at: <https://www.mckinsey.com/industries/travel-logistics-and-transport-infrastructure/our-insights/will-airline-hubs-recover-from-covid-19#> **2020**.
2. Liu, J. et al.. Will the aviation industry have a bright future after the COVID-19 outbreak? Evidence from Chinese airport shipping sector. *Journal of Risk and Financial Management* **2020**, *13*, 276.
3. Owen, B.; Lee, D.S.; Lim, L. Flying into the future: Aviation emissions scenarios to 2050. *Environ. Sci. Technol.* **2010**, *44*, 2255–2260.
4. Kroo, I. Drag due to lift: Concepts for prediction and reduction. *Annu. Rev. Fluid Mech.* **2001**, *33*, 587–617.
5. UN General Assembly. Transforming our world: the 2030 Agenda for Sustainable Development. *Resolut. A RES 70 1* **2015**, *25*, 1–35.
6. Vinuesa, R.; Azizpour, H.; Leite, I.; Balaam, M.; Dignum, V.; Domisch, S.; Felländer, A.; Langhans, S.D.; Tegmark, M.; Fuso Nerini, F. The role of artificial intelligence in achieving the Sustainable Development Goals. *Nat. Commun.* **2020**, *11*, 233.
7. Gupta, S.; Langhans, S.D.; Domisch, S.; Fuso-Nerini, F.; Felländer, A.; Battaglini, M.; Tegmark, M.; Vinuesa, R. Assessing whether artificial intelligence is an enabler or an inhibitor of sustainability at indicator level. *Transp. Eng.* **2021**, *4*, 100064.

8. Fahland, G.; Stroh, A.; Frohnafel, B.; Atzori, M.; Vinuesa, R.; Schlatter, P.; Gatti, D. Investigation of blowing and suction for turbulent flow control on airfoils. *AIAA J.*, <https://doi.org/10.2514/1.J060211> **2021**.
9. Bechert, D.W.; Bartenwerfer, M. The viscous flow on surfaces with longitudinal ribs. *J. Fluid Mech.* **1989**, *206*, 105–112.
10. Viswanath, P.R. Aircraft viscous drag reduction using riblets. *Prog. Aerosp. Sci.* **2002**, *38*, 571–600.
11. Gad-el Hak, M. Flow control – Passive, active and reactive flow management. *Cambridge University Press, Cambridge*. **2000**.
12. Choi, H.; Jeon, W.P.; Kim, J. Control of flow over a bluff body. *Annu. Rev. Fluid Mech.* **2008**, *40*, 113–139.
13. Quadrio, M. Drag reduction in turbulent boundary layers by in-plane wall motion. *Philos. Trans. Royal Soc. A* **2011**, *369*, 1428–1442.
14. Choi, H.; P., M.; Kim, J. Active turbulence control for drag reduction in wall-bounded flows. *J. Fluid Mech.* **1994**, *262*, 75–110.
15. Stroh, A.; Frohnafel, B.; Schlatter, P.; Hasegawa, Y. A comparison of opposition control in turbulent boundary layer and turbulent channel flow. *Phys. Fluids* **2015**, *27*, 075101.
16. Stroh, A.; Frohnafel, B.; Schlatter, P.; Hasegawa, Y. Status and perspectives of laminar flow. *Aeronaut. J.* **2005**, *109*, 639–644.
17. Wood, R.M. Impact of advanced aerodynamic technology on transportation energy consumption. *SAE Technical Papers* **2004**, *113*, 854–874.
18. Abbas, A.; De Vicente, J.; Valero, E. Aerodynamic technologies to improve aircraft performance. *Aerosp. Sci. Technol.* **2013**, *28*, 100–132.
19. Hwang, D. A proof of concept experiment for reducing skin friction by using a micro-blowing technique. *35th Aerospace Sciences Meeting and Exhibit* **1997**.
20. Hwang, D. Review of research into the concept of the microblowing technique for turbulent skin friction reduction. *Prog. Aerosp. Sci.* **2004**, *40*, 559–575.
21. Kornilov, V.I. Current state and prospects of researches on the control of turbulent boundary layer by air blowing. *Prog. Aerosp. Sci.* **2015**, *76*, 1–23.
22. Park, J.; Choi, H. Effects of uniform blowing or suction from a spanwise slot on a turbulent boundary layer flow. *Phys. Fluids* **1999**, *11*, 3095–3105.
23. Kametani, Y.; Fukagata, K. Direct numerical simulation of spatially developing turbulent boundary layers with uniform blowing or suction. *J. Fluid Mech.* **2011**, *681*, 154–172.
24. Kametani, Y.; Fukagata, K.; Örlü, R.; Schlatter, P. Effect of uniform blowing/suction in a turbulent boundary layer at moderate Reynolds number. *Int. J. Heat Fluid Flow* **2015**, *55*, 132–142.
25. Noorani, A.; Vinuesa, R.; Brandt, L.; Schlatter, P. Aspect ratio effect on particle transport in turbulent duct flows. *Phys. Fluids* **2016**, *28*, 115103.
26. Bobke, A.; Vinuesa, R.; Örlü, R.; Schlatter, P. History effects and near equilibrium in adverse-pressure-gradient turbulent boundary layers. *J. Fluid Mech.* **2017**, *820*, 667–692.
27. Vinuesa, R.; Negi, P.S.; Atzori, M.; Hanifi, A.; Henningson, D.S.; Schlatter, P. Turbulent boundary layers around wing sections up to $Re_c = 1,000,000$. *Int. J. Heat Fluid Flow* **2018**, *72*, 86–99.
28. Tanarro, A.; Vinuesa, R.; Schlatter, P. Effect of adverse pressure gradients on turbulent wing boundary layers. *J. Fluid Mech.* **2020**, *883*, A8.
29. Eto, K.; Kondo, Y.; Fukagata, K.; Tokugawa, N. Assessment of friction drag reduction on a clark-Y airfoil by uniform blowing. *AIAA J.* **2019**, *57*, 2774–2782.
30. Kornilov, V.I.; Kavun, I.N.; Popkov, A.N. Modification of turbulent airfoil section flow using a combined control action. *Thermophys. Aeromechanics* **2019**, *26*, 165–178.
31. Vinuesa, R.; Schlatter, P. Skin-friction control of the flow around a wing section through uniform blowing. *Proc. European Drag Reduction and Flow Control Meeting (EDRFCM)* **2017**.
32. Atzori, M.; Vinuesa, R.; Fahland, G.; Stroh, A.; Gatti, D.; Frohnafel, B.; Schlatter, P. Aerodynamic effects of uniform blowing and suction on a NACA4412 airfoil. *Flow Turbul. Combust.* **2020**, *105*, 735–759.
33. Albers, M.; Meysonnat, P.S.; Schröder, W. Actively reduced airfoil drag by transversal surface waves. *Flow Turbul. Combust.* **2019**, *102*, 865–886.
34. Jeong, J.; Hussain, F. On the identification of a vortex. *J. Fluid Mech.* **1995**, *285*, 69–94.
35. Hosseini, S.M.; Vinuesa, R.; Schlatter, P.; Hanifi, A.; Henningson, D.S. Direct numerical simulation of the flow around a wing section at moderate Reynolds number. *Int. J. Heat Fluid Flow* **2016**, *61*, 117–128.

36. Vinuesa, R.; Hosseini, S.M.; Hanifi, A.; Henningson, D.S.; Schlatter, P. Pressure-gradient turbulent boundary layers developing around a wing section. *Flow Turbul. Combust.* **2017**, *99*, 613–641.
37. Batikh, A.; Baldas, L.; Colin, S. Application of Active Flow Control on Aircrafts - State of the Art. *International Workshop on Aircraft System Technologies* **2017**, pp. 1–10.
38. Khodadoust, A.; Washburn, A. Active control of flow separation on a high-lift system with slotted flap at high Reynolds number. Collection of Technical Papers - AIAA Applied Aerodynamics Conference, 2007, Vol. 3, pp. 1645–1669. doi:10.2514/6.2007-4424.
39. Khün, T.; Ciobaca, V.; Rudnik, R.; Gölling, B.; Breitenstein, W. Active flow separation control on a high-lift wing-body configuration part 1: Baseline flow and constant blowing. 29th AIAA Applied Aerodynamics Conference 2011, 2011, pp. 1–11. doi:10.2514/6.2011-3168.
40. Radespiel, R.; Burnazzi, M.; Casper, M.; Scholz, P. Active flow control for high lift with steady blowing. *Aeronautical Journal* **2016**, *120*, 171–200. doi:10.1017/aer.2015.7.
41. Fricke, S.; Ciobaca, V.; Kröhnert, A.; Wild, J.; Blesbois, O. Active Flow Control Applied at the Engine-Wing Junction. 5th CEAS Air & Aerospace conference, 2015, number 249, pp. 1–12.
42. Schloesser, P.; Soudakov, V.; Bauer, M.; Wild, J. Active separation control at the pylon-wing junction of a real-scale model. *AIAA Journal* **2019**, *57*, 132–141. doi:10.2514/1.J057345.
43. Hue, D.; François, C.; Dandois, J.; Gebhardt, A. Simulations of an aircraft with constant and pulsed blowing flow control at the engine/wing junction. *Aerospace Science and Technology* **2017**, *69*, 659–673. doi:10.1016/j.ast.2017.07.031.
44. Shmilovich, A.; Yadlin, Y.; Whalen, E.A. Active flow control computations: From a single actuator to a complete airplane. *AIAA Journal* **2018**, *56*, 4730–4740. doi:10.2514/1.J056307.
45. Andino, M.Y.; Lin, J.C.; Roman, S.; Graff, E.C.; Gharib, M.; Whalen, E.A.; Wygnanski, I.J. Active Flow Control on Vertical Tail Models. *AIAA Journal* **2019**, *57*, 3322–3338. doi:10.2514/1.J057876.
46. Whalen, E.A.; Lacy, D.; Lin, J.C.; Andino, M.Y.; Washburn, A.E.; Graff, E.; Wygnanski, I.J. Performance enhancement of a full-scale vertical tail model equipped with active flow control. *53rd AIAA Aerospace Sciences Meeting* **2015**. doi:10.2514/6.2015-0784.
47. Kara, K.; Kim, D.; Morris, P.J. Flow-separation control using sweeping jet actuator. *AIAA Journal* **2018**, *56*, 4604–4613. doi:10.2514/1.J056715.
48. Chen, S.; Li, W.; Meng, Q.; Zhou, Z.; Wang, S. Effects of a sweeping jet actuator on aerodynamic performance in a linear turbine cascade with tip clearance. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering* **2019**, *233*, 4468–4481. doi:10.1177/0954410018823606.
49. Aley, K.S.; Guha, T.K.; Kumar, R. Active Flow Control of a High-Lift Supercritical Airfoil with Microjet Actuators. *AIAA Journal* **2020**, pp. 1–17. doi:10.2514/1.j058939.
50. Glezer, A. Some aspects of aerodynamic flow control using synthetic-jet actuation. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **2011**, *369*, 1476–1494. doi:10.1098/rsta.2010.0374.
51. Hartwich, P.M.; Shmilovich, A.; Lacy, D.S.; Dickey, E.D.; Scalafani, A.J.; Sundaram, P.; Yadlin, Y. Refined AFC-Enabled High-Lift System Integration Study. Technical Report CR-2016-219170, 2016.
52. McCormick, D.C. Boundary Layer Separation control with Directed Synthetic Jets. 38th Aerospace Sciences Meeting & Exhibit, 2000, number AIAA 2000-0519. doi:10.2514/6.2000-519.
53. Amitay, M.; Glezer, a. Role of actuation frequency in controlled flow reattachment over a stalled airfoil. *AIAA Journal* **2002**, *40*, 209–216. doi:10.2514/3.15052.
54. Gilarranz, J.L.; Traub, L.W.; Rediniotis, O.K. A New Class of Synthetic Jet Actuators—Part I: Design, Fabrication and Bench Top Characterization. *Journal of Fluids Engineering* **2005**, *127*, 367. doi:10.1115/1.1839931.
55. You, D. and Moin, P. Active control of flow separation over an airfoil using synthetic jets. *Journal of Fluids and Structures* **2008**, *24*, 1349–1357. doi:10.1016/j.jfluidstructs.2008.06.017.
56. Jabbal, M.; Liddle, S.C.; Crowther, W.J. Active flow control systems architectures for civil transport aircraft. *Journal of Aircraft* **2010**, *47*, 1966–1981. doi:10.2514/1.C000237.
57. Shmilovich, A.; Yadlin, Y. Active Flow Control for Practical High-Lift Systems. *Journal of Aircraft* **2009**, pp. 1354–1364.
58. Bauer, M.; Lohse, J.; Haucke, F.; Nitsche, W. High-Lift Performance Investigation of a Two-Element Configuration with a Two-Stage Actuator System. *AIAA Journal* **2014**, *52*, 1307–1313. doi:10.2514/1.j052639.

59. Lin, J.C.; Pack Melton, L.G.; Viken, S.; Andino, M.Y.; Koklu, M.; Hannon, J.; Vatsa, V.N. High Lift Common Research Model for Wind Tunnel Testing: An Active Flow Control Perspective. 55th AIAA Aerospace Sciences Meeting, 2017. doi:10.2514/6.2017-0319.
60. Lin, J.C.; Melton, L.G.P.; Hannon, J.; Andino, M.; Koklu, M.; Paschal, K.; Vatsa, V.N. Wind Tunnel Testing of Active Flow Control on the High Lift Common Research Model. AIAA Aviation 2019 Forum. doi:10.2514/6.2019-3723.
61. Jansen, K.E.; Rasquin, M.; Farnsworth, J.A.; Rathay, N.; Monastero, M.C.; Amitay, M. Interaction of a Synthetic Jet with Separated Flow over a Vertical Tail. *AIAA Journal* **2018**, *56*, 2653–2668. doi:10.2514/1.j056751.
62. Lehmkuhl, O.; Lozano-Durán, A.; Rodriguez, I. Active flow control for external aerodynamics: from micro air vehicles to a full aircraft in stall. *Journal of Physics: Conference Series* **2020**, *1522*, 012017. doi:10.1088/1742-6596/1522/1/012017.
63. Johnson, F.T.; Tinoco, E.N.; Yu, N.J. Thirty years of development and application of CFD at Boeing Commercial Airplanes, Seattle. *Computers & Fluids* **2005**, *34*, 1115–1151.
64. Abbas-Bayoumi, A.; Becker, K. An industrial view on numerical simulation for aircraft aerodynamic design. *Journal of Mathematics in Industry* **2011**, *1*, 1–14.
65. Spalart, P.R.; Venkatakrishnan, V. On the role and challenges of CFD in the aerospace industry. *The Aeronautical Journal* **2016**, *120*, 209–232. doi:10.1017/aer.2015.10.
66. Clark, A.M.; Slotnick, J.P.; Taylor, N.J.; Rumsey, C.L. Requirements and challenges for CFD validation within the High-Lift Common Research Model ecosystem. AIAA Aviation Forum, 2020, [<https://arc.aiaa.org/doi/pdf/10.2514/6.2020-2772>]. doi:10.2514/6.2020-2772.
67. Recommended Practice: When Flight Modelling Is Used to Reduce Flight Testing Supporting Aircraft Certification (AIAA R-154-2021). In *Recommended Practice: When Flight Modelling Is Used to Reduce Flight Testing Supporting Aircraft Certification (AIAA R-154-2021)*; 2021; [<https://arc.aiaa.org/doi/pdf/10.2514/4.106231.001>]. doi:10.2514/4.106231.001.
68. Slotnick, J.; Khodadoust, A.; Alonso, J.; Darmofal, D.; Gropp, W.; Lurie, E.; Mavriplis, D. *CFD Vision 2030 Study: A Path to Revolutionary Computational Aerosciences*; Tech. Rep. CR-2014-218178, NASA, 2014.
69. Lozano-Durán, A.; Bose, S.T.; Moin, P. Performance of Wall-Modeled LES with Boundary-Layer-Conforming Grids for External Aerodynamics. *AIAA J.* **2021**, p. J061041.
70. Spalart, P.R.; Jou, W.H.; Strelets, M.; Allmaras, S.R. Comments on the feasibility of LES for wings, and on a hybrid RANS/LES approach. *Advances in DNS/LES* **1997**, *1*, 4–8.
71. Spalart, P.R. Detached-eddy simulation. *Annu. Rev. Fluid Mech.* **2009**, *41*, 181–202.
72. Casey, M.; Wintergerste, T.; European Research Community on Flow, T.; Combustion.. *ERCOFTAC best practice guidelines: ERCOFTAC special interest group on “quality and trust in industrial CFD”*; ERCOFTAC, 2000.
73. Rumsey, C.L.; Slotnick, J.P.; Sclafani, A.J. Overview and Summary of the Third AIAA High Lift Prediction Workshop. *J. Aircr.* **2019**, *56*, 621–644. doi:10.2514/1.C034940.
74. Rumsey, C.L.; Carlson, J.; Ahmad, N. FUN3D Junction Flow Computations Compared with Experimental Data. AIAA Scitech 2019 Forum, 2019. doi:10.2514/6.2019-0079.
75. Choi, H.; Moin, P. Grid-point requirements for large eddy simulation: Chapman’s estimates revisited. *Phys. Fluids* **2012**, *24*, 011702.
76. Piomelli, U.; Balaras, E. Wall-layer models for large-eddy simulations. *Annu. Rev. Fluid Mech.* **2002**, *34*, 349–374. doi:10.1146/annurev.fluid.34.082901.144919.
77. Cabot, W.H.; Moin, P. Approximate wall boundary conditions in the large-eddy simulation of high Reynolds number flow. *Flow Turbul. Combust.* **2000**, *63*, 269–291. doi:10.1023/A:1009958917113.
78. Larsson, J.; Kawai, S.; Bodart, J.; Bermejo-Moreno, I. Large eddy simulation with modeled wall-stress: recent progress and future directions. *Mech. Eng. Rev.* **2016**, *3*, 1–23. doi:10.1299/mer.15-00418.
79. Bose, S.T.; Park, G.I. Wall-Modeled LES for Complex Turbulent Flows. *Ann. Rev. Fluid Mech.* **2018**, *50*, 535–561. doi:10.1146/annurev-fluid-122316-045241.
80. Deardorff, J. A numerical study of three-dimensional turbulent channel flow at large Reynolds numbers. *J. Fluid Mech.* **1970**, *41*, 453–480.
81. Schumann, U. Subgrid scale model for finite difference simulations of turbulent flows in plane channels and annuli. *J. Comp. Phys.* **1975**, *18*, 376–404. doi:10.1016/0021-9991(75)90093-5.

82. Piomelli, U.; Ferziger, J.; Moin, P.; Kim, J. New approximate boundary conditions for large eddy simulations of wall-bounded flows. *Phys. Fluids* **1989**, *1*, 1061–1068. doi:10.1063/1.857397.
83. Balaras, E.; Benocci, C.; Piomelli, U. Two-layer approximate boundary conditions for large-eddy simulations. *AIAA J.* **1996**, *34*, 1111–1119. doi:10.2514/3.13200.
84. Wang, M.; Moin, P. Dynamic wall modeling for large-eddy simulation of complex turbulent flows. *Phys. Fluids* **2002**, *14*, 2043–2051.
85. Chung, D.; Pullin, D.I. Large-eddy simulation and wall modelling of turbulent channel flow. *J. Fluid Mech.* **2009**, *631*, 281–309. doi:10.1017/S0022112009006867.
86. Bodart, J.; Larsson, J. Wall-modeled large eddy simulation in complex geometries with application to high-lift devices. *Center for Turbulence Research - Annual Research Briefs* **2011**, pp. 37–48.
87. Kawai, S.; Larsson, J. Dynamic non-equilibrium wall-modeling for large eddy simulation at high Reynolds numbers. *Phys. Fluids* **2013**, *25*, 015105.
88. Bermejo-Moreno, I.; Campo, L.; Larsson, J.; Bodart, J.; Helmer, D.; Eaton, J.K. Confinement effects in shock wave/turbulent boundary layer interactions through wall-modelled large-eddy simulations. *J. Fluid Mech.* **2014**, *758*, 5–62. doi:10.1017/jfm.2014.505.
89. Park, G.I.; Moin, P. An improved dynamic non-equilibrium wall-model for large eddy simulation. *Phys. Fluids* **2014**, *26*, 015108. doi:10.1063/1.4861069.
90. Yang, X.I.A.; Sadique, J.; Mittal, R.; Meneveau, C. Integral wall model for large eddy simulations of wall-bounded turbulent flows. *Phys. Fluids* **2015**, *27*, 025112. doi:10.1063/1.4908072.
91. Bose, S.T.; Moin, P. A dynamic slip boundary condition for wall-modeled large-eddy simulation. *Phys. Fluids* **2014**, *26*, 015104. doi:10.1063/1.4849535.
92. Lozano-Durán, A.; Bae, H.; Bose, S.; Moin, P. Dynamic wall models for the slip boundary condition. *Center for Turbulence Research - Annual Research Briefs* **2017**, pp. 229–242.
93. Bae, H.J.; Lozano-Durán, A.; Bose, S.T.; Moin, P. Dynamic slip wall model for large-eddy simulation. *J. Fluid Mech.* **2019**, *859*, 400–432. doi:10.1017/jfm.2018.838.
94. Yang, X.I.A.; Zafar, S.; Wang, J.X.; Xiao, H. Predictive large-eddy-simulation wall modeling via physics-informed neural networks. *Phys. Rev. Fluids* **2019**, *4*, 034602. doi:10.1103/PhysRevFluids.4.034602.
95. Lozano-Durán, A.; Bae, H.J. Self-critical machine-learning wall-modeled LES for external aerodynamics. *Center for Turbulence Research - Annual Research Briefs* **2020**, pp. 197–210.
96. Radhakrishnan, S.; Gyamfi, L.A.; Miro, A.; Font, B.; J., C.; Lehmkühl, O. A data-driven wall-shear stress model for LES using gradient boosted decision trees. *International Conference on High Performance Computing* **2021**, pp. 105–121.
97. Eivazi, H.; Tahani, M.; Schlatter, P.; Vinuesa, R. Physics-informed neural networks for solving Reynolds-averaged Navier–Stokes equations. *Preprint arXiv:2107.10711* **2021**.
98. Koza, J.R. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*; MIT Press: Cambridge, MA, USA, 1992.
99. Rudin, C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat. Mach. Intell.* **2019**, *1*, 206–215.
100. Vinuesa, R.; Sirmacek, B. Interpretable deep-learning models to help achieve the Sustainable Development Goals. *Nat. Mach. Intell.* **2021**, *3*, 926.
101. Li, R.; Noack, B.R.; Cordier, L.; Borée, J.; Harambat, F. Drag reduction of a car model by linear genetic programming control. *Exp. Fluids* **2017**, *58*, 103.
102. Minelli, G.; Dong, T.; Noack, B.R.; Krajnović, S. Upstream actuation for bluff-body wake control driven by a genetically inspired optimization. *J. Fluid Mech.* **2020**, *893*, A1.
103. Rasmussen, C.E. Gaussian processes for machine learning. *Journal für Urologie und Urogynäkologie. Berlin: Springer International Publishing* **2003**, *7*.
104. Morita, Y.; Rezaeiravesh, S.; Tabatabaei, N.; Vinuesa, R.; Fukagata, K.; Schlatter, P. Applying Bayesian optimization with Gaussian-process regression to Computational Fluid Dynamics problems. *Preprint arXiv:2101.09985* **2021**.
105. Mahfoze, O.A.; Moody, A.; Wynn, A.; Whalley, R.D.; Laizet, S. Reducing the skin-friction drag of a turbulent boundary-layer flow with low-amplitude wall-normal blowing within a Bayesian optimization framework. *Phys. Rev. Fluids* **2019**, *4*, 094601.

106. Kornilov, V.I.; Boiko, A.V. Efficiency of air microblowing through microperforated wall for flat plate drag reduction. *AIAA J.* **2012**, *50*, 724–732.
107. 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. *J. Fluid Mech.* **2021**, *928*, A27.
108. Güemes, A.; Discetti, S.; Ianiro, A.; Sirmacek, B.; Azizpour, H.; Vinuesa, R. From coarse wall measurements to turbulent velocity fields through deep learning. *Phys. Fluids* **2021**, *3*, 075121.
109. Abreu, L.I.; Cavalieri, A.V.G.; Schlatter, P.; Vinuesa, R.; Henningson, D.S. Spectral proper orthogonal decomposition and resolvent analysis of near-wall coherent structures in turbulent pipe flows. *J. Fluid Mech.* **2020**, *900*, A11.
110. Abreu, L.I.; Tanarro, A.; Cavalieri, A.V.; Schlatter, P.; Vinuesa, R.; Hanifi, A.; Henningson, D.S. Spanwise-coherent hydrodynamic waves around flat plates and airfoils. *J. Fluid Mech.* **2021**, *927*, A1.
111. Eivazi, H.; Le Clainche, S.; Hoyas, S.; Vinuesa, R. Towards extraction of orthogonal and parsimonious non-linear modes from turbulent flows. *Preprint arXiv:2109.01514* **2021**.
112. François-Lavet, V.; Henderson, P.; Islam, R.; Bellemare, M.G. An Introduction to Deep Reinforcement Learning. *Foundations and Trends in Machine Learning, II* **2018**, *3*, 1–140.
113. 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. *J. Fluid Mech.* **2019**, *865*, 281–302.
114. Hornik, K.; Stinchcombe, M.; White, H. Multilayer feedforward networks are universal approximators. *Neural networks* **1989**, *2*, 359–366.
115. Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; Klimov, O. Proximal policy optimization algorithms. *Preprint ArXiv:1707.06347* **2017**.
116. Mnih, V. et al.. Human-level control through deep reinforcement learning. *Nature* **2015**, *518*, 529–533.
117. Gu, S.; Lillicrap, T.; Sutskever, U.; Levine, S. Continuous deep Q-learning with model-based acceleration. *33rd International Conference on Machine Learning, ICML* **2016**, *6*, 4135–4148.
118. Rabault, J.; Ren, F.; Zhang, W.; Tang, H.; Xu, H. Deep reinforcement learning in fluid mechanics: A promising method for both active flow control and shape optimization. *Journal of Hydrodynamics* **2020**, *32*, 234–246.
119. Garnier, P.; Viquerat, J.; Rabault, J.; Larcher, A.; Kuhnle, A.; Hachem, E. A review on deep reinforcement learning for fluid mechanics. *Computers & Fluids* **2021**, *225*, 104973.
120. Bucci, M.A.; Semeraro, O.; Allauzen, A.; Wisniewski, G.; Cordier, L.; Mathelin, L. Control of chaotic systems by deep reinforcement learning. *Proceedings of the Royal Society A* **2019**, *475*, 20190351.
121. Paris, R.; Beneddine, S.; Dandois, J. Robust flow control and optimal sensor placement using deep reinforcement learning. *Journal of Fluid Mechanics* **2021**, *913*.
122. Beintema, G.; Corbetta, A.; Biferale, L.; Toschi, F. Controlling Rayleigh–Bénard convection via reinforcement learning. *Journal of Turbulence* **2020**, *21*, 585–605.
123. 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* **2020**, *32*, 053605, [<https://doi.org/10.1063/5.0006492>]. doi:10.1063/5.0006492.
124. Xu, H.; Zhang, W.; Deng, J.; Rabault, J. Active flow control with rotating cylinders by an artificial neural network trained by deep reinforcement learning. *Journal of Hydrodynamics* **2020**, *32*, 254–258.
125. 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* **2020**, *117*, 26091–26098, [<https://www.pnas.org/content/117/42/26091.full.pdf>]. doi:10.1073/pnas.2004939117.
126. Ren, F.; Rabault, J.; Tang, H. Applying deep reinforcement learning to active flow control in weakly turbulent conditions. *Physics of Fluids* **2021**, *33*, 037121.
127. Rabault, J.; Kuhnle, A. Accelerating deep reinforcement learning strategies of flow control through a multi-environment approach. *Phys. Fluids* **2019**, *31*, 094105.
128. Belus, V.; Rabault, J.; Viquerat, J.; Che, Z.; Hachem, E.; Reglade, U. Exploiting locality and translational invariance to design effective deep reinforcement learning control of the 1-dimensional unstable falling liquid film. *AIP Advances* **2019**, *9*, 125014, [<https://doi.org/10.1063/1.5132378>]. doi:10.1063/1.5132378.