

# A Bayesian optimisation framework for drag reduction and net energy saving in a turbulent boundary layer using wall-normal blowing

Omar A. Mahfoze<sup>a</sup>, Joseph O'Connor<sup>a†</sup>, Andrew Wynn<sup>a</sup>, Richard D. Whalley<sup>b</sup> and Sylvain Laizet<sup>a</sup>

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

<sup>b</sup>School of Engineering, Newcastle University, Newcastle, UK

†joseph.oconnor@imperial.ac.uk

## Abstract

Skin friction drag is a major engineering concern, with implications across many industries. Active flow control strategies aimed at reducing skin friction have the potential to offer substantial improvements in aerodynamic efficiency, reduce operating costs, and help meet emission targets. However, they are difficult to design and optimise. Furthermore, any performance benefits must be weighed against the power required to drive the actuation. Bayesian optimisation is one such technique that is well 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 wall-normal blowing as a drag reduction strategy for turbulent boundary layers by combining a high-order flow solver (Xcompact3d) with a Bayesian optimisation framework. Two separate optimisation campaigns are performed. The first extends the input space of a previous study to allow greater flexibility in the blowing strategy, as well as including wall-normal suction. The second examines the effect of increasing the Reynolds number. In both campaigns the key metric is the net energy saving, which accounts for both the drag reduction and the input power required to drive the actuation. The results from the first campaign are mixed, with the optimal solution achieving only a small net energy saving. Furthermore, suction is shown to be of no benefit in this particular setup. The results from the second campaign are more promising, indicating that the efficiency of the wall-normal blowing strategy increases with Reynolds number. In both campaigns the best configuration is shown to be an intermittent blowing strategy. The results demonstrate that wall-normal blowing can be an effective strategy for skin friction drag reduction, under certain conditions, and that Bayesian optimisation is an effective approach for optimising such control strategies.

## 1. Introduction

Skin friction drag reduction is a topic of great interest due to its importance in many engineering applications. In the aviation industry, for example, skin friction drag is estimated to account for approximately half of the total drag.<sup>1</sup> Similar figures are also observed across the transport sector, such as in the maritime industry.<sup>4</sup> Clearly, strategies aimed at reducing skin friction drag 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 practical and affordable method for skin friction drag reduction is yet to be implemented in real-world applications.

There are a variety of flow control strategies that have been investigated over the years for reducing skin friction drag. This includes passive strategies, such as polymer additives,<sup>20</sup> riblets,<sup>2</sup> and large eddy breakup devices,<sup>7</sup> to name a few. However, passive approaches often suffer from parasitic drag effects and may lose effectiveness over the full range of operating conditions when in service. On the other hand, active control strategies, such as surface jets or wall deformation, are among the most promising approaches to tackle skin friction drag due to their increased performance, robustness, and flexibility. However, designing effective active control policies is a difficult task, owing to the nonlinear, high-dimensional, and chaotic dynamics of fluids. Furthermore, active control strategies require a power supply. Therefore, any performance benefits must be weighed against the power required to drive the actuation.

Of the available active control strategies, mass flow injection (e.g. surface jets) is one of the most extensively studied, owing to its practicality and performance. Multiple studies, both experimental and numerical, have demonstrated impressive reductions in local skin friction drag (approximately 55%–70%) for a turbulent boundary layer with low

intensity wall-normal blowing (up to 1 % of the free-stream velocity).<sup>8,19</sup> One of the strengths of this approach is its flexibility, with a vast range of choices in terms of design, from spatial 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.

The challenges associated with designing wall-normal blowing strategies, and active flow control policies in general, for turbulent boundary layers have led to a trend towards adopting optimisation techniques. Mathematical optimisation is a mature field and there are a vast range of different techniques available. However, 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 and where the objective function is non-convex and expensive to evaluate, such as with large eddy simulations (LES) or direct numerical simulations (DNS). Bayesian optimisation (BO) is one such approach that has recently gained significant popularity for hyperparameter tuning of machine learning models.<sup>17,18</sup> Thanks to this success, BO has also started to find applications in engineering optimisation, such as turbulent drag reduction. Mahfoze et al.<sup>14</sup> applied BO to optimise the control parameters of a wall-normal blowing configuration over a turbulent boundary layer. Rather than targetting drag reduction, they optimised for net energy saving by taking into account both the power saved due to drag reduction and the input power required to drive the actuation, which was estimated via experimental data. Their approach was able to find combinations of control parameters leading to between 1.2 %–5 % net energy saving, depending on the experimental data they used to estimate power consumption. More recently, Morita et al.<sup>15</sup> applied BO to a wide variety of fluid problems, showing that it was capable of finding globally optimal solutions with relatively few evaluations of the objective function (simulations).

In light of the above, the purpose of this work is to extend the work of Mahfoze et al.<sup>14</sup> to investigate the potential of wall-normal blowing as a drag reduction strategy for turbulent boundary layers over a greater range of blowing configurations and flow conditions. In this endeavour, the high-order flow solver Xcompact3d is combined with a BO framework to optimise the control parameters for different blowing configurations. Two separate optimisation campaigns are carried out. The first extends the input space of Mahfoze et al.<sup>14</sup> by including a second blowing amplitude to be optimised, as well as allowing both blowing and suction for both of the amplitudes. The second campaign extends the work of Mahfoze et al.<sup>14</sup> to higher Reynolds numbers to investigate the performance of the control strategy in flow conditions more relevant to real-world applications. The key performance metric for both campaigns is the net energy saving, which takes into account both the power saving due to drag reduction and the power required to drive the actuation (which is estimated using experimental data). In the following sections the flow solver and BO algorithm are presented, followed by a description of the two optimisation campaigns. The results are then presented for each campaign before a discussion on the conclusions and suggestions for future work.

## 2. Methods

### 2.1 Numerical Methods

The governing equations are the incompressible Navier-Stokes equations:

$$\nabla \cdot \mathbf{u} = 0 \quad (1)$$

$$\frac{\partial \mathbf{u}}{\partial t} + \frac{1}{2} [\nabla(\mathbf{u} \otimes \mathbf{u}) + (\mathbf{u} \cdot \nabla)\mathbf{u}] = -\frac{1}{\rho} \nabla p + \nu \nabla^2 \mathbf{u} \quad (2)$$

where  $\mathbf{u}$  is the velocity vector,  $t$  is time,  $\rho$  is mass density,  $p$  is pressure, and  $\nu$  is the kinematic viscosity. Note that Equation 2 is given in skew-symmetric form to reduce aliasing errors.<sup>9</sup>

Equations 1 and 2 are solved on a Cartesian mesh with a sixth-order compact finite-difference scheme for spatial discretisation and a third-order Adams-Bashforth scheme for time advancement. This explicit scheme is additionally combined with an implicit Crank-Nicolson approach for the diffusive term in the wall-normal direction. This allows a sufficiently fine mesh resolution in the near-wall region to properly resolve the viscous sublayer, without leading to intractable timestep constraints. The Poisson equation, which enforces incompressibility, is fully solved in spectral space via the use of relevant 3D fast Fourier transforms (FFTs). Using the concept of the modified wavenumber,<sup>12</sup> the divergence-free condition is ensured up to machine accuracy. The mesh for the pressure field is staggered with respect to the velocity field to avoid the spurious pressure oscillations typically observed in a fully collocated approach.<sup>10</sup> The simple mesh structure allows easy implementation within a 2D domain decomposition strategy, based on the message passing interface (MPI).<sup>11</sup> The computational domain is split into several sub-domains (pencils), which are each assigned an MPI process. The derivatives and interpolations in the  $x$ ,  $y$ , and  $z$  directions are performed within the  $x$ ,  $y$ , and  $z$  pencils, respectively. The 3D FFTs, required by the Poisson solver, are broken down into a series of 1D FFTs computed in one direction at a time. Global transpositions to switch from one pencil to another are performed with the MPI

command `MPI_ALLTOALL(V)`. This approach has been shown to scale well on hundreds of thousands of cores for simulations with several billion mesh nodes.<sup>11</sup> Further details and validations of Xcompact3d can be found in Laizet & Lamballais.<sup>10</sup>

This work adopts an implicit LES scheme, via an original approach, whereby the coefficients of the stencils for the second-order derivatives are modified to add numerical dissipation at the small scales using the numerical error of the finite-difference scheme, while still retaining sixth-order accuracy. This numerical dissipation mimics a spectral vanishing viscosity and can be used to model the dissipation of the non-resolved scales at no additional cost (see Dairay et al.<sup>3</sup> and Mahfoze & Laizet<sup>13</sup> for more details).

## 2.2 Bayesian Optimisation

This work adopts a BO framework to optimise different configurations of low-amplitude wall-normal blowing (e.g. spatial arrangement, amplitude, etc.) for net energy saving. BO is a derivative-free, non-intrusive approach that is well suited to non-convex objective functions that are expensive to evaluate (e.g. LES) and have a moderate number of input dimensions.<sup>5,18</sup> The goal of BO is to minimise/maximise a user-defined objective function over a given set of input parameter ranges through repeated evaluation of the objective function at different locations within the parameter space.

There are two main stages to BO. Firstly, a surrogate model of the objective function is generated based on previous observations (evaluations) of the objective function. This surrogate model is typically chosen to be a Gaussian process (GP), which provides a probabilistic model of the true objective function based on the available training data. The GP model provides a relatively inexpensive way to estimate the value of the objective function in locations of the parameter space that have not yet been sampled (via the mean of the GP model) while also providing a measure of the uncertainty associated with those estimates (via the covariance of the GP model). The second stage in BO requires selecting new locations in the parameter space to sample the objective function. This is done by optimising a user-defined acquisition function, which relies on efficient estimation of the objective function via the surrogate model. There are a variety of acquisition functions available in the literature, most of which typically rely on a balance between exploitation (favouring areas where the mean of the GP model is optimal) and exploration (favouring areas where the uncertainty/covariance of the GP model is maximised). This work adopts the popular Expected Improvement acquisition function, given for the minimisation problem by:

$$EI(\mathbf{x}) = \mathbb{E} [\max (f(\mathbf{x}^+) - f(\mathbf{x}), 0)] \quad (3)$$

which has a closed form solution when combined with GPs:

$$EI(\mathbf{x}) = \begin{cases} (f(\mathbf{x}^+) - \mu(\mathbf{x}) + \xi) \Phi(Z) + \sigma(\mathbf{x})\phi(Z) & \sigma(\mathbf{x}) > 0 \\ 0 & \sigma(\mathbf{x}) = 0 \end{cases} \quad (4)$$

where  $\mathbf{x}^+$  is the location of the best observation so far,  $f(\mathbf{x}^+)$  is the value of the objective function at that location,  $\mu(\mathbf{x})$  and  $\sigma(\mathbf{x})$  are the mean and standard deviation of the GP model at  $\mathbf{x}$ ,  $\xi$  is a tunable parameter that can be used to control the balance between exploration and exploitation, and  $\Phi(Z)$  and  $\phi(Z)$  are the cumulative and probability density functions of the standard normal distribution for  $Z$ , given by:

$$Z = \begin{cases} (f(\mathbf{x}^+) - \mu(\mathbf{x}) + \xi) / \sigma(\mathbf{x}) & \sigma(\mathbf{x}) > 0 \\ 0 & \sigma(\mathbf{x}) = 0 \end{cases} \quad (5)$$

Figure 1 illustrates the process of applying BO to a simple 1D test function. Initially, an observation is drawn by randomly sampling the input space and evaluating the objective function at this location. This provides the first training point with which to build the GP model. Figure 1a shows that at the location where the sample was drawn the mean of the GP model matches the underlying objective function and the uncertainty (covariance) is zero. However, away from the sample location the mean of the GP model departs significantly from the true objective function and the covariance increases to indicate the level of uncertainty. The location of the next sample is selected by optimising the acquisition function, which balances both exploration and exploitation by considering both the mean and the covariance of the GP model. Once the objective function is evaluated at the next sample location the result is appended to the training data and the GP model is recomputed. As can be seen from Figure 1b, the updated GP model now matches the objective function at the new sample location and again the covariance becomes zero. These steps are repeated until the user-defined stopping criteria are satisfied. These criteria can be based on performance metrics or on the available resources, for example. After each iteration the GP model becomes more and more accurate and, as a result, it can be used as a much cheaper surrogate model for the true objective function.




Figure 1: Illustration of applying Bayesian optimisation to a 1D test function. The black line indicates the true objective function, the black dots indicate evaluations of the objective function, the red line indicates the mean of the Gaussian process posterior and the red shaded region indicates the covariance (uncertainty).

### 3. Simulation Setup

Two separate configurations are investigated in this work, each representing a canonical turbulent boundary layer flow with low-amplitude wall-normal blowing. Figure 2 shows a schematic of the simulation setup. The streamwise, wall-normal, and spanwise directions are denoted by  $x$ ,  $y$ , and  $z$ , respectively. A laminar Blasius solution is prescribed at the inlet with a boundary layer height of  $\delta_0$ . The remaining boundary conditions are a convective condition at the outlet, a homogenous Neumann condition in the far-field, and periodic conditions in the spanwise direction. To accelerate the transition to turbulence, a random volume forcing approach,<sup>16</sup> located at  $x = 3.5\delta_0$ , is used to trip the boundary layer. Wall-normal blowing is implemented as a velocity boundary condition.

#### 3.1 Configuration I

The first configuration is an extension of the work conducted by Mahfoze et al.<sup>14</sup> The domain dimensions are  $L_x \times L_y \times L_z = 750\delta_0 \times 80\delta_0 \times 15\delta_0$ . The Reynolds number at the inlet is  $Re_\theta = 170$ , based on the momentum thickness. The mesh resolution is  $n_x \times n_y \times n_z = 1536 \times 128 \times 32$ , resulting in a resolution of  $\Delta x^+ = 34$ ,  $0.60 \leq \Delta y^+ \leq 780$ , and  $\Delta z^+ = 16$ , where the wall units are with respect to the friction velocity before the start of the blowing region at  $Re_\theta = 365$ . The flow is initialised to a laminar Blasius solution through the entire domain and allowed to develop until  $t^+ = 4250$  (inner scaling) before the statistics begin recording. The spanwise-averaged skin friction profile along the domain of interest is monitored at each time step, and the mean squared difference between successive time steps is used as a stopping criteria for each simulation. The blowing region starts at  $x = 68\delta_0$  and extends to  $x = 145\delta_0$ , corresponding to  $470 \leq Re_\theta \leq 700$ , and covers the entire spanwise extent of the domain. The blowing configuration



Figure 2: Schematic of the simulation setup. Reproduced from Mahfoze et al.<sup>14</sup>

Figure 3: Blowing profiles for both configurations.

Table 1: Input ranges for each configuration. The blowing amplitudes are expressed as a percentage of the free-stream velocity.

(a) Input ranges for Configuration I.

$A_1$	$A_2$	$\alpha$	$N_B$
$[-1, 1]$	$[-1, 1]$	$[0, 1]$	$[1, 10]$

(b) Input ranges for Configuration II.

$A$	$\alpha$	$N_B$
$[0, 1]$	$[0, 1]$	$[1, 10]$

consists of alternating strips of wall-normal velocity, as denoted in Figure 3a. The main difference between the present work and that of Mahfoze et al.<sup>14</sup> is that the present work includes two separate amplitudes for the alternating strips, whereas Mahfoze et al.<sup>14</sup> only included one (with the second amplitude always zero). Furthermore, in the present work suction is also permitted. The optimisation is thus parameterised by four inputs: the two blowing/suction amplitudes ( $A_1$  and  $A_2$ ), the surface area ratio between the two alternating strips ( $\alpha = \lambda_1 N_B / L_B$ ), and the number of strips ( $N_B$ ). The ranges for each input are given in Table 1a.

### 3.2 Configuration II

The second configuration is also an extension of Mahfoze et al.<sup>14</sup> However, this time the blowing configuration is kept the same and instead the domain is extended to investigate the effect of the Reynolds number. Everything is therefore the same as the previous configuration, except now the domain dimensions are  $L_x \times L_y \times L_z = 2000\delta_0 \times 150\delta_0 \times 30\delta_0$ . The mesh resolution is  $n_x \times n_y \times n_z = 4097 \times 512 \times 128$ , resulting in a resolution of  $\Delta x^+ = 33$ ,  $0.57 \leq \Delta y^+ \leq 500$ , and  $\Delta z^+ = 16$  (at  $Re_\theta = 365$ ). With these dimensions the maximum Reynolds number reaches approximately  $Re_\theta \approx 4300$  at the end of the computational domain. The control region extends from  $Re_\theta = 500$ – $1300$ . The parameterisation of the optimisation is the same as that given in Mahfoze et al.<sup>14</sup> and is illustrated in Figure 3b. There are three inputs this time: the blowing amplitude ( $A$ ), the surface area ratio between the blowing/non-blowing regions ( $\alpha = \lambda N_B / L_B$ ), and the number of strips ( $N_B$ ). The ranges for each input are given in Table 1b.

### 3.3 Power Estimation

The net energy saving ( $S$ ) generated by each blowing control strategy is assessed by taking into account the input power required to generate the wall-normal blowing, plus any power saving due to a reduction in skin friction drag. Owing to the long-lasting downstream effects of the blowing control, a global skin friction drag coefficient ( $\overline{C}_f$ ) is defined as

## A BAYESIAN OPTIMISATION FRAMEWORK FOR DRAG REDUCTION AND NET ENERGY SAVING

the streamwise-averaged skin friction coefficient from  $Re_\theta = 385$ – $1900$  for Configuration I and  $Re_\theta = 500$ – $4300$  for Configuration II. The reduction in the global skin friction drag coefficient is referred to as the global drag reduction (GDR) and the net energy saving is defined as the reduction in the gross power input, which is the sum of the mean viscous power to overcome the shear stress and the blowing power.

In this work, the wall-normal blowing is imposed as a velocity boundary condition. Therefore, the blowing power is estimated using experimental data from two different blowing systems. The first system ( $S_1$ ) is based on a pressurised chamber on one side of a micro-drilled plate.<sup>8</sup> A relationship between the pressure drop through the plate and the blowing amplitude was provided. However, the authors did not indicate the power required to compress the air, potentially leading to an underestimation of the total power requirements. The second system ( $S_2$ ) is currently being developed at Newcastle University (UK) and is based on miniature electromagnetic speakers to blow air across a micro-perforated plate. Here, the instantaneous total power input for the speakers as a function of the blowing velocity is used to compute the net energy saving. See Mahfoze et al.<sup>14</sup> for more details.

## 4. Results

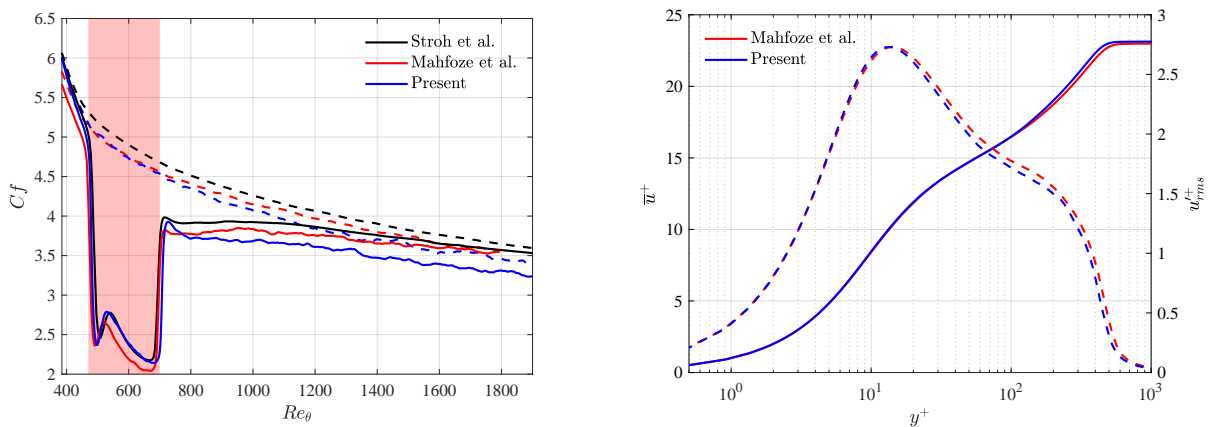
### 4.1 Validation

In this section the present implicit LES scheme is validated against benchmark DNS data from Stroh et al.<sup>19</sup> and Mahfoze et al.<sup>14</sup> Figure 4a shows the skin friction coefficient with respect to the momentum Reynolds number for both a canonical turbulent boundary layer and the same case with uniform blowing of amplitude  $A = 0.5\%$  of the free-stream. The agreement is good, with the reduction in skin friction over the blowing region properly captured. There is a slight drift away from the benchmark DNS data in the latter half of the domain, which is consistent across both the canonical and blowing cases. However, overall the agreement with the benchmark data is good. Figure 4b shows the mean and fluctuating streamwise velocity with respect to the wall-normal coordinate for the canonical case at  $Re_\theta = 1410$ . The agreement is excellent and demonstrates the validity of the present scheme.

### 4.2 Configuration I

As described in Section 3.1, the purpose of this configuration is to extend the work of Mahfoze et al.<sup>14</sup> by including a second controllable amplitude in the input space. Furthermore, the effect of suction is also investigated by extending the input ranges of both amplitudes to also allow for negative wall-normal velocity. The optimisation campaign was initialised with four seed points. These were manually selected based on preliminary work indicating high-performing combinations, as well providing sufficient coverage of the input space. The objective function for optimisation is the net energy saving ( $S_2$ ) of the second experimental system described in Section 3.3. A total of 24 separate simulations were conducted. An overview of the results are presented in Table 2.

Table 2 shows that most of the cases do not lead to any net energy saving (when considering only  $S_2$ ). In fact,



(a) Skin friction coefficient ( $\times 1000$ ) for blowing (solid) and canonical (dashed) cases. Shaded area indicates blowing region. (b) Mean (solid) and fluctuating (dashed) streamwise velocity for the canonical case at  $Re_\theta = 1410$ .

Figure 4: Validation of the present implicit LES scheme against DNS data of Stroh et al.<sup>19</sup> and Mahfoze et al.<sup>14</sup>

## A BAYESIAN OPTIMISATION FRAMEWORK FOR DRAG REDUCTION AND NET ENERGY SAVING

Table 2: Control parameters, maximum local drag reduction, global drag reduction and net energy saving for Configuration I. Drag reduction and net energy saving are indicated as positive. The red row indicates the case with highest drag reduction, the blue row indicates the case with highest net energy saving, and the magenta row indicates the case with both the lowest drag reduction and net energy saving.

	$A_1$	$A_2$	$\alpha$	$N_B$	MLDR	GDR	$S_1$	$S_2$		$A_1$	$A_2$	$\alpha$	$N_B$	MLDR	GDR	$S_1$	$S_2$
1	0.5	0	1	1	53.8	13.2	3.52	-0.59	13	0.46	1	0.52	10	73.1	18.4	-4.25	-1.33
2	-0.5	0	1	1	18.8	-17.5	-27.2	-31.3	14	0.29	0.75	0.32	10	65.1	15.5	-0.23	-0.73
3	0.5	0	0.5	10	40.3	7.2	2.36	0.3	15	1	0.47	0.68	10	77.4	20.8	-8.4	-2.11
4	0.5	-0.5	0.5	10	33.1	-0.23	-9.92	-14	16	0.63	0.76	0.68	10	66.6	16.9	-0.63	-0.92
5	0.44	-0.01	0.11	3	29	1.12	0.31	-0.4	17	-0.94	1	0	9	82.1	23.6	-15	-3.75
6	0.66	0.14	0.68	10	58.7	13	1.05	-0.53	18	0.17	0.52	0.32	10	50.5	10.9	3.46	-0.54
7	0.64	-0.07	0.61	10	52.1	9.84	-0.04	-1.54	19	0.75	0.48	0	10	52.1	12.7	3.62	-0.73
8	0.5	0.45	0.47	10	51.1	12.4	3.84	-0.67	20	0.24	1	0.9	1	70.5	8.05	2.22	-1.36
9	0.99	1	0	10	82	23.8	-14.6	-3.47	21	-0.43	1	0.71	2	76.4	-3.22	-19.4	-19.7
10	0.4	0.21	0.42	10	38.5	7.84	4.29	-0.76	22	-0.1	0.75	1	10	1.69	-3.04	-3.42	-6.42
11	0.06	1	0.32	10	74.1	17.8	-8.49	-1.51	23	-1	-1	0.4	7	36.24	-36.6	-75	-63.8
12	0.45	0.16	0.32	10	36.9	7.09	3.83	-0.61	24	1	0.16	0	10	21.5	4.82	3.77	-0.53

Case 3, one of the initial seed points, is the only case that exhibits positive net energy saving, with  $S_2 = 0.3\%$ . This case is characterised by 10 rows of intermittent blowing, alternating between equally spaced regions of 0.5% and 0% of the free-stream velocity. It is interesting to note that in terms of drag reduction this case is relatively unexceptional, with a maximum local drag reduction (MLDR) and GDR of 40.3% and 7.2%, respectively. Nevertheless, the intermittent spatial arrangement combined with the moderate blowing amplitude works to reduce the required power consumption, while still achieving appreciable drag reduction. Further examination of Table 2 shows that as the BO algorithm progresses it tends to favour a high number of blowing strips, similar to Case 3, and while it finds many combinations that exhibit similar levels of performance it is not able to find anything better than Case 3. It is also interesting to note that this optimal solution is different to the optimal solution found in Mahfoze et al.,<sup>14</sup> which was characterised by a short single strip of moderate-to-high intensity blowing. This can be explained by the fact both optimisation campaigns tend to favour different regions of the parameter space right from the start, which can be attributed to different initial seeding. It is likely that running both campaigns for longer would lead to both converging on the same globally optimal solution (to within minor tolerances associated with differences in the numerical methods).

Further examination of Table 2 shows that Case 9 exhibits the largest drag reduction, with an MLDR and GDR of 82% and 23.8%, respectively. With  $\alpha = 0$ , this combination of control parameters degenerates into uniform blowing with maximum blowing amplitude ( $A_2 = 1\%$ ). Based on previous work this type of strategy is expected to provide the greatest drag reduction. However, the power required to drive the high intensity blowing means that this strategy is not performant when considering the net energy saving.

An additional goal of this section was to assess the potential benefits of including suction or combining blowing with suction for net energy saving. From Table 2, the cases where either or both amplitudes are negative (e.g. Case 23) show that suction does not perform well, both in terms of drag reduction and net energy saving. This is also expected from the literature, where suction is known to increase skin friction drag in zero-pressure gradient boundary layers.<sup>6</sup>

Figure 5 shows the spanwise-averaged profiles of the skin friction coefficient with respect to the momentum Reynolds number for the cases with maximum drag reduction (Case 9), maximum net energy saving (Case 3), and the case with both the minimum drag reduction and minimum net energy saving (Case 23). The canonical case is also shown for reference. Focussing on Case 3 first, the intermittent blowing strategy leads to significant drag reduction over each individual strip, with an abrupt rise in skin friction between each blowing region. However, it can be seen that even with this rise the skin friction does not return to the canonical value between the blowing regions. Furthermore, there is a slight drag reduction that persists somewhat beyond the blowing region. For Case 9, the high intensity uniform blowing leads to a drastic and sustained drag reduction over the blowing region. Furthermore, in contrast to Case 3, its effect persists much further downstream of the blowing region and to a much greater extent. Yet, even with these significant drag reductions, the power required to drive this strategy is too much to achieve net energy saving. Finally, Case 23 is the worst performing case, both in terms of drag reduction and net energy saving. The reason for this is evident in Figure 5. The high intensity suction leads to a drastic increase in skin friction, while also requiring a significant amount of power to drive the actuation. It can be concluded that, at least for the present setup, suction is not beneficial for drag reduction or net energy saving.

## A BAYESIAN OPTIMISATION FRAMEWORK FOR DRAG REDUCTION AND NET ENERGY SAVING

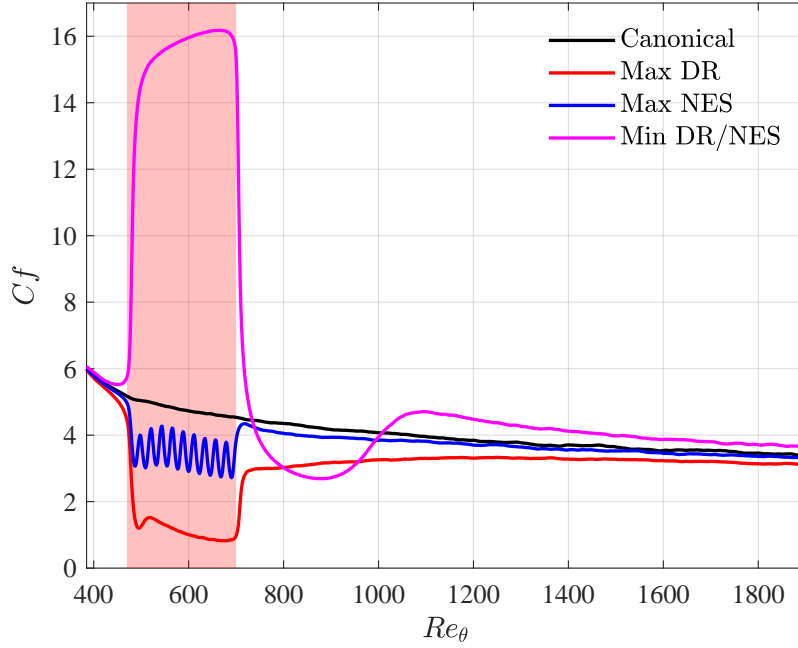


Figure 5: Profiles of the spanwise-averaged skin friction coefficient ( $\times 1000$ ) with respect to the momentum Reynolds number for Case 9 (max drag reduction), Case 3 (max net energy saving) and Case 23 (min drag reduction/net energy saving) of Configuration I. For reference, the canonical case is also shown. Red shaded area indicates the blowing region.

### 4.3 Configuration II

As illustrated in Mahfoze et al.,<sup>14</sup> net energy saving, along with substantial drag reduction, is possible for a low Reynolds number turbulent boundary layer. This section investigates the effect of increasing the Reynolds number on the achievable net energy saving. The streamwise length of the domain and the control region are approximately 2.5 and 4 times that of Mahfoze et al.<sup>14</sup> The objective function is the net energy saving ( $S_2$ ) based on the second system described in Section 3.3. A total of 24 simulations are conducted to get a satisfactory convergence. An overview of the results are presented in Table 3.

Table 3 shows that, in contrast to Configuration I, most of the cases for this configuration achieve net energy saving. Furthermore, the BO algorithm tends to favour intermittent blowing, similar to Configuration I. A maximum net energy saving of  $S_2 = 3.24\%$  is seen for Case 20. The MLDR and GDR for this case are 82.1% and 24.5%, respectively.

Table 3: Control parameters, maximum local drag reduction, global drag reduction and net energy saving for Configuration II. Drag reduction and net energy saving are indicated as positive. The red row indicates the case with highest drag reduction whereas the blue row indicates the case with highest net energy saving.

	$A$	$\alpha$	$N_B$	MLDR	GDR	$S_1$	$S_2$		$A$	$\alpha$	$N_B$	MLDR	GDR	$S_1$	$S_2$
1	0.93	0.23	5	78.4	8.8	-3.2	1.1	13	0.7	0.75	2	79.7	20.6	-1.4	2.41
2	0.29	1	1	46.7	12.5	7.4	0.97	14	0.75	0.78	8	82.5	23.3	-2.7	3.21
3	0.5	0.78	10	66.6	15.8	4.3	1.89	15	0	0	10	0	0	0	0
4	1	1	1	94.5	34.7	-25	-1.76	16	0.725	0.4	2	74.8	12.1	-0.5	2.06
5	0.8	1	1	88.6	29.4	-8.8	1.71	17	0.725	1	10	85.1	27.6	-3.8	2.51
6	1	0.78	1	93.6	26.8	-19.5	-1.47	18	0.675	0.88	10	79.7	23.5	-0.3	3.09
7	1	0.5	1	90.4	18.6	-11.3	0.34	19	0.725	0.68	10	79.2	19.9	-1.3	3.01
8	0.75	0.98	1	85.6	27.6	-5.2	2.32	20	0.725	0.9	10	82.1	24.5	-2.2	3.24
9	0.5	0.33	10	56.8	6.8	2	1.02	21	0.7	0.83	8	81.0	22.6	-1.6	2.6
10	0.8	0.68	3	83.3	21.1	-4.7	2.34	22	0.7	0.85	8	81.3	22.9	-2.0	2.31
11	0.575	0.8	9	72.2	18.8	3.0	2.65	23	0.725	0.8	9	81.6	22.8	-2.3	2.77
12	0.675	1	8	82.5	26.4	-0.8	3.04	24	0.725	0.83	8	82.0	23.2	-2.7	2.57



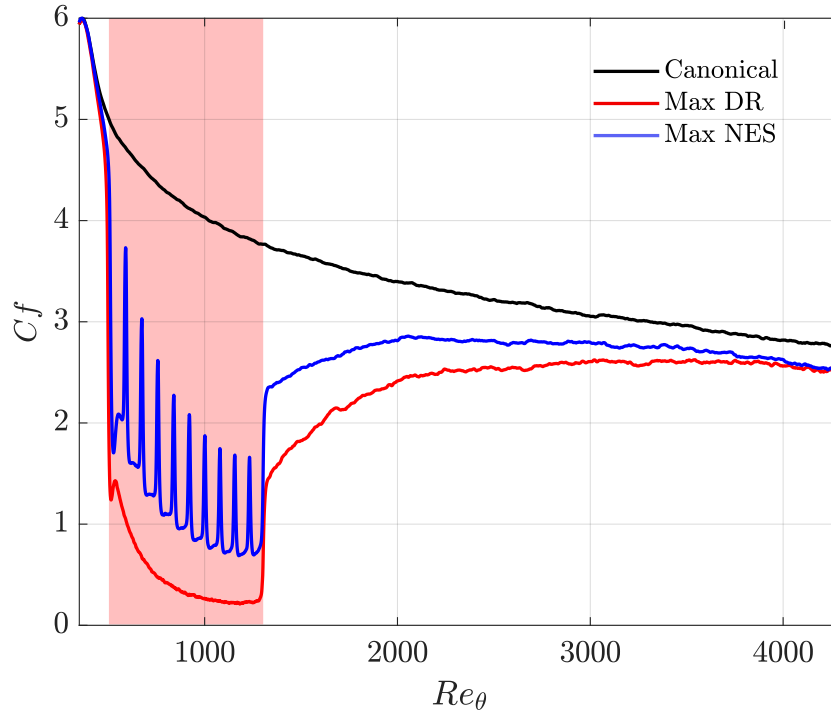


Figure 6: Profiles of the spanwise-averaged skin friction coefficient ( $\times 1000$ ) with respect to the momentum Reynolds number for Case 4 (max drag reduction) and Case 20 (max net energy saving) of Configuration II. For reference, the canonical case is also shown. Red shaded area indicates the blowing region.

It is important to highlight that the maximum net energy saving reported here is significantly more than the value of  $S_2 = 1.2\%$  reported in Mahfoze et al.,<sup>14</sup> indicating that increasing the Reynolds number leads to an increase in the effectiveness of the control strategy. Further examination of Table 3 shows that a uniform blowing strategy (Case 4) achieves the maximum MLDR and GDR (94.5% and 34.7%, respectively). Although the power required to drive such a control strategy is more than the power saved due to drag reduction, the difference is relatively small (1.76%). It is also interesting to note that Case 2, which is similar to the optimum case found in Mahfoze et al.,<sup>14</sup> shows an improvement in the first net energy saving metric ( $S_1$ ) from 5% to 7.4%.

These results suggest that the proposed blowing strategy becomes more efficient when the Reynolds number is increased, which is a promising sign for implementing such a strategy in real-world applications. Furthermore, increasing the streamwise length of the control region shows that intermittent blowing has the potential to enhance net energy saving. The increased performance with Reynolds number shown here can be attributed to the decrease in the friction velocity with Reynolds number. As a result, the wall-normal velocity scaled by the local friction velocity actually increases with Reynolds number, leading to increased drag reduction.

Figure 6 shows the spanwise-averaged profiles of the skin friction coefficient with respect to the momentum Reynolds number for the cases with maximum drag reduction (Case 4) and net energy saving (Case 20). The canonical case is also shown for reference. Focussing on Case 20 first, similar to Case 3 from Configuration I, the intermittent blowing strategy leads to significant drag reduction over each individual strip which does not fully recover to the canonical level in between each strip. Furthermore, the drag reduction persists far beyond the blowing region. The combination of reduced drag between the blowing regions and far downstream of the control region is the key factor which leads to improved net energy savings for this strategy. For Case 4, an MLDR of approximately 95% over the control area suggests that the boundary layer may experience near separation-like behaviour. A topic for future work would be to investigate wall-normal blowing strategies for turbulent boundary layers with non-zero pressure gradients.

## 5. Conclusions

To investigate the potential of wall-normal blowing as a control strategy for net energy saving in turbulent boundary layers, this work presents the results from two separate optimisation campaigns. These campaigns are an extension of the work by Mahfoze et al.<sup>14</sup> to include a greater range of blowing configurations and flow conditions. To achieve this, a high-order implicit LES solver (Xcompact3d) is combined with a BO framework to find optimal combinations of control

## A BAYESIAN OPTIMISATION FRAMEWORK FOR DRAG REDUCTION AND NET ENERGY SAVING

parameters that balance the drag reduction with the power required to drive the actuation, in order to achieve net energy saving.

The purpose of the first campaign was to investigate the effect of increasing the input space provided to the BO framework, allowing greater flexibility in the blowing configuration, as well as also including wall-normal suction. The results showed that net energy saving is difficult to achieve with this approach to power estimation, with only one case showing a positive net energy saving ( $S_2 = 0.3\%$ ). This case was characterised by 10 strips of equally-spaced intermittent blowing regions of moderate intensity, which provided moderate drag reduction but with a relatively low power requirement. The results also showed that suction significantly degrades the performance of the control strategy for this setup, with suction leading to increases in skin friction drag while still requiring significant power to drive the actuation. The purpose of the second campaign was to investigate the effect of Reynolds number, with a view towards real-world applications. The results were promising, with most cases showing some level of net energy saving, indicating that the efficiency of the control strategy improves with increasing Reynolds number. Intermittent blowing was again found to be the best strategy for net energy saving, with a value ( $S_2 = 3.24\%$ ) that is significantly higher than what has previously been demonstrated at lower Reynolds numbers.

There are a number of possible avenues for future work. Firstly, a more detailed analysis of the present results, particularly with regards to understanding the flow physics of high-performing cases, is necessary. Secondly, more optimisation campaigns should be conducted to investigate further the effects of higher Reynolds numbers, non-zero pressure gradients, and different blowing strategies (e.g. travelling waves of wall-normal blowing), to name a few.

## Acknowledgments

The authors would like to thank the Engineering and Physical Sciences Research Council (EPSRC) for the computational time made available on the UK supercomputing facility ARCHER/ARCHER2, via the UK Turbulence Consortium (EP/L000261/1). The authors also acknowledge PRACE for awarding access to Hazel Hen at HLRS, Germany (Project No. 2018184381). Omar Mahfoze would like to thank Imperial College London for funding his PhD via the Imperial College President Scholarship. The work has been supported by EPSRC under grant numbers EP/T021144/1 and EP/T020946/1.

## References

- [1] A. Abbas, G. Bugea, E. Ferrer, S. Fu, J. Periaux, J. Pons-Prats, E. Valero, and Y. Zheng. Drag reduction via turbulent boundary layer flow control, 2017.
- [2] H. Choi, P. Moin, and J. Kim. Direct numerical simulation of turbulent flow over riblets. *J. Fluid Mech.*, 255:503–539, 1993.
- [3] T. Dairay, E. Lamballais, S. Laizet, and J.C. Vassilicos. Numerical dissipation vs. subgrid-scale modelling for large eddy simulation. *J. Comput. Phys.*, 337:252–274, 2017.
- [4] Y.F. Fu, C.Q. Yuan, and X.Q. Bai. Marine drag reduction of shark skin inspired riblet surfaces. *Biosurface and Biotribology*, 3(1):11–24, 2017.
- [5] M.A. Gelbart, J. Snoek, and R.P. Adams. Bayesian optimization with unknown constraints. In *Proc. 30th Uncertain. Artif. Intell. Conf.*, pages 250–259, 2014.
- [6] Y. Kametani and K. Fukagata. Direct numerical simulation of spatially developing turbulent boundary layers with uniform blowing or suction. *J. Fluid Mech.*, 681:154–172, 2011.
- [7] J.S. Kim, J. Hwang, M. Yoon, J. Ahn, and H.J. Sung. Influence of a large-eddy breakup device on the frictional drag in a turbulent boundary layer. *Phys. Fluids*, 29(6):065103, 2017.
- [8] V.I. Kornilov and A.V. Boiko. Efficiency of air microblowing through microperforated wall for flat plate drag reduction. *AIAA J.*, 50(3):724–732, 2012.
- [9] A.G. Kravchenko and P. Moin. On the Effect of Numerical Errors in Large Eddy Simulations of Turbulent Flows. *J. Comput. Phys.*, 131(2):310–322, 1997.
- [10] S. Laizet and E. Lamballais. High-order compact schemes for incompressible flows: A simple and efficient method with quasi-spectral accuracy. *J. Comput. Phys.*, 228(16):5989–6015, 2009.

- [11] S. Laizet and N. Li. 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.
- [12] S.K. Lele. Compact finite difference schemes with spectral-like resolution. *J. Comput. Phys.*, 103(1):16–42, 1992.
- [13] O.A. Mahfoze and S. Laizet. Non-explicit large eddy simulations of turbulent channel flows from  $Re_\tau=180$  up to  $Re_\tau=5,200$ . *Comput. Fluids*, page 105019, 2021.
- [14] O.A. Mahfoze, A. Moody, A. Wynn, R.D. Whalley, and S. Laizet. 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.
- [15] Y. Morita, S. Rezaeiravesh, N. Tabatabaei, R. Vinuesa, K. Fukagata, and P. Schlatter. Applying Bayesian optimization with Gaussian process regression to computational fluid dynamics problems. *J. Comput. Phys.*, 449:110788, 2022.
- [16] P. Schlatter and R. Örlü. Assessment of direct numerical simulation data of turbulent boundary layers. *J. Fluid Mech.*, 659:116–126, 2010.
- [17] B. Shahriari, K. Swersky, Z. Wang, R.P. Adams, and N. De Freitas. Taking the human out of the loop: A review of Bayesian optimization, 2016.
- [18] J. Snoek, H. Larochelle, and R.P. Adams. Practical Bayesian optimization of machine learning algorithms. In *Adv. Neural Inf. Process. Syst.*, volume 4, 2012.
- [19] A. Stroh, Y. Hasegawa, P. Schlatter, and B. Frohnapfel. Global effect of local skin friction drag reduction in spatially developing turbulent boundary layer. *J. Fluid Mech.*, 805:303–321, 2016.
- [20] C.M. White and M.G. Mungal. Mechanics and prediction of turbulent drag reduction with polymer additives, 2008.