



# Optimisation and Analysis of Streamwise-Varying Wall-Normal Blowing in a Turbulent Boundary Layer

Joseph O'Connor<sup>1</sup> · Mike Diessner<sup>2</sup> · Kevin Wilson<sup>3</sup> · Richard D. Whalley<sup>4</sup> · Andrew Wynn<sup>1</sup> · Sylvain Laizet<sup>1</sup>

Received: 14 October 2022 / Accepted: 9 March 2023  
© The Author(s) 2023

## Abstract

Skin-friction drag is a major engineering concern, with wide-ranging consequences across many industries. Active flow-control techniques targeted at minimising skin friction have the potential to significantly enhance aerodynamic efficiency, reduce operating costs, and assist in meeting emission targets. However, they are difficult to design and optimise. Furthermore, any performance benefits must be balanced against the input power required to drive the control. Bayesian optimisation is a technique that is ideally suited to problems with a moderate number of input dimensions and where the objective function is expensive to evaluate, such as with high-fidelity computational fluid dynamics simulations. In light of this, this work investigates the potential of low-intensity wall-normal blowing as a skin-friction drag reduction strategy for turbulent boundary layers by combining a high-order flow solver (*Incompact3d*) with a Bayesian optimisation framework. The optimisation campaign focuses on streamwise-varying wall-normal blowing, parameterised by a cubic spline. The inputs to be optimised are the amplitudes of the spline control points, whereas the objective function is the net-energy saving (NES), which accounts for both the skin-friction drag reduction and the input power required to drive the control (with the input power estimated from real-world data). The results of the optimisation campaign are mixed, with significant drag reduction reported but no improvement over the canonical case in terms of NES. Selected cases are chosen for further analysis and the drag reduction mechanisms and flow physics are highlighted. The results demonstrate that low-intensity wall-normal blowing is an effective strategy for skin-friction drag reduction and that Bayesian optimisation is an effective tool for optimising such strategies. Furthermore, the results show that even a minor improvement in the blowing efficiency of the device used in the present work will lead to meaningful NES.

**Keywords** Bayesian optimisation · Skin-friction drag reduction · Large-eddy simulation · Net-energy saving · Turbulent boundary layer · Wall-normal blowing

---

✉ Joseph O'Connor  
joseph.oconnor@imperial.ac.uk

Extended author information available on the last page of the article

## 1 Introduction

Skin-friction drag is a major engineering concern, with direct consequences for several different industries and applications. In the aviation industry, for example, skin-friction drag is estimated to account for approximately half of the overall drag (Abbas et al. 2017). Comparable figures are also observed across the transport sector, and can be even higher in maritime applications (Fu et al. 2017), where the proportion of skin-friction drag can be up to 90% for submarines (Gad-El-Hak 1994). It is clear that flow control strategies that specifically target skin-friction drag reduction have the potential to offer substantial improvements in aerodynamic/hydrodynamic efficiency, reduce operating costs, and help meet emission targets. However, despite many decades of extensive research, a performant, practical and affordable method for skin-friction drag reduction is yet to find widespread adoption in real-world applications.

Over the years there have been a variety of flow-control strategies put forward, with the aim of reducing skin-friction drag. This includes passive techniques, such as polymer additives (White and Mungal 2008), riblets (Choi et al. 1993; García-Mayoral and Jiménez 2011), and large-eddy breakup devices (Kim et al. 2017), to name a few. The main advantage of passive strategies is that they do not require a power source, and so any drag reduction obtained is translated directly into energy savings. However, passive approaches often suffer from parasitic drag effects and are typically designed for a specific flow condition. This inflexibility means that passive devices are prone to losses in efficiency over the full range of operating conditions when deployed in real-world applications. On the other hand, active flow-control strategies—such as surface jets (Kametani and Fukagata 2011; Kametani et al. 2015; Kornilov and Boiko 2012; Stroh et al. 2016), wall motion (Quadrio et al. 2009; Marusic et al. 2021), and plasma actuators (Wang et al. 2013; Mahfoze and Laizet 2017)—are among the most promising approaches to tackle skin-friction drag, due to their increased performance, robustness, and flexibility. However, designing optimal active control policies is a challenge, owing to the nonlinear, high-dimensional, and chaotic nature of fluid dynamics. Furthermore, active control strategies require a power supply. Therefore, any performance benefits must be weighed against the power required to drive the control.

Of the available active flow-control strategies, mass-flow injection (e.g. surface jets) is one of the most extensively studied, owing to its practicality and performance. Several studies, both experimental and numerical, have demonstrated impressive reductions in skin-friction drag for a turbulent boundary layer (TBL) with low intensity wall-normal blowing (up to 1% of the free-stream velocity). For example, Kametani and Fukagata (2011) performed direct numerical simulations (DNS) of a zero-pressure gradient TBL with both uniform blowing and suction. Their findings showed that blowing leads to a reduction in skin-friction drag, whereas the reverse is true for suction. Their results also demonstrated a globally-averaged drag reduction of 75% with a blowing intensity of 1% of the freestream velocity. A later study by Kametani et al. (2015) showed similar findings using large-eddy simulations (LES) and investigated the effect of both the position and length of the control region. Their results showed that longer control regions placed further upstream are more beneficial for skin-friction drag reduction. The experimental study of Kornilov and Boiko (2012) used wall-normal blowing through a micro-perforated plate to demonstrate a local drag reduction of approximately 70% with a blowing velocity of 0.287% of the freestream velocity. Stroh et al. (2016) used DNS to study the effect of two different types of control, numerical body force damping and wall-normal blowing, on a TBL. While the body force damping was shown to provide a greater drag reduction in the vicinity of the control

region, the long-lasting downstream effect of the wall-normal blowing led to improved globally-averaged drag reductions, when integrated over the downstream length of the domain. More recently, Atzori et al. (2020) used LES to investigate the effect of both uniform blowing and suction on the pressure and suction sides of a NACA4412 aerofoil. Their results support the aforementioned studies by showing that blowing leads to a reduction in skin-friction drag, whereas suction leads to an increase. However, their results also demonstrate that skin-friction drag reduction does not automatically lead to improved efficiencies when accounting for other forms of drag and performance metrics.

It is important to highlight that the reported drag reduction values in all of the above-mentioned studies are highly dependent on the flow conditions (e.g. Reynolds number), as well as the approach for measuring the drag reduction (e.g. local or globally-averaged over a defined streamwise length). Therefore, care must be taken when making comparisons between different studies. It is also worth highlighting that while many studies provide estimates for the overall net-energy saving (NES) by estimating the input power required to drive the control (Kametani and Fukagata 2011; Kornilov and Boiko 2012, 2014; Kametani et al. 2015, 2016; Mahfoze et al. 2019), these are typically based on idealised assumptions and neglect losses associated with real-world devices (e.g. mechanical losses). To accelerate the translation of active flow-control strategies to real-world applications, it is noted that more rigorous estimates of the input power are required.

One of the strengths of mass-flow injection is its flexibility, with a vast range of choices in terms of design, from placement and number of jets, temporal actuation and blowing intensity. However, this flexibility has meant that it is still not clear what the optimal configuration for such a control strategy and particular flow condition is. This has led to a trend towards adopting optimisation techniques. Mathematical optimisation is a mature field and there is a vast range of different techniques available, such as gradient descent (Ruder 2016), genetic algorithms (Katoch et al. 2021), and particle swarm optimisation (Freitas et al. 2020), to name a few. As a result, selecting the most suitable optimisation approach for a particular task is a challenge and depends heavily on the characteristics of the underlying problem. For the purpose of designing control strategies for turbulent flows, the optimisation scheme should be suitable for problems with a moderate number of input dimensions (e.g. amplitude, blowing frequency) and where the objective function (e.g. drag reduction, NES) is non-convex and expensive to evaluate (e.g. LES and DNS). These characteristics make the aforementioned approaches (e.g. gradient descent, evolutionary algorithms) unsuitable, due to the large number of function evaluations required to estimate gradients or iterate over enough generations. Instead, a more sample-efficient approach is required.

Bayesian optimisation is a global optimisation approach that has recently gained significant popularity for hyperparameter tuning of machine learning models (Snoek et al. 2012; Shahriari et al. 2016). The basic approach in Bayesian optimisation is to build a surrogate model of the objective function through repeated evaluations over the input space. New sample points are then chosen by optimising an acquisition function, which balances both exploration and exploitation, applied over the surrogate model. The sample efficiency of Bayesian optimisation comes from the fact it considers both exploration and exploitation in the sampling strategy, making it well suited to problems with expensive-to-evaluate objective functions. Thanks to its success in machine learning, Bayesian optimisation has also started to find applications in engineering optimisation, such as turbulent drag reduction. Mahfoze et al. (2019) recently applied Bayesian optimisation to optimise the control parameters of a wall-normal blowing configuration within a TBL. Rather than targeting skin-friction drag reduction, they optimised for NES by taking into account both the power

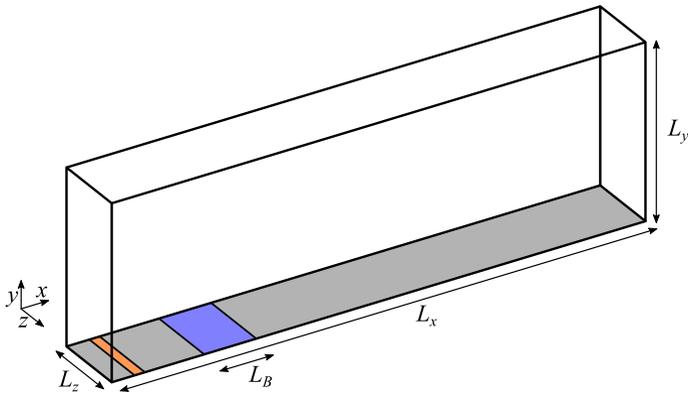
saved due to drag reduction and the input power required to drive the control, which was estimated via both an idealised approximation and real-world experimental data. Their approach was able to find combinations of control parameters leading to between 1.2 and 5% NES, depending on the approach used to estimate the power input. More recently, Morita et al. (2022) applied Bayesian optimisation to a wide variety of fluid dynamics problems, showing that it was capable of finding globally optimal solutions with relatively few evaluations of the objective function (e.g. approximately 90 evaluations required for up to 8 inputs). Finally, Larroque et al. (2022) provided a detailed analysis of the performance characteristics of Bayesian optimisation by using it to optimise an active flow-control strategy for the flow around a cylinder. Their results showed that, in this instance, Bayesian optimisation significantly outperformed a variety of other derivative-free optimisation techniques, such as particle swarm optimisation and the explorative gradient method, in terms of both overall drag reduction and the number of function evaluations required.

Several previous studies have shown that low-intensity wall-normal blowing can have a long-lasting downstream effect on the skin-friction drag (Kornilov and Boiko 2014; Stroh et al. 2016; Kametani et al. 2016; Mahfoze et al. 2019). These works show a subtle but sustained drag reduction that persists well beyond the control region. This ‘inertia’ in the skin friction recovery is a key contributor to the overall NES (Mahfoze et al. 2019) and suggests a potential pathway to improving NES via a streamwise-varying blowing configuration (e.g. via intermittent blowing or a continuous profile with streamwise modulation). In light of this, the purpose of this work is to investigate the potential of a time-independent streamwise-varying blowing strategy, parameterised by a cubic spline, to achieve skin-friction drag reduction and NES in a TBL. In this endeavour, the high-order flow solver `Incompact3d` is combined with a newly designed Bayesian optimisation framework to optimise the three control parameters that define the wall-normal blowing profile. The key performance metric is the NES, which takes into account both the power saving due to drag reduction and the power required to drive the actuation. The input power requirements are estimated via real-world experimental to provide a realistic estimate of the NES. Note that this work follows on from Mahfoze et al. (2019). Therefore, both studies share some similarities, such as the same flow problem, solver and power estimation method. However, the key novelties of this work are: (1) investigation of a novel blowing strategy; (2) the validation and use of implicit LES data, instead of DNS data, to extract the NES values; (3) a detailed analysis of the flow physics for selected cases, including a comparison between effective and ineffective configurations; (4) the use of a state-of-the-art in-house Bayesian optimisation framework (Diessner et al. 2022). In the following sections the case setup is introduced, followed by a description of the flow solver and Bayesian optimisation framework. The results are then presented by first summarising the optimisation campaign. Selected cases are then chosen for further analysis to highlight the drag reduction mechanisms and flow physics. Finally, the main findings are summarised and suggestions for future work are proposed.

## 2 Case Description

### 2.1 Simulation Setup

The case setup is based on the work of Mahfoze et al. (2019), which represents a canonical zero-pressure gradient TBL flow with low-intensity wall-normal blowing. Figure 1 shows

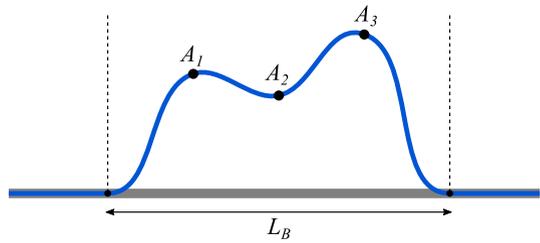


**Fig. 1** Schematic of the simulation setup. The blue and orange shaded regions indicate the blowing and tripping regions, respectively

a schematic of the computational domain, with the control region highlighted by the blue shaded area. The streamwise, wall-normal, and spanwise directions are denoted by  $x$ ,  $y$ , and  $z$ , respectively. Equivalently,  $u$ ,  $v$ , and  $w$  denote the instantaneous velocity components in the three spatial directions, each of which can be decomposed into a time-averaged ( $\bar{u}$ ) and fluctuating component ( $u'$ ), such that  $u = \bar{u} + u'$ . A laminar Blasius solution is prescribed at the inlet, with a boundary layer height of  $\delta_0$  and freestream velocity of  $u_\infty$ . Note that in the following sections all quantities are non-dimensionalised with respect to  $\delta_0$  and  $u_\infty$  (outer-scaling). However, where it is more convenient to use inner-scaled quantities (denoted by the plus superscript notation), these are non-dimensionalised with respect to the canonical friction velocity at the start of the blowing region (unless otherwise stated). The remaining boundary conditions are composed of a convective condition at the outlet, a homogenous Neumann condition in the far-field, and periodic conditions in the spanwise direction. The domain dimensions are  $L_x \times L_y \times L_z = 750 \times 80 \times 30$ . Compared to Mahfoze et al. (2019), the spanwise (periodic) extent of the domain is doubled as it was found to provide better agreement with benchmark data during validation, with only a modest increase in compute time (owing to an acceleration in the convergence of the statistics, which are averaged in the spanwise direction). The Reynolds number at the inlet is  $Re_{\delta_0} = u_\infty \delta_0 / \nu = 1250$ , based on the boundary layer thickness at the inlet, where  $\nu$  is the kinematic viscosity. This is equivalent to a momentum Reynolds number of  $Re_{\theta_0} = u_\infty \theta_0 / \nu \approx 169$ , where  $\theta_0$  is the momentum thickness at the inlet. Under these conditions, the momentum Reynolds number increases to approximately  $Re_\theta \approx 2025$  at the outlet for the canonical case.

The mesh size is chosen to be  $n_x \times n_y \times n_z = 1537 \times 257 \times 128$ , with a uniform spacing in the streamwise and spanwise directions and non-uniform spacing in the wall-normal direction to properly resolve the near-wall effects. This results in a mesh resolution of  $\Delta x^+ = 31$ ,  $0.54 \leq \Delta y^+ \leq 705$ , and  $\Delta z^+ = 15$  in viscous (inner) units. The time step is chosen to be  $\Delta t = 0.0064$ , which corresponds to approximately  $\Delta t^+ \approx 0.02$ . The flow is initialised to a laminar Blasius solution through the entire domain and allowed to develop until  $t = 1500$  ( $t^+ \approx 4755$ ) before the recording of the statistics begins. The spanwise-averaged skin friction profile along most of the streamwise extent of the domain ( $x = 35\text{--}650$ ) is monitored at each time step and the mean squared difference between successive time steps is used as a stopping criterion. To accelerate the transition to turbulence, a random volume

**Fig. 2** Illustration of the wall-normal blowing profile and parameterisation. The locations of the control points (spline knots) are fixed, whereas their amplitudes are controllable



forcing approach (Schlatter and Örlü 2010), located at  $x = 3.5$ , is used to trip the boundary layer (orange shaded region in Fig. 1).

## 2.2 Blowing Configuration

The wall-normal blowing is implemented as a velocity boundary condition and extends from  $x = 68$ – $145$  in the streamwise direction, resulting in a length of  $L_B = 77$ . This corresponds to a Reynolds number range of  $Re_\theta \approx 479$ – $703$  for the canonical case. The blowing configuration consists of a streamwise-varying profile, parameterised by a cubic spline with five control points (knots), as shown in Fig. 2. The streamwise locations of the control points are equally spaced from the start to the end of the blowing region and are fixed throughout the entire optimisation campaign. Furthermore, the amplitudes of the first and last control points are set to zero to ensure a smooth transition from the non-blowing to blowing region (and vice-versa). Therefore, the amplitudes of the three inner control points ( $A_1$ ,  $A_2$ , and  $A_3$ ) are controllable and comprise the only inputs to the Bayesian optimisation campaign. The permitted ranges for each input during the optimisation campaign are chosen to be 0%–1% of the freestream velocity. This range was selected based on preliminary work, as well as the work performed in Mahfoze et al. (2019), which indicated that there is an optimum in the NES within this range for the uniform blowing configuration. Once values for the input amplitudes are suggested by the Bayesian optimisation algorithm, boundary conditions for the outer control points are all that is required to completely define the cubic spline profile and close the system. In the present work, zero-gradient boundary conditions are applied to ensure a smooth transition in the blowing velocity at the edges of the control region. Note that in the present work, the control amplitudes are independent of time, so that the resulting blowing profile is static. Furthermore, compared to the work of Mahfoze et al. (2019), which focussed on an intermittent blowing configuration, the present work adopts a streamwise-varying profile of wall-normal blowing. While intermittent blowing is known to perform well in terms of NES, the present setup provides more flexibility in terms of the shape of the blowing profile.

While the amplitudes of the control points are enforced to always be positive (corresponding to blowing), it is still possible that for certain inputs the spline will become negative in regions between the control points (corresponding to suction). This is a feature of the cubic spline interpolation and arises due to the fact it generates curves that are  $C^2$  smooth. Although this smoothness is beneficial from a numerical perspective, suction is known to increase skin-friction drag in a zero-pressure gradient TBL (Kametani and Fukagata 2011). Nevertheless, no action was taken to prevent this behaviour as it only occurs under certain combinations of inputs and therefore its effect is expected to be minimal. Furthermore, implementing such a constraint would add unnecessary complexity since if

such behaviour is truly of no benefit then the Bayesian optimisation algorithm will automatically learn to avoid these regions of the parameter space. Another option would be to replace the standard cubic spline interpolation with a monotonicity-preserving approach, such as the monotonic cubic interpolation proposed by Fritsch and Carlson (1980). However, these methods generally sacrifice continuity in the second derivative to prevent the overshoot. Therefore, the effect of this reduction in smoothness on the numerical stability of the flow solver is left for future work.

### 2.3 Performance Metrics

While skin-friction drag reduction is of benefit in many applications, the key metric for the present work is NES, which takes into account both the drag reduction and the input power required to drive the control device. Therefore, a global scalar measure of the NES is required to pass to the Bayesian optimisation algorithm as an objective function. The dimensionless skin-friction coefficient is given by:

$$\overline{C_f}(x) = \frac{\overline{\tau_w}(x)}{\frac{1}{2}\rho u_\infty^2} \tag{1}$$

where  $\overline{\tau_w}(x)$  is the local wall shear stress and  $\rho$  is the mass density. Note that spanwise (periodic) averaging is implicitly assumed. Based on Eq. (1), the local skin-friction drag reduction (LDR) is defined as:

$$\text{LDR}(x) = \frac{\overline{C_{f,0}}(x) - \overline{C_f}(x)}{\overline{C_{f,0}}(x)} \tag{2}$$

where the zero subscript refers to the canonical case (i.e. the baseline case with no blowing). The blowing control has a long-lasting downstream effect on the skin-friction drag that is not just restricted to the control region. To capture this effect, a global skin-friction drag reduction (GDR) is obtained by integrating the skin-friction coefficient along most of the streamwise extent of the domain, from  $x = 35$ –650:

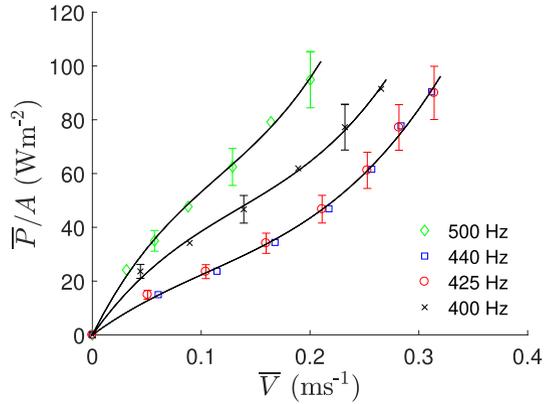
$$\hat{C}_f = \frac{1}{L_{GDR}} \int_{x=35}^{x=650} \overline{C_f}(x) dx \tag{3}$$

where the hat notation denotes a globally-averaged quantity and  $L_{GDR} = 615$  is the length over which the GDR is calculated, which covers 82% of the streamwise length of the domain and corresponds to approximately  $Re_{\theta,0} \approx 387$ –1919 for the canonical case. The GDR is thus given by:

$$\text{GDR} = \frac{\hat{C}_{f,0} - \hat{C}_f}{\hat{C}_{f,0}} \tag{4}$$

To calculate the NES, an estimate for the input power required to drive the control device is necessary. In numerical studies, the control is often applied as a simple boundary condition, with no possibility to directly evaluate the power required. Therefore, an alternative approach to estimating the input power is necessary. Several previous studies have estimated this using first principles, where the input power is related to the cube of the blowing velocity and the pressure difference across the blowing surface (Kametani and Fukagata

**Fig. 3** Time-averaged power per unit area with respect to the time-averaged wall-normal velocity for the present blowing device, operating over 400–500 Hz. Reproduced from Mahfoze et al. (2019)



2011; Kametani et al. 2015, 2016; Mahfoze et al. 2019). However, this is an idealised estimate as it does not take into account the losses associated with a real-world device (e.g. mechanical losses). Therefore, the present work follows Mahfoze et al. (2019) by estimating the input power from experimental data of a real-world blowing device. The device in question uses miniature electromagnetic speakers to blow air through a micro-perforated plate. Figure 3 shows experimental measurements of the time-averaged input power for the blowing device, as a function of the time-averaged blowing velocity, at different operating frequencies (Mahfoze et al. 2019). For this work, the lower uncertainty limits of the 425 Hz curve is selected and a third-order polynomial is fitted to provide a relationship between the wall-normal velocity and input power per unit area. To retain consistency with Mahfoze et al. (2019), a reference velocity of  $u_{ref} = 21$  m/s is chosen to provide the scaling between physical and non-dimensional units. For more details the reader is referred to Mahfoze et al. (2019).

The local power per unit area ( $\bar{P}/A$ ) can be calculated from the cubic spline velocity profile and the third-order polynomial fitted to the experimental data of Mahfoze et al. (2019). This can then be integrated over the control region to obtain the total input power:

$$\bar{P} = L_z \int_{x=68}^{x=145} \bar{P}/A(x) dx \tag{5}$$

which in dimensionless form is:

$$\hat{C}_w = \frac{\bar{P}}{\frac{1}{2} \rho L_{GDR} L_z u_\infty^3} \tag{6}$$

Note that the normalisation here is with respect to  $L_{GDR}$ , as opposed to  $L_B$ , to be consistent with Eq. (3). Finally, the NES is defined as:

$$NES = \frac{\hat{C}_{f,0} - (\hat{C}_f + \hat{C}_w)}{\hat{C}_{f,0}} \tag{7}$$

which is the objective function that is to be maximised by the Bayesian optimisation algorithm.

## 3 Methods

### 3.1 Flow Solver

The numerical simulations presented in this work are performed with the high-order compact finite-difference flow solver `Incompact3d`<sup>1</sup> (Laiet and Lamballais 2009), which is part of the open-source framework of flow solvers `Xcompact3d` (Bartholomew et al. 2020). The implicit large-eddy simulation (ILES) approach adopted in this work introduces targeted numerical dissipation at the small scales through the discretisation of the second derivatives of the viscous term and does not involve the convective term (e.g. through upwinding). The governing equations are the unsteady three-dimensional incompressible Navier–Stokes equations, given in non-dimensional form as:

$$\frac{\partial u_i}{\partial x_i} = 0 \quad (8)$$

$$\frac{\partial u_i}{\partial t} = -\frac{1}{2} \left( u_j \frac{\partial u_i}{\partial x_j} + \frac{\partial(u_i u_j)}{\partial x_j} \right) - \frac{\partial p}{\partial x_i} + \nu \mathcal{D} + F_i \quad (9)$$

where  $i = 1, 2, 3$  corresponds to the  $x, y,$  and  $z$  directions, respectively,  $u_i$  is the velocity vector,  $p$  is pressure, and  $F_i$  accounts for any additional forcing (e.g. random volume forcing for the tripping). Since Eqs. (8) and (9) do not make any reference to a particular explicit filter, both  $u_i$  and  $p$  can be interpreted as mesh-resolved quantities. The sub-grid scales are modelled via the hyper-viscous momentum diffusion term,  $\mathcal{D}$ , which is defined as:

$$\mathcal{D} = \frac{\partial^2 u_i}{\partial x_j \partial x_j} + \frac{\partial}{\partial x_j} \left( Q_c * \frac{\partial u_i}{\partial x_j} \right) \quad (10)$$

where the star notation denotes a convolution and  $Q_c$  is a hyper-viscous kernel used to construct an ILES operator via the convolution. For brevity, further details are not provided here. However, the reader is referred to Dairay et al. (2017), Deskos et al. (2019), Frantz et al. (2021), Mahfoze and Laiet (2021) for more information about the present ILES strategy.

The degree of numerical dissipation is controlled by a single parameter ( $\nu_0/\nu$ ), which corresponds to the added dissipation at the cutoff wavenumber, and the range of scales over which the numerical dissipation is added is controlled by the parameter  $c_1$  (see Dairay et al. (2017) for more details). The present work follows the recommendation of Mahfoze and Laiet (2021), who determined suitable ranges for these parameters for wall-bounded turbulent flows. Therefore, for this work  $\nu_0/\nu = 20$  and  $c_1 = 0.1755$  (low Reynolds numbers)  $-0.351$  (high Reynolds numbers).

Sixth-order compact finite-difference stencils are used in this work to discretise the governing equations. Furthermore, an explicit third-order Adams-Bashforth scheme is adopted for time integration, which is combined with an implicit Crank-Nicolson scheme for the diffusive terms in the wall-normal direction to circumvent the stability constraints imposed by the non-uniform mesh resolution used to properly resolve the near-wall effects. The governing

<sup>1</sup> See <https://www.incompact3d.com> and <https://github.com/xcompact3d/Incompact3d>.

equations are solved in skew-symmetric form for increased stability and to reduce aliasing errors (Kravchenko and Moin 1997). The pressure Poisson equation (PPE), which enforces incompressibility, is solved entirely in spectral space via the use of relevant three-dimensional fast Fourier transforms (FFTs). Through the use of a modified wavenumber (Lele 1992), the divergence-free condition is ensured up to machine accuracy. To avoid the spurious pressure oscillations observed in fully-collocated approaches (Laizet and Lamballais 2009), the pressure field is defined offset (by half a mesh width) with respect to the velocity field. The simplicity of the structured mesh allows easy implementation of a two-dimensional domain decomposition strategy, based on pencils, using the message passing interface (MPI) (Laizet and Li 2011). The computational domain is split into several subdomains (pencils), each of which are assigned to an MPI process. The derivatives and interpolations in the  $x$ ,  $y$ , and  $z$  direction are performed from within the  $X$ ,  $Y$ , and  $Z$  pencils, respectively. The three-dimensional FFTs required by the PPE solver are performed as a series of one-dimensional FFTs, computed in one direction at a time. Global transpositions to switch from one pencil to another are performed via the MPI command `MPI_ALLTOALL(V)`. `Incompact3d` has been extensively validated on a variety of turbulent flows (Diaz-Daniel et al. 2017; Mahfoze and Laizet 2017; Deskos et al. 2019) and shown to scale well on up to hundreds of thousands cores on CPU-based supercomputers (Laizet and Li 2011; Bartholomew et al. 2020).

### 3.2 Bayesian Optimisation

The present work adopts the Bayesian optimisation framework described in Diessner et al. (2022), a brief description of which is given here. Bayesian optimisation is an optimisation technique that relies on inexpensive surrogate models to locate the global optimum of black-box functions that are expensive to evaluate (Mockus 1994; Jones et al. 1998; Jones 2001; Brochu et al. 2010; Shahriari et al. 2016). Black-box functions are characterised by a non-existent or unknown mathematical expression. Therefore, the only possible way of gathering information about the underlying function is to provide a set of parameter inputs and observe their output (Gramacy 2020). This renders conventional optimisation methods that rely on gradient information (e.g. gradient descent) impractical, due to the large number of function evaluations typically required to estimate multidimensional gradients. Similarly, evolutionary algorithms that use a large number of function evaluations, such as differential evolution (Storn and Price 1997), cannot be used, due to the cost of each evaluation. Bayesian optimisation, on the other hand, does not require any information about the underlying objective function (e.g. gradients) and is designed to be sample efficient (i.e. it aims to find an approximate solution in as few function evaluations as possible).

Bayesian optimisation takes a sequential approach to optimisation, where the algorithm is run for many iterations, with each iteration suggesting the next point to evaluate from the objective function. Each iteration (optimisation step) consists of two parts. Firstly, a surrogate model, usually a Gaussian process (GP), is fitted to some observations to represent the objective function. A GP is a non-parametric model that is defined by a prior mean function  $\mu_0(\mathbf{x}_i)$  and a covariance kernel  $k(\mathbf{x}_i, \mathbf{x}_j)$ , where  $\mathbf{x}_i$  and  $\mathbf{x}_j$  are individual data points. Using a set of  $n$  initial observation pairs  $\mathcal{D}_n = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$  as training data, the posterior mean  $\mu_n(\cdot)$  and variance  $\sigma_n^2(\cdot)$  can be computed for a new point  $\mathbf{x}$  as:

$$\mu_n(\mathbf{x}) = \mu_0(\mathbf{x}) + \mathbf{k}(\mathbf{x})^T (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} (\mathbf{y} - \mathbf{m}) \quad (11)$$

$$\sigma_n^2(\mathbf{x}) = k(\mathbf{x}, \mathbf{x}) - \mathbf{k}(\mathbf{x})^T (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{k}(\mathbf{x}) \quad (12)$$

where  $\mathbf{K} = k(\mathbf{x}_{1:n}, \mathbf{x}_{1:n})$ ,  $\mathbf{k}(\mathbf{x}) = k(\mathbf{x}, \mathbf{x}_{1:n})$ ,  $\mathbf{m} = \mu_0(\mathbf{x}_{1:n})$ , and  $\sigma^2$  is the observational noise. These can be used to compute both the prediction and the corresponding uncertainty for each point  $\mathbf{x}$  of the parameter space (Mockus 1994; Brochu et al. 2010; Shahriari et al. 2016). The parameters in the mean function and the covariance kernel (e.g. constant mean value, one or multiple length-scales, signal variance) can be estimated from the training data  $\mathcal{D}_n$  via maximum likelihood estimation, maximum a posteriori estimation, or a fully Bayesian approach. The exact parameters depend on the choice of prior mean function and covariance kernel. For a more detailed discussion on GPs the reader is referred to Rasmussen and Williams (2005); Gramacy (2020).

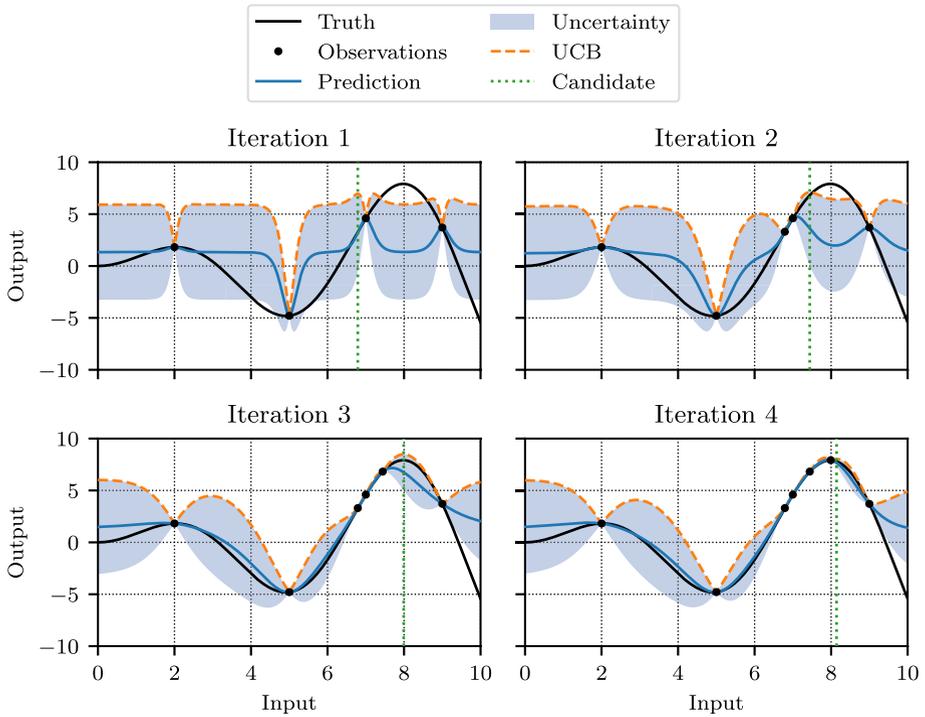
The second part of the optimisation step is to select the next point to evaluate from the objective function by optimising a user-defined acquisition function. There are several available acquisition functions to choose from, each with different properties, and all of them rely on the predictions and/or corresponding uncertainties of the surrogate model. One popular acquisition function is Upper Confidence Bound (UCB) (Srinivas et al. 2010), given by:

$$\alpha_{UCB}(\mathbf{x}; \mathcal{D}_n) = \mu_n(\mathbf{x}) + \sqrt{\beta} \sigma_n(\mathbf{x}) \quad (13)$$

Equation (13) shows how both the posterior mean  $\mu_n(\mathbf{x})$  and variance  $\sigma_n(\mathbf{x})$  are used to compute the acquisition of point  $\mathbf{x}$ . UCB assumes that the uncertainty associated with the prediction of the GP is true and selects the maximum value as the next point to evaluate from the objective function (Brochu et al. 2010; Shahriari et al. 2016). An important property of most acquisition functions is the exploitation-exploration trade-off. By balancing these two aspects, the Bayesian optimisation algorithm not only selects points where the prediction is high (exploitation), but also points where the uncertainty is high (exploration) (Jones et al. 1998). If the focus relies too heavily on exploitation the algorithm is more likely to get stuck in a local optimum, whereas if the focus is on exploration the algorithm is less likely to focus on promising areas, making it less likely to find optimal solutions. To control the balance between exploitation and exploration, UCB includes a tunable hyperparameter ( $\beta$ ), where high values favour more exploration and low values favour more exploitation.

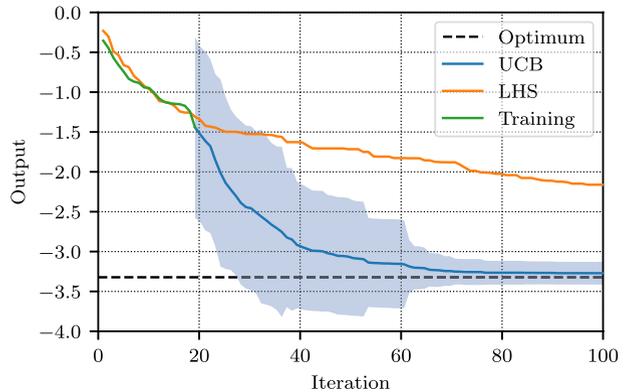
It is worth noting that the optimisation setup in this work follows the Monte Carlo approach discussed in Snoek et al. (2012) and Wilson et al. (2018). Instead of directly computing the acquisition function in Equation (13), samples are drawn from the predictive multivariate normal distribution  $\mathcal{N}(\mu_n(\mathbf{x}), \sigma_n^2(\mathbf{x}))$  and are used to compute the acquisition. This allows for the selection of multiple points at each iteration, for which the analytical acquisition functions would be intractable. In particular, a greedy optimisation strategy is implemented that computes a batch of points by sequentially selecting them. Points in the batch that have already been selected are assumed to be fixed by including them in the predictive distribution. Batches of points can then be evaluated in parallel through multiple simultaneous simulations. In summary, with respect to the present work, the Bayesian optimisation algorithm suggests different blowing configurations to run. These are then executed by `Incompact3d` and the resulting NES for each configuration is calculated. Finally, these new samples are used to update the Bayesian optimisation algorithm and the next set of blowing configurations is suggested.

Figure 4 shows four iterations of the Bayesian optimisation algorithm on a simple test function, given by  $f(x) = -x \sin(x)$ . In this example, the GP is initialised with a constant



**Fig. 4** Four iterations of the Bayesian optimisation algorithm on a simple 1D test function

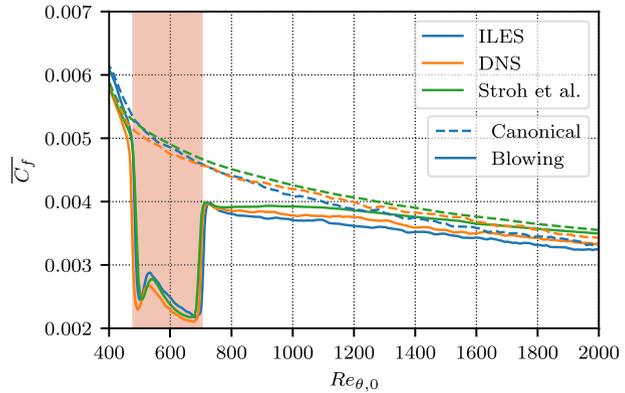
**Fig. 5** Ensemble-averaged value for the objective function at each Bayesian optimisation iteration on the 6D Hartmann function. For the UCB data, the solid line represent the mean across 30 separate Bayesian optimisation runs, whereas the shaded region indicates the 95% confidence interval. The baseline (LHS) is also shown for comparison



mean function and a Matérn 5/2 kernel. The trade-off parameter ( $\beta$ ) is chosen to correspond to the upper range of the 95% confidence interval, depicted with the blue area around the predicted mean of the GP. The prediction of the GP approaches the true objective function with each Bayesian optimisation iteration and locates a solution close to the optimum after three iterations.

Figure 5 shows the results of Bayesian optimisation applied to the more complex Hartmann function. This function accepts six inputs and has six local minima, in addition to

**Fig. 6** Spanwise-averaged skin friction coefficient with respect to the canonical momentum Reynolds number for both the canonical and uniform blowing cases using both ILES and DNS. The DNS data of Stroh et al. (2016) is also shown for comparison. The red shaded area indicates the blowing region

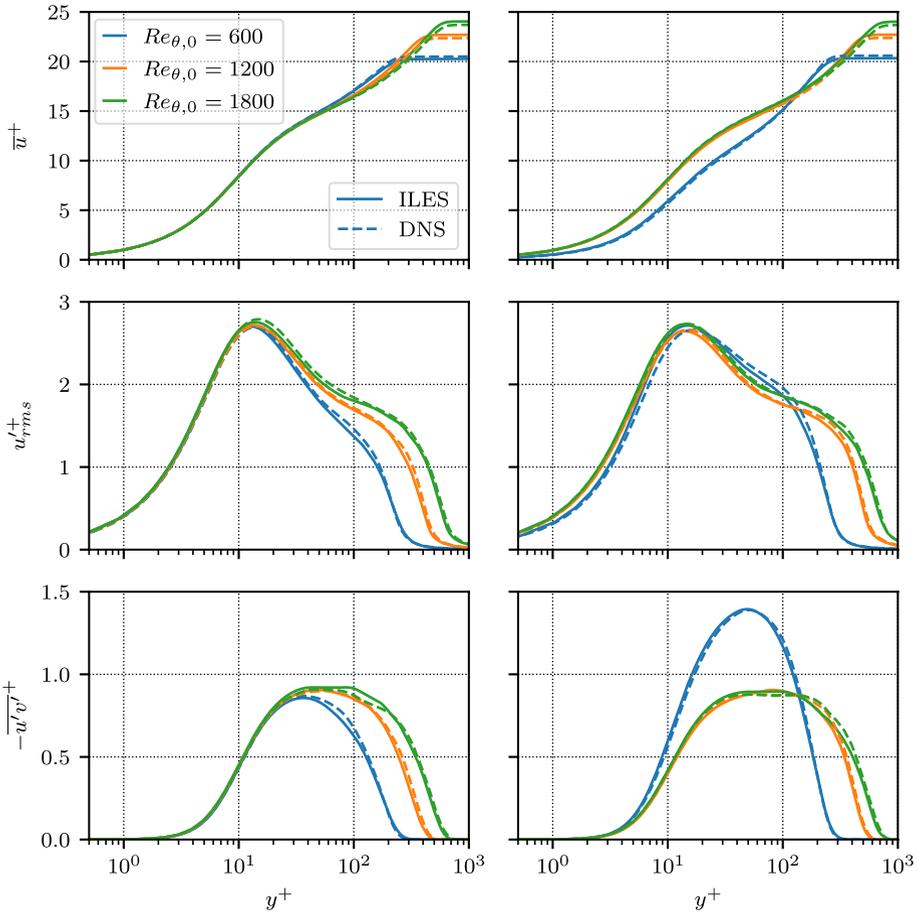


one global minimum at  $\mathbf{x}^*$  with a value of  $f(\mathbf{x}^*) = -3.32237$ . To ensure robust results, the experiment is repeated 30 times, each with an evaluation budget of 100. As a baseline, the mean of 30 maximin Latin hypercube designs, which aim to cover the most space with the available evaluations, is also shown. All of the 30 Bayesian optimisation runs allocate 18 evaluations of the budget as training data, which are also sampled by a Latin hypercube, and use the remaining budget to perform 82 iterations of the Bayesian optimisation algorithm. Figure 5 shows that Bayesian optimisation performs significantly better than the baseline. Furthermore, while the variability between runs is initially quite large immediately following the training data, it decreases with the number of evaluations. This suggests that in all 30 Bayesian optimisation runs a close-to-optimal solution is found. For more details on the Bayesian optimisation framework and its validation, the reader is referred to Diessner et al. (2022).

## 4 Validation

The purpose of this section is to validate the present ILES scheme against the existing DNS scheme within *Incompact3d* and benchmark DNS data from Stroh et al. (2016). The case setup is exactly the same as described in Sect. 2.1, except the control is modified to a uniform blowing profile with an amplitude of 0.5% of the freestream velocity. For the DNS, the mesh resolution is doubled so that  $n_x \times n_y \times n_z = 3073 \times 513 \times 256$  and the time step is halved to  $\Delta t = 0.0032$ . All of the simulations in this section, and the following sections, were performed on the ARCHER2 UK national supercomputer, which has 5860 nodes housing two AMD EPYC 7742 CPUs ( $2 \times 64$  total cores @ 2.25GHz) in a non-uniform memory access arrangement. Network communication is via the HPE Slingshot interconnect, which provides  $2 \times 100$ Gbit/s of bidirectional bandwidth. On this system, the ILES simulations took between 8 and 14 hours on 24 nodes (3072 cores), depending on how long the statistics took to converge, whereas the DNS simulations took between 47 and 48 hours on 128 nodes (16384 cores). This clearly demonstrates the performance benefits of the present ILES approach over DNS.

Figure 6 shows the skin-friction coefficient with respect to the canonical momentum Reynolds number for both the canonical and blowing cases and ILES and DNS schemes. The red-shaded area indicates the blowing region. The agreement between the ILES, DNS, and Stroh et al. (2016) data is good, despite the use of different numerical methods, with



**Fig. 7** Spanwise-averaged wall-normal profiles of selected velocity moments at selected canonical momentum Reynolds numbers for both the canonical (left) and uniform blowing (right) cases using both ILES and DNS. Note that for this figure the inner scaling is with respect to the friction velocity of the canonical case, calculated via each method, at the given momentum Reynolds number

both the behaviour over the control region and its lasting downstream effect properly captured. For the blowing cases, there is a sharp decrease in the skin-friction drag at the start of the control, followed by a small recovery about 25% along the blowing region, before a further decrease towards the global minimum just before the end of the control. Following the blowing region, there is a rapid recovery of the skin-friction drag. However, this is only a partial recovery, with a subtle but sustained long-lasting downstream effect that persists across the length of the domain. Both the ILES and DNS results display a slight drift in the skin friction coefficient towards the end of the domain when compared with Stroh et al. (2016). However, overall the agreement is good and demonstrates the validity of the present ILES scheme.

Figure 7 shows the mean and fluctuating streamwise velocity and the Reynolds stress with respect to the wall-normal coordinate at various streamwise locations (expressed in terms of the canonical momentum Reynolds number) for both the canonical and blowing

cases and ILES and DNS schemes. Note that for this figure, the inner scaling for each method (ILES, DNS) is with respect to the friction velocity of the canonical case, at the given momentum Reynolds number, as calculated using the respective method. The agreement is excellent and further demonstrates the validity of the present ILES approach, with both the main and secondary peaks properly captured for both the canonical and blowing cases. The effect of the control is especially evident at  $Re_{\theta,0} = 600$ , which is approximately in the centre of the blowing region, where both the magnitude of the peak in Reynolds stress, and its wall-normal location, is significantly modified. Of particular note is the reduction in mean streamwise velocity close to the wall, which must be properly captured to ensure the computed skin-friction coefficient is accurate. In conclusion, these results demonstrate the validity of the present ILES approach and justify its use over the more expensive DNS in the following Bayesian optimisation campaign.

## 5 Results

This section presents the results from the optimisation campaign, followed by a detailed analysis of selected cases to investigate the key drag reduction mechanisms and flow physics. For the Bayesian optimisation algorithm, the GP is configured with a zero mean function and a Matérn 5/2 kernel. The Matérn 5/2 kernel was chosen over the very popular squared exponential kernel as it is able to represent functions that are less smooth and is therefore better suited to real-world problems (Snoek et al. 2012). The signal variance and the length-scales of the kernel are estimated via maximum likelihood estimation. The acquisition function for this work is chosen to be UCB, with a trade-off parameter  $\beta = 1$ , as it was found to perform well on a variety of different test problems (Diessner et al. 2022). The parallel strategy briefly described in Sect. 3.2 is adopted to accelerate the overall optimisation time, with four candidate points sampled at each iteration. To seed the optimisation campaign, 12 initial training points are selected using maximin Latin hypercube sampling. All simulations are performed with the present ILES approach. In all cases the flow is initialised to a laminar Blasius solution through the entire domain and the blowing velocity is ramped up from  $t = 0$ –100 to prevent any numerical instabilities. Once the recordings of the statistics begin (at  $t = 1500$ ), the spanwise-averaged skin friction profile along the GDR region ( $x = 35$ –650) is monitored at each time step and the mean squared difference between successive time steps is used as a stopping criterion for each individual case. The time for collecting statistics is therefore different for each case, but all cases fall within the range  $t_{stat} \approx 2100$ –4500 ( $t_{stat}^+ \approx 6660$ –14260), resulting in a compute time of 8–14 hours on 24 nodes (3072 cores).

### 5.1 Optimisation Campaign

Table 1 summarises the results of the Bayesian optimisation campaign. Each row represents an individual simulation and displays the associated iteration number, input amplitudes, maximum LDR (MLDR), GDR and NES. Note that all quantities are given as a percentage, with the input amplitudes given as a percentage of the freestream velocity and constrained to be within 0–1% (as discussed in Sect. 2.2). The grey shaded area denotes the initial training (seed) data obtained via Latin hypercube sampling.

The objective function for the optimisation campaign is the NES, which is to be maximised. Focussing on this column first, it can be seen that, while many cases come close to

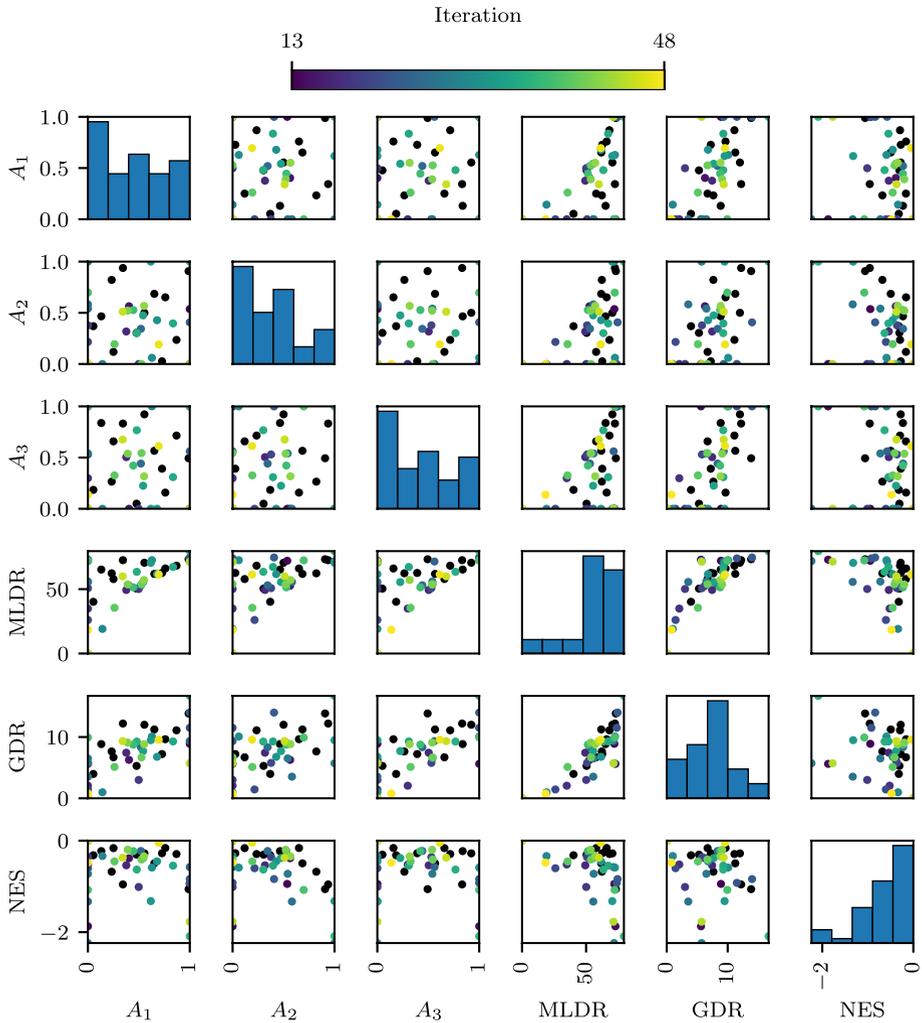
**Table 1** Summary of the Bayesian optimisation campaign showing the iteration ID, input amplitudes, MLDR, GDR, and NES

ID	$A_1$	$A_2$	$A_3$	MLDR	GDR	NES	ID	$A_1$	$A_2$	$A_3$	MLDR	GDR	NES
1	0.25	0.12	0.66	57.85	6.69	-0.16	25	0.52	0.34	0.44	53.01	8.34	-0.23
2	0.55	0.50	0.92	70.52	12.01	-0.28	26	0.00	0.58	0.00	55.60	3.53	-1.33
3	0.05	0.37	0.18	40.18	3.97	-0.32	27	0.00	0.54	0.54	55.91	6.50	-0.72
4	0.34	0.94	0.83	72.51	12.20	-0.96	28	0.00	1.00	0.00	71.81	5.72	-2.24
5	0.87	0.24	0.72	68.23	11.10	-0.29	29	0.14	0.00	0.00	19.11	0.98	-0.34
6	0.48	0.30	0.05	50.40	5.29	-0.31	30	0.63	0.06	1.00	72.51	10.08	-0.59
7	0.13	0.47	0.84	65.21	8.85	-0.23	31	0.68	0.43	0.00	61.67	6.77	-0.45
8	0.76	0.65	0.16	66.05	9.64	-0.16	32	0.62	1.00	0.00	70.43	9.49	-1.32
9	0.99	0.91	0.49	73.22	13.85	-1.06	33	0.54	0.48	0.22	55.50	7.70	-0.38
10	0.73	0.03	0.39	62.59	7.22	-0.28	34	0.84	0.40	0.31	67.23	9.28	-0.55
11	0.65	0.69	0.57	61.83	11.24	-0.50	35	0.48	0.30	0.77	63.47	9.26	-0.63
12	0.23	0.82	0.26	62.30	7.64	-0.68	36	1.00	1.00	1.00	79.52	16.72	-2.09
13	1.00	0.54	0.00	72.06	8.90	-0.95	37	0.45	0.53	0.42	51.39	8.76	-0.23
14	0.00	0.00	1.00	72.88	5.62	-1.87	38	0.00	0.70	1.00	72.32	9.93	-1.08
15	0.40	0.56	0.00	52.65	6.27	-0.38	39	0.26	0.19	0.33	35.53	5.12	-0.45
16	0.00	0.00	0.00	0.00	0.00	0.00	40	0.53	0.00	0.54	53.83	6.68	-0.41
17	0.38	0.32	0.50	49.35	7.44	-0.54	41	0.39	0.52	0.54	53.65	9.09	-0.20
18	0.00	0.00	0.30	34.92	2.05	-0.52	42	0.55	0.57	0.32	57.10	8.83	-0.34
19	0.50	0.00	0.00	50.11	2.99	-1.01	43	0.00	0.00	0.00	0.00	0.00	0.00
20	0.00	0.38	0.53	50.77	5.77	-0.36	44	1.00	0.00	0.00	71.35	5.72	-1.77
21	0.00	0.22	0.00	26.04	1.44	-0.60	45	0.34	0.51	0.68	59.98	9.34	-0.37
22	1.00	0.00	1.00	73.88	11.51	-0.92	46	0.00	0.00	0.00	0.00	0.00	0.00
23	1.00	0.41	1.00	74.57	14.00	-0.84	47	0.69	0.19	0.61	61.40	9.56	-0.05
24	1.00	0.00	0.56	71.48	9.34	-0.58	48	0.00	0.00	0.14	18.37	0.79	-0.48

Note that all quantities are given as a percentage, with the input amplitudes given as a percentage of the freestream velocity. The grey shaded area denotes the initial training (seed) data whereas the blue, orange and green rows highlight the best NES (after neglecting the trivial cases), best GDR and worst NES/GDR cases, respectively

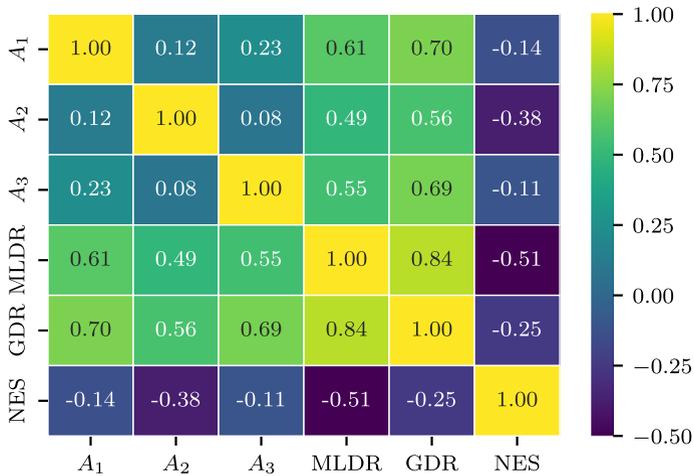
achieving a positive NES, there are no cases that improve on the trivial (canonical) case. This is in spite of the fact that many cases exhibit significant GDR compared to the canonical case (up to nearly 80% for the MLDR). Reassuringly, by examining the GDR and NES columns together, it can be seen that even a minor improvement in the efficiency of the blowing device would lead to positive NES. However, this would also modify the shape of the objective function, and so would likely change the trajectory of the optimisation campaign. It should be noted that reducing the reference velocity used to estimate the power input for the blowing device would also improve the NES estimates. However, as discussed in Sect. 2.3, for the purpose of retaining consistency with Mahfoze et al. (2019), this is not done. The main conclusion from Table 1 is that it is certainly possible to obtain meaningful NES just by improving the efficiency of the blowing device. For example, improving the efficiency by approximately 10% would lead to NES values of approximately 1% for certain cases in Table 1.

Overall, the best performing cases are 16, 43, and 46, which are all characterised by zero blowing across the entire control region and are therefore equivalent to the trivial (canonical) case. Note that Bayesian optimisation algorithms are often designed to work with stochastic objective functions, and so are often permitted to sample the same input point multiple times. For the present work, where the objective function is deterministic, this results in the exact same output and is therefore unnecessary. However, rather than



**Fig. 8** Scatter plot matrix showing each input, MLDR, GDR, and NES for each iteration. The colour bar indicates the iteration number in the scatter plots, with the training data given in black

implementing some additional logic to skip repeated samples, they are presented here as an indicator of the convergence of the Bayesian optimisation algorithm. The first time the optimisation algorithm suggests the trivial case it continues to explore the parameter space. However, towards the end of the campaign it suggests the same point twice in quick succession. This indicates a shift from exploration to exploitation, which itself suggests a relatively low degree of uncertainty in the underlying GP model. While the stopping criterion for the present study was ultimately determined by the computational budget, this observation can be used as an indicator (but not a guarantee) of at least some degree of convergence. After neglecting the trivial cases, the next best NES case (47) is selected for further analysis in Sect. 5.2, as is the best GDR case (36) and the worst NES/GDR case (28), all of which are highlighted in Table 1.

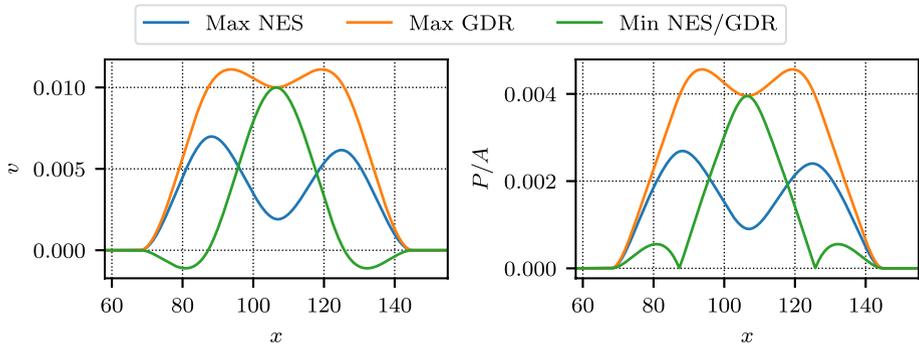


**Fig. 9** Correlation matrix illustrating the Pearson correlation coefficient for each input, MLDR, GDR, and NES

To aid with the analysis of Table 1, Fig. 8 displays a scatter plot matrix showing the pairwise relationship between each input and output of interest. Histograms of each quantity are also shown along the diagonal. The colour scale in the scatter plots indicates the iteration number, with the training data given in black. Examining the histograms first, it can be seen that the NES distribution, although negative, is skewed towards positive values. This is expected since the goal of the Bayesian optimisation campaign is to maximise the NES. On the other hand, examination of the input amplitudes shows a slight clustering towards smaller values. This is most likely a result of the nonlinear relationship between the blowing amplitude and the NES, which is a function of both the GDR and input power (which are themselves nonlinear functions of the blowing velocity). The scatter plots show a strong correlation between the MLDR and GDR, as well as a moderate correlation between both the MLDR and GDR and the input amplitudes. This is also expected since previous works on uniform blowing have shown a direct (almost linear) relationship between the blowing velocity and the drag reduction at low intensities (Kametani and Fukagata 2011). To quantify this relationship, Fig. 9 shows the pairwise correlation coefficients for each input and output. These results confirm there is a strong correlation between the MLDR and GDR, and that both the MLDR and GDR are moderately correlated with the input amplitudes. On the other hand, there does not seem to be any relationship between the NES and any of the other quantities. This is perhaps not surprising, given the nonlinear relationship between the blowing amplitudes and blowing profile, as well as the nonlinear relationships relating the blowing profile to the power input and GDR (on which the NES depends).

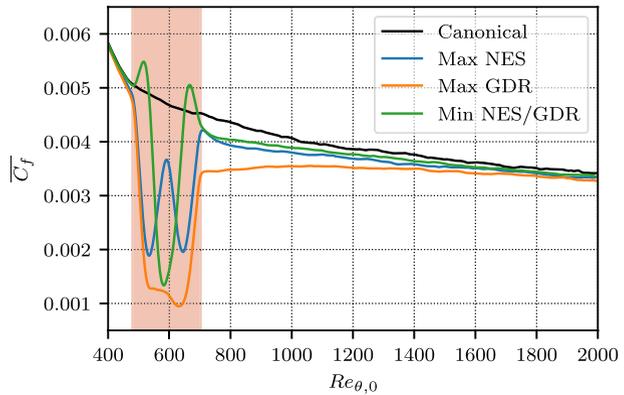
## 5.2 Flow Physics and Analysis

In this section, selected cases are chosen for further analysis. From Table 1, these are the best NES (after neglecting the trivial cases), best GDR and worst NES/GDR cases and are highlighted in blue, orange, and green respectively. Figure 10 shows the blowing velocity profile and associated power profile for each of these cases. The maximum NES case is characterised by moderate intensity blowing at the edges ( $A_1$  and  $A_3$ ) with a large dip



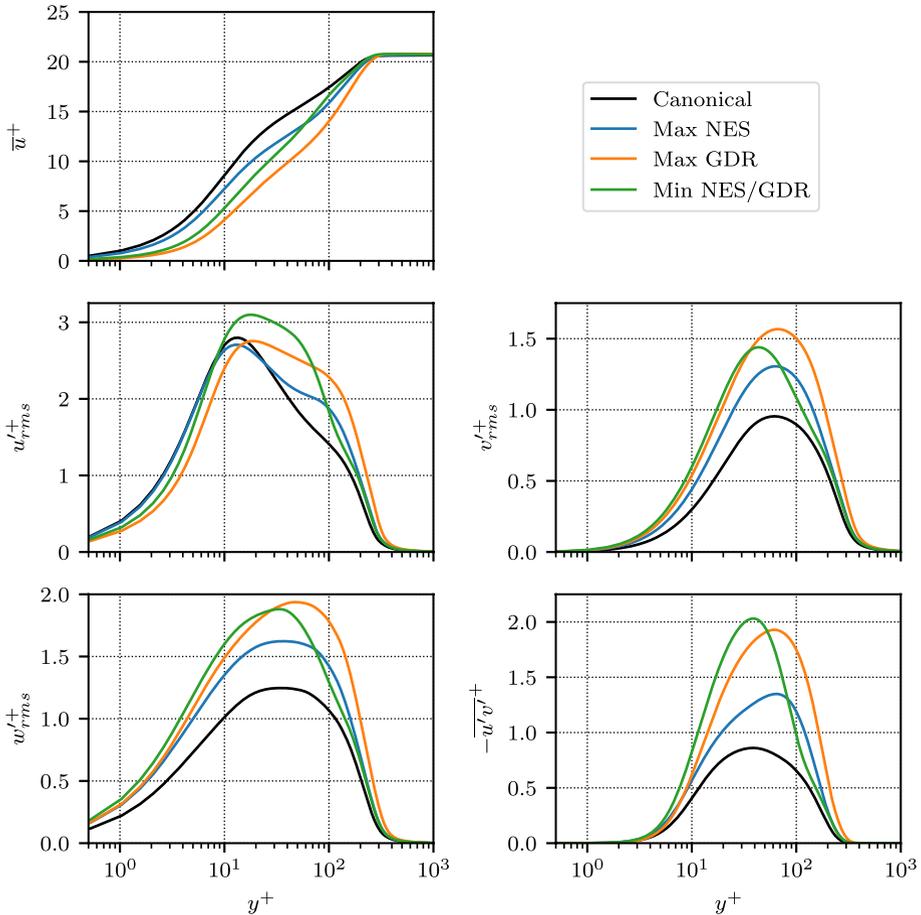
**Fig. 10** Blowing velocity profile (left) and the associated power profile (right) for selected cases

**Fig. 11** Spanwise-averaged skin friction coefficient with respect to the canonical momentum Reynolds number for selected cases. The canonical case is also shown for comparison. The red shaded area indicates the blowing region



in the centre of the blowing region at  $A_2$ . The maximum GDR case, on the other hand, is characterised by maximum intensity blowing for all amplitudes and is resemblant of a uniform blowing profile. Finally, the worst NES/GDR case is characterised by maximum intensity blowing in the centre of the blowing region, with zero blowing at  $A_1$  and  $A_3$ . Note that for this combination of inputs the blowing profile is negative (corresponding to suction) towards the edges of the blowing region. The power profiles are mostly similar to the velocity profiles, apart from some distortion, due to the nonlinear relationship between the blowing velocity and blowing power, and the change in sign when the blowing profile is negative (reflecting the fact that both blowing and suction require power).

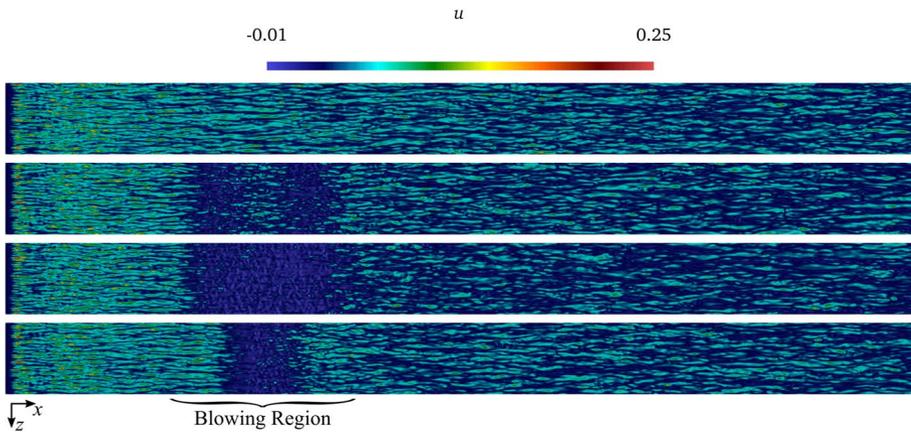
Figure 11 shows the skin friction coefficient with respect to the canonical momentum Reynolds number for both the canonical and selected blowing configurations. The red-shaded area indicates the blowing region. Starting with the maximum NES case, there is a sharp reduction in skin friction towards the start of the blowing region, which is associated with the moderate value for  $A_1$ . This is followed by a partial recovery of the skin friction in the centre of the blowing region, owing to the dip in blowing velocity at  $A_2$ . Finally, there is a sharp reduction in skin friction again at  $A_3$ . The effect of the control persists far downstream from the blowing region, which is important for its performance in terms of NES. For the maximum GDR case, there is a strong and sustained reduction in the skin friction over the entire length of the blowing region. Furthermore, the downstream effect is much



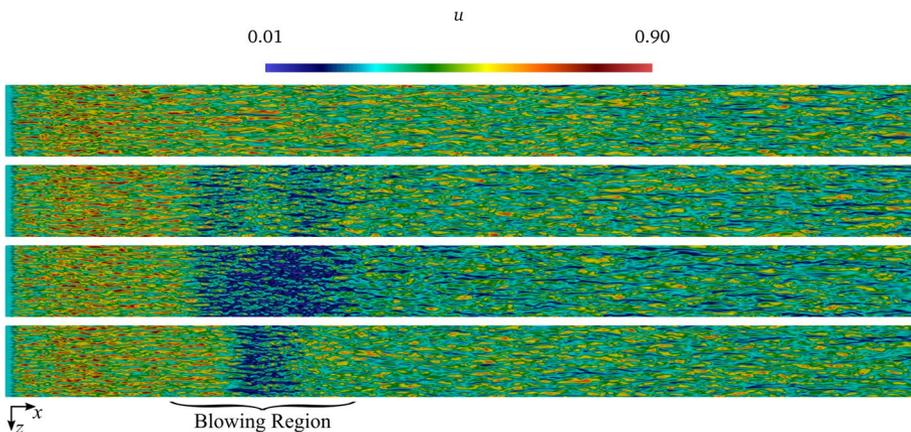
**Fig. 12** Spanwise-averaged wall-normal profiles of velocity moments for selected cases at a canonical Reynolds number of  $Re_{\theta,0} = 600$  (approximately in the centre of the blowing region). Note that for this figure the inner scaling is with respect to the friction velocity of the canonical case at the given momentum Reynolds number ( $Re_{\theta,0} = 600$ )

more prominent compared to the best NES case. However, even with significant GDR, this case does not perform well in terms of NES, owing to the large input power required to operate the blowing device at the maximum permitted intensity. Finally, for the worst NES/GDR case, the suction at the edges of the blowing region leads to a minor increase in the local skin friction, while still requiring power to drive the control. Although this case still exhibits a positive GDR, owing to the  $A_2$  amplitude in the middle of the control region, the suction regions are ultimately the main factor that leads to the poor performance, both in terms of GDR and NES, for this case.

Figure 12 shows the velocity moments and the Reynolds stress with respect to the wall-normal coordinate at a canonical momentum Reynolds number of  $Re_{\theta,0} = 600$  (approximately in the centre of the blowing region) for both the canonical and selected blowing configurations. Note that, for this figure, the inner scaling is with respect to the local friction velocity of the canonical case at  $Re_{\theta,0} = 600$ . There is a wide variation in behaviour



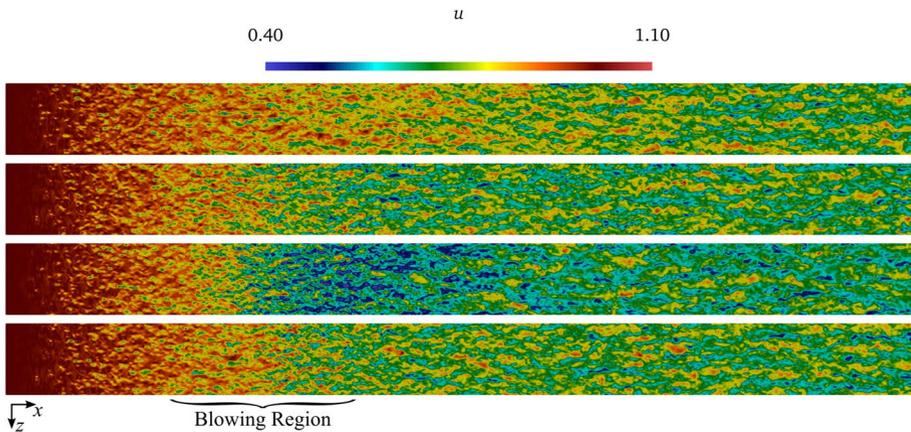
**Fig. 13** Slices of instantaneous streamwise velocity at  $y^+ = 1$ . From top to bottom the cases are the canonical case, maximum NES case, maximum GDR case, and minimum NES/GDR case



**Fig. 14** Slices of instantaneous streamwise velocity at  $y^+ = 10$ . From top to bottom the cases are the canonical case, maximum NES case, maximum GDR case, and minimum NES/GDR case

across each case. However, there are some commonalities that can be extracted. For example, close to the wall ( $y^+ < 10$ ) there is a reduction in the mean and fluctuating components of the streamwise velocity compared to the canonical case. For the mean streamwise velocity, this reduction is sustained across the entire height of the boundary layer, whereas in the case of the fluctuating component the behaviour reverses at approximately  $y^+ \approx 10$ . For the other velocity moments, the effect of the control is to increase their intensity across the entire height of the boundary layer, particularly the intensity of the peaks at around  $50 < y^+ < 100$ . This increase in turbulent activity is associated with the injection of kinetic energy into the boundary layer and has been reported in previous works (Kametani and Fukagata 2011; Kametani et al. 2015).

Figures 13, 14, and 15 show slices of the instantaneous streamwise velocity at  $y^+ = 1$  (Fig. 13),  $y^+ = 10$  (Fig. 14), and  $y^+ = 100$  (Fig. 15) over the first half of the streamwise

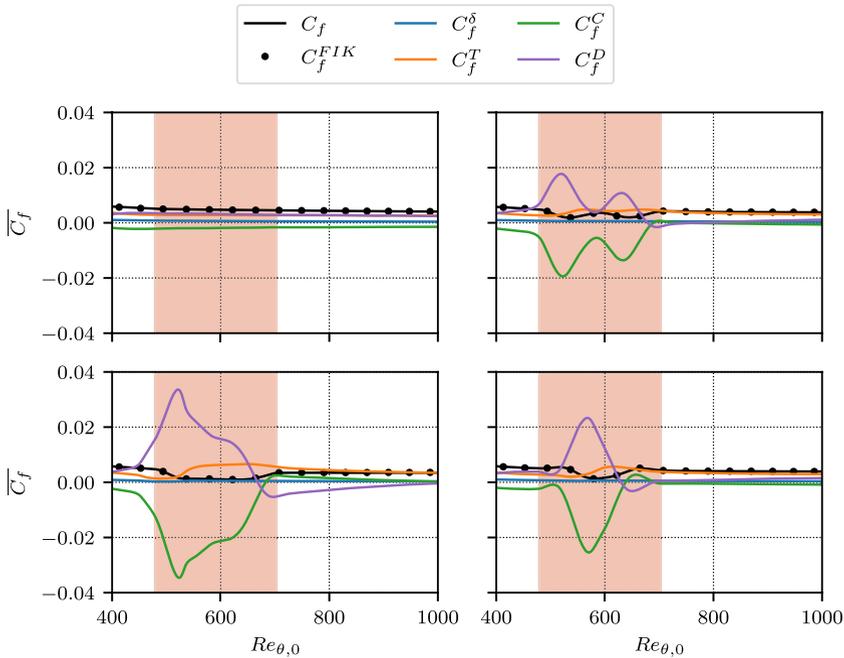


**Fig. 15** Slices of instantaneous streamwise velocity at  $y^+ = 100$ . From top to bottom the cases are the canonical case, maximum NES case, maximum GDR case, and minimum NES/GDR case

extent of the domain, for both the canonical and selected blowing configurations. From top to bottom the cases are the canonical case, best NES case, best GDR case, and worst NES/GDR case. As can be seen in Fig. 13, the effect of the control is clear, with a significant reduction in the streamwise velocity in regions where the blowing velocity is relatively high. The signature of the blowing profile is also imprinted on the velocity contours, with the best NES case displaying a slight increase in velocity in the centre of the blowing region (where the blowing velocity is reduced) and the worst NES/GDR case showing a short reduced-velocity region, with a slight increase in velocity at the start/end points of the blowing region (where the control corresponds to suction). While the streamwise velocity mostly recovers quite soon after the end of the blowing region, a subtle but lasting effect can be observed much further downstream, particularly for the maximum GDR case, where low momentum regions are more prevalent.

Looking at Fig. 14, the shape of the blowing profile is still evident. Furthermore, the lasting downstream effect is more obvious, particularly for the best GDR case, where the appearance of low-momentum streaks is more prominent. Finally, focusing on Fig. 15 the effect of the control is still clear, although the overall shape of the blowing profile is harder to discern. Furthermore, the effect of the control is shifted downstream with respect to the blowing region. This is most likely due to convection, leading to the effects of the control not appearing in higher regions of the boundary layer until further downstream. For the best GDR case, where the convective effect is very prominent, this could explain why the skin friction coefficient exhibits a very different downstream behaviour with respect to the other two cases (see Fig. 11).

To investigate the quantitative effect of different contributions to the drag reduction, the local skin friction coefficient is decomposed into four separate components using the so-called Fukagata-Iwamoto-Kasagi (FIK) identity (Fukagata et al. 2002). For a TBL, the FIK identity is defined as in Strohh et al. (2015):

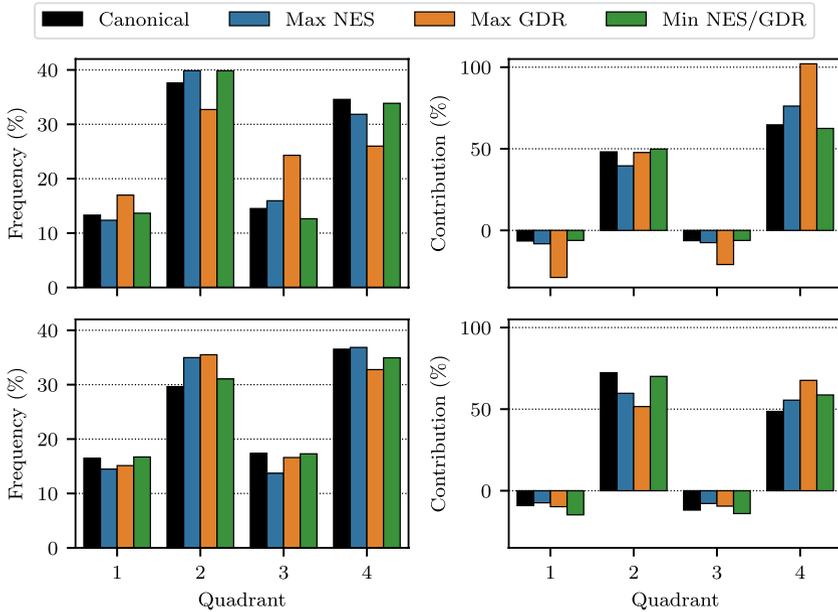


**Fig. 16** FIK identity with respect to the canonical momentum Reynolds number for the canonical case (top left), maximum NES case (top right), maximum GDR case (bottom left), and minimum NES/GDR case (bottom right)

$$\begin{aligned}
 \overline{C_f^{FIK}}(x) = & \underbrace{\frac{4(1 - \delta^*)}{Re_\delta}}_{C_f^\delta} + 4 \underbrace{\int_0^1 (1 - y)(-\overline{u'v'})dy}_{C_f^T} + 4 \underbrace{\int_0^1 (1 - y)(-\overline{u}\overline{v})dy}_{C_f^C} \\
 & - 2 \underbrace{\int_0^1 (1 - y)^2 \left( \frac{\partial \overline{u}^2}{\partial x} + \frac{\partial \overline{u'u'}}{\partial x} - \frac{1}{Re_\delta} \frac{\partial^2 \overline{u}}{\partial x^2} + \frac{\partial \overline{p}}{\partial x} \right) dy}_{C_f^D}
 \end{aligned} \tag{14}$$

where all quantities are non-dimensionalised with respect to the local boundary layer thickness ( $\delta$ ) and the freestream velocity ( $u_\infty$ ), and  $\delta^*$  is the local displacement thickness. The terms  $C_f^\delta$ ,  $C_f^T$ ,  $C_f^C$ , and  $C_f^D$  refer to the contributions arising from the boundary layer thickness, Reynolds shear stress (turbulence), mean convection, and spatial development, respectively.

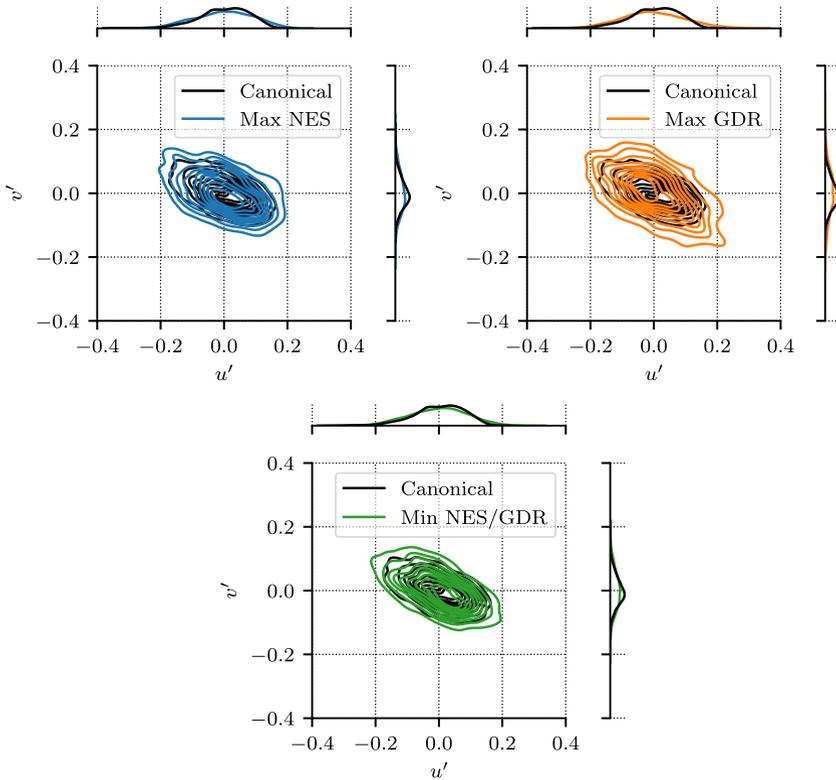
Figure 16 shows the decomposition of the local skin friction coefficient, using the FIK identity, with respect to the canonical momentum Reynolds number for both the canonical and selected blowing configurations. For comparison, the skin friction coefficient calculated directly, via Eq. (1), is also shown and matches the sum of the FIK components. The results show a large negative contribution from the mean convective term over the control region. This is also accompanied by a large positive contribution from the spatial development term. However, the mean convection term is more dominant and is ultimately the key factor in the skin friction drag reduction over the



**Fig. 17** Quadrant analysis at two probe locations for selected cases. The first probe (top) is located in the middle of the blowing region ( $x = 106.5$ ) at  $y^+ = 10$ , whereas the second probe (bottom) is located at the end of the control region ( $x = 145$ ) at  $y^+ = 100$ . Both the frequency of occurrence (left) and overall contribution to the time-averaged Reynolds stress ( $-\overline{u'v'}$ ) (right) is shown

control region. This same behaviour has been reported in previous works (Kametani and Fukagata 2011; Mahfoze et al. 2019) and is expected, since  $-\overline{u'v'} < 0$ . Interestingly, this behaviour is reversed downstream of the control for the best GDR case, where the spatial development term becomes negative and dominates the drag reduction, which is also reported in Mahfoze et al. (2019). For the best NES and worst NES/GDR cases, there is an initial reversal between the convective and spatial development terms towards the end of the control region. However, this reversal is short-lived and subsequently the convective term returns to a drag-reducing effect whereas the converse is true for the spatial development term.

Figure 17 shows the results of a quadrant analysis at two probe locations. The first probe is located in the middle of the blowing region ( $x = 106.5$ ,  $Re_{\theta,0} \approx 591$ ) at  $y^+ = 10$ , whereas the second probe location is at the end of the control region ( $x = 145$ ,  $Re_{\theta,0} \approx 703$ ) at  $y^+ = 100$ . Quadrant analysis is a data-processing technique that categorises each time-resolved measurement into a particular quadrant, according to the relative signs of the fluctuating streamwise and wall-normal velocities. The quadrants are labelled according to: Q1 ( $u' > 0$ ,  $v' > 0$ ), Q2 ( $u' < 0$ ,  $v' > 0$ ), Q3 ( $u' < 0$ ,  $v' < 0$ ), and Q4 ( $u' > 0$ ,  $v' < 0$ ). Both the frequency of occurrence of each quadrant, as well as its total contribution to the time-averaged Reynolds stress ( $-\overline{u'v'}$ ), is shown. The Q2 and Q4 quadrants are related to the drag-inducing ejection and sweep events, and work to increase the time-averaged Reynolds stress, whereas the Q1 and Q3 quadrants have the opposite effect and work to decrease the time-averaged Reynolds stress. The best GDR case shows the greatest difference with respect to the canonical case, with an increase in



**Fig. 18** Joint and marginal probability distributions at the second probe location ( $x = 145$ ,  $y^+ = 100$ ) for the streamwise and wall-normal fluctuating velocities for selected cases. The contours are spaced at intervals of 0.1 in the range (0, 1)

both the frequency and overall contribution of the Q1 and Q3 events. This is accompanied by a decrease in the frequency of the Q2 and Q4 events. However, the overall contribution to the time-averaged Reynolds stress from Q4 is increased significantly compared to the canonical case. While such behaviour would typically be associated with an increase in skin-friction drag, the effect of the increase in contribution from the Q1 and Q3 quadrants works to counter this effect and ultimately leads to the reduced drag for this case.

Close to the wall ( $y^+ < 15$ ), Q4 is known to have the largest contribution to the Reynolds stress, whereas Q2 typically dominates further from the wall (Wallace 2016). Figure 17 confirms this behaviour close to the wall, where Q4 dominates for all cases. However, while Q2 becomes dominant for the canonical, best NES, and worst NES/GDR cases away from the wall, Q4 remains dominant for the best GDR case. To investigate this behaviour further, Fig. 18 shows the joint and marginal probability distributions for the streamwise and wall-normal fluctuating velocities at the second probe location ( $x = 145$ ,  $y^+ = 100$ ). In all cases, the principal component of the distribution is aligned diagonally along the Q2 and Q4 quadrants. Furthermore, all cases display a slight dilation of the distribution when compared with the canonical case. This is especially true for the best GDR case and is evidence of the increase in turbulent kinetic energy associated with the control.

For the best GDR case, the joint distribution is skewed more heavily towards Q4, which is directly related to the fact it remains the dominant contributor to the Reynolds stress at  $y^+ = 100$ .

## 6 Summary and Conclusions

To investigate the potential of a continuous streamwise-varying profile of wall-normal blowing as a control strategy for turbulent skin-friction drag reduction, this work presents the results from a Bayesian optimisation campaign aimed at maximising the NES of a blowing configuration parameterised by a cubic spline. In this endeavour, a high-order ILES solver (*Incompact3d*) is combined with a newly designed Bayesian optimisation framework to find optimal combinations of control parameters that balance the drag reduction with the power required to drive the actuation. Furthermore, to provide realistic estimates of NES, the input power for the control is estimated from real-world experimental data, rather than an idealised first principle-based approximation.

The results show that while significant GDR is relatively straightforward to achieve, NES is more difficult and in fact no improvement over the trivial (canonical) case was observed. After neglecting the trivial (non-blowing) cases, the next best NES was observed to be  $-0.05\%$ . This configuration was characterised by moderate-intensity blowing at the start and end points of the control region, with a reduction in blowing intensity in the middle of the control region (resemblant of an intermittent blowing strategy). The GDR for this case was found to be  $9.56\%$ , indicating the potential for significant NES gains with more efficient blowing devices. The overall best GDR case was characterised by maximum intensity blowing across the entire control region, which is resemblant of uniform blowing for this configuration. While a GDR of  $16.72\%$  was reported for this case, the NES was found to be relatively poor, at  $-2.09\%$ , which is a result of the input power required to sustain maximum intensity blowing across the entire control region. Although the present work was not able to demonstrate positive NES for this blowing configuration, the results show that even a small improvement in the efficiency of the present blowing device would lead to meaningful NES.

Compared to the work of Mahfoze et al. (2019), which demonstrated a NES of  $1.2\%$  using a uniform blowing profile, it seems that the present control strategy is not as effective at achieving NES. Having said that, the present study is limited by the computational budget available for the optimisation campaign. Furthermore, the input space was restricted by focussing on one specific spline type with a fixed number of knots, each of which with a fixed location. Moreover, this work focusses on just one specific flow condition. Therefore, this study is not sufficient to discount streamwise-varying blowing as an effective strategy for achieving NES in a turbulent boundary layer. This is especially true since such an approach is well placed to exploit the well-documented ‘inertia’ in the skin-friction recovery that occurs in the vicinity of the control region.

There are a number of possible avenues for future work. Firstly, resource limitations prevented further iterations of the Bayesian optimisation algorithm. It is possible that positive NES could still be achieved with this configuration with a longer optimisation campaign. Secondly, it would be interesting to repeat this study with a more efficient blowing device. While the results demonstrate that even a small improvement in efficiency would lead to positive NES, introducing a different blowing device would also change the shape of the objective function and would therefore likely change the optimal configuration. Thirdly,

to translate this technology into real-world applications it is necessary to reproduce this work experimentally. In the experimental setting it will also be possible to perform a larger number of optimisation iterations compared with the numerical simulations. Finally, more optimisation campaigns should be conducted to investigate the effects of Reynolds number, non-zero pressure gradients, and different blowing strategies (e.g. travelling waves of wall-normal blowing), to name a few.

**Author Contributions** All authors conceived, discussed, and planned the numerical simulations, as well as analysed and interpreted the results. JO and MD carried out the numerical simulations. SL provided the high-order flow solver, whereas MD, KW, and RW developed the Bayesian optimisation framework. RW, SL, KW, and AW obtained the funding for the present work. JO lead the writing of the manuscript, with MD and SL also contributing. All authors reviewed the manuscript and provided feedback before submission.

**Funding** The authors would like to thank the Engineering and Physical Sciences Research Council (EPSRC) for funding this work under Grant Numbers EP/T021144/1 and EP/T020946/1, and the EPSRC Centre for Doctoral Training (CDT) in Cloud Computing for Big Data under Grant Number EP/L015358/1. The authors would also like to thank the UK Turbulence Consortium (EP/R029326/1) for the computational time made available on the ARCHER2 UK National Supercomputing Service.

**Data Availability** Data and scripts to reproduce the figures are available at <https://doi.org/10.5281/zenodo.7687850>, with additional data available on request.

**Code Availability** A fork of `Incompact3d` with the implementation for the wall-normal blowing boundary condition has been made available at <https://github.com/joconnor22/Incompact3d>.

## Declarations

**Conflict of interest** The authors declare that they have no competing interests.

**Ethical Approval** Not applicable.

**Informed Consent** Not applicable.

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

## References

- Abbas, A., Bugeda, G., Ferrer, E., et al.: Drag reduction via turbulent boundary layer flow control. *Sci. China Technol. Sci.* **60**(9), 1281–1290 (2017)
- Atzori, M., Vinuesa, R., Fahland, G., et al.: Aerodynamic effects of uniform blowing and suction on a NACA4412 airfoil. *Flow Turbul. Combust.* **105**(3), 735–759 (2020)
- Bartholomew, P., Deskos, G., Frantz, R., et al.: Xcompact3D: An open-source framework for solving turbulence problems on a Cartesian mesh. *SoftwareX* **12**(100), 550 (2020)
- Brochu, E., Cora, V., de Freitas, N.: A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning. arXiv Preprint [arXiv: 1012.2599](https://arxiv.org/abs/1012.2599) [cs.LG] (2010)
- Choi, H., Moin, P., Kim, J.: Direct numerical simulation of turbulent flow over riblets. *J. Fluid Mech.* **255**, 503–539 (1993)

- Dairay, T., Lamballais, E., Laizet, S., et al.: Numerical dissipation vs. subgrid-scale modelling for large eddy simulation. *J. Comput. Phys.* **337**, 252–274 (2017)
- Deskos, G., Laizet, S., Piggott, M.: Turbulence-resolving simulations of wind turbine wakes. *Renew. Energy* **134**, 989–1002 (2019)
- Diaz-Daniel, C., Laizet, S., Vassilicos, J.: Wall shear stress fluctuations: mixed scaling and their effects on velocity fluctuations in a turbulent boundary layer. *Phys. Fluids* **29**(5), 055,102 (2017)
- Diessner, M., O'Connor, J., Wynn, A., et al.: Investigating Bayesian optimization for expensive-to-evaluate black box functions: Application in fluid dynamics. *Front Appl. Math. Stat.* **8**(1076), 296 (2022)
- Frantz, R., Deskos, G., Laizet, S., et al.: High-fidelity simulations of gravity currents using a high-order finite-difference spectral vanishing viscosity approach. *Comput. Fluids* **221**(104), 902 (2021)
- Freitas, D., Lopes, L., Morgado-Dias, F.: Particle Swarm Optimisation: a historical review up to the current developments. *Entropy* **22**(3), 362 (2020)
- Fritsch, F., Carlson, R.: Monotone piecewise cubic interpolation. *SIAM J. Numer. Anal.* **17**(2), 238–246 (1980)
- Fu, Y., Yuan, C., Bai, X.: Marine drag reduction of shark skin inspired riblet surfaces. *Biosurf. Biotribol.* **3**(1), 11–24 (2017)
- Fukagata, K., Iwamoto, K., Kasagi, N.: Contribution of Reynolds stress distribution to the skin friction in wall-bounded flows. *Phys. Fluids* **14**(11), L73 (2002)
- Gad-El-Hak, M.: Interactive control of turbulent boundary layers - A futuristic overview. *AIAA J.* **32**(9), 1753–1765 (1994)
- García-Mayoral, R., Jiménez, J.: Drag reduction by riblets. *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* **369**(1940), 1412–1427 (2011)
- Gramacy, R.: Surrogates: Gaussian process modeling, design, and optimization for the applied sciences. Chapman and Hall (2020)
- Jones, D.: A taxonomy of global optimization methods based on response surfaces. *J. Glob. Optim.* **21**, 345–383 (2001)
- Jones, D., Schonlau, M., Welch, W.: Efficient global optimization of expensive black-box functions. *J. Glob. Optim.* **13**, 455–492 (1998)
- Kametani, T., Fukagata, K., Örlü, R., et al.: Drag reduction in spatially developing turbulent boundary layers by spatially intermittent blowing at constant mass-flux. *J. Turbul.* **17**(10), 913–929 (2016)
- Kametani, Y., Fukagata, K.: Direct numerical simulation of spatially developing turbulent boundary layers with uniform blowing or suction. *J. Fluid Mech.* **681**, 154–172 (2011)
- Kametani, Y., Fukagata, K., Örlü, R., et al.: Effect of uniform blowing/suction in a turbulent boundary layer at moderate Reynolds number. *Int. J. Heat Fluid Flow* **55**, 132–142 (2015)
- Katoch, S., Chauhan, S., Kumar, V.: A review on genetic algorithm: past, present, and future. *Multimed. Tools Appl.* **80**(5), 8091–8126 (2021)
- Kim, J., Hwang, J., Yoon, M., et al.: Influence of a large-eddy breakup device on the frictional drag in a turbulent boundary layer. *Phys. Fluids* **29**(6), 065,103 (2017)
- Kornilov, V., Boiko, A.: Efficiency of air microblowing through microperforated wall for flat plate drag reduction. *AIAA J.* **50**(3), 724–732 (2012)
- Kornilov, V., Boiko, A.: Flat-plate drag reduction with streamwise noncontinuous microblowing. *AIAA J.* **52**(1), 93–103 (2014)
- Kravchenko, A., Moin, P.: On the effect of numerical errors in large Eddy simulations of turbulent flows. *J. Comput. Phys.* **131**(2), 310–322 (1997)
- Laizet, S., Lamballais, E.: High-order compact schemes for incompressible flows: a simple and efficient method with quasi-spectral accuracy. *J. Comput. Phys.* **228**(16), 5989–6015 (2009)
- Laizet, S., Li, N.: Incompact3d: A powerful tool to tackle turbulence problems with up to  $O(10^5)$  computational cores. *Int. J. Numer. Methods Fluids* **67**(11), 1735–1757 (2011)
- Larroque, A., Fosas de Pando, M., Lafuente, L.: Cylinder drag minimization through wall actuation: a Bayesian Optimization approach. *Comput. Fluids* **240**(105), 370 (2022)
- Lele, S.: Compact finite difference schemes with spectral-like resolution. *J. Comput. Phys.* **103**(1), 16–42 (1992)
- Mahfoze, O., Laizet, S.: Skin-friction drag reduction in a channel flow with streamwise-aligned plasma actuators. *Int. J. Heat Fluid Flow* **66**, 83–94 (2017)
- Mahfoze, O., Laizet, S.: Non-explicit large eddy simulations of turbulent channel flows from  $Re_{\tau}=up$  to  $Re_{\tau}=5,200$ . *Comput. Fluids* **228**, 105019 (2021)
- Mahfoze, O., Moody, A., Wynn, A., et al.: 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* **4**(9), 094601 (2019)

- Marusic, I., Chandran, D., Rouhi, A., et al.: An energy-efficient pathway to turbulent drag reduction. *Nat. Commun.* **12**, 5805 (2021)
- Mockus, J.: Application of Bayesian approach to numerical methods of global and stochastic optimization. *J. Glob. Optim.* **4**, 347–365 (1994)
- Morita, Y., Rezaeiravesh, S., Tabatabaei, N., et al.: Applying Bayesian optimization with Gaussian process regression to computational fluid dynamics problems. *J. Comput. Phys.* **449**(110), 788 (2022)
- Quadrio, M., Ricco, P., Viotti, C.: Streamwise-travelling waves of spanwise wall velocity for turbulent drag reduction. *J. Fluid. Mech.* **627**, 161–178 (2009)
- Rasmussen, C., Williams, C.: *Gaussian processes for machine learning*. The MIT Press (2005)
- Ruder, S.: An overview of gradient descent optimization algorithms. *arXiv Preprint arXiv:1609.04747* [cs.LG] (2016)
- Schlatter, P., Örlü, R.: Assessment of direct numerical simulation data of turbulent boundary layers. *J. Fluid Mech.* **659**, 116–126 (2010)
- Shahriari, B., Swersky, K., Wang, Z., et al.: Taking the human out of the loop: a review of Bayesian optimization. *Proc. IEEE* **104**(1), 148–175 (2016)
- Snoek, J., Larochelle, H., Adams, R.: Practical Bayesian optimization of machine learning algorithms. *Adv. Neural Inf. Process. Syst.* (2012)
- Srinivas, N., Krause, A., Kakade, S., et al.: Gaussian process optimization in the bandit setting: no regret and experimental design. *arXiv Preprint arXiv:0912.3995* [cs.LG] (2010)
- Storn, R., Price, K.: Differential evolution—a simple and efficient Heuristic for global optimization over continuous spaces. *J. Glob. Optim.* **11**, 341–359 (1997)
- Stroh, A., Frohnapfel, B., Schlatter, P., et al.: A comparison of opposition control in turbulent boundary layer and turbulent channel flow. *Phys. Fluids* **27**(7), 075,101 (2015)
- Stroh, A., Hasegawa, Y., Schlatter, P., et al.: Global effect of local skin friction drag reduction in spatially developing turbulent boundary layer. *J. Fluid Mech.* **805**, 303–321 (2016)
- Wallace, J.: Quadrant analysis in turbulence research: history and evolution. *Annu. Rev. Fluid Mech.* **48**, 131–158 (2016)
- Wang, J., Choi, K., Feng, L., et al.: Recent developments in DBD plasma flow control. *Prog. Aerosp. Sci.* **62**, 52–78 (2013)
- White, C., Mungal, M.: Mechanics and prediction of turbulent drag reduction with polymer additives. *Annu. Rev. Fluid Mech.* **40**, 235–256 (2008)
- Wilson, J., Hutter, F., Deisenroth, M.: Maximizing acquisition functions for Bayesian optimization. *Adv. Neural Inf. Process. Syst.* 31 (2018)

## Authors and Affiliations

Joseph O'Connor<sup>1</sup> · Mike Diessner<sup>2</sup> · Kevin Wilson<sup>3</sup> · Richard D. Whalley<sup>4</sup> · Andrew Wynn<sup>1</sup> · Sylvain Laizet<sup>1</sup>

Mike Diessner  
m.diessner2@newcastle.ac.uk

Kevin Wilson  
kevin.wilson@newcastle.ac.uk

Richard D. Whalley  
richard.whalley@newcastle.ac.uk

Andrew Wynn  
a.wynn@imperial.ac.uk

Sylvain Laizet  
s.laizet@imperial.ac.uk

<sup>1</sup> Department of Aeronautics, Imperial College London, London, UK

<sup>2</sup> School of Computing, Newcastle University, Newcastle upon Tyne, UK

<sup>3</sup> School of Mathematics, Statistics and Physics, Newcastle University, Newcastle upon Tyne, UK

<sup>4</sup> School of Engineering, Newcastle University, Newcastle upon Tyne, UK