



# Data-driven optimisation of wind farm layout and wake steering with large-eddy simulations

Nikolaos Bempedelis<sup>1</sup>, Filippo Gori<sup>1</sup>, Andrew Wynn<sup>1</sup>, Sylvain Laizet<sup>1</sup>, and Luca Magri<sup>1,2</sup>

<sup>1</sup>Department of Aeronautics, Imperial College London, SW7 2AZ London, UK

<sup>2</sup>The Alan Turing Institute, London NW1 2DB, UK

**Correspondence:** Nikolaos Bempedelis (n.bempedelis20@imperial.ac.uk); Luca Magri (l.magri@imperial.ac.uk)

**Abstract.** Maximising the power production of large wind farms is key to the transition towards net zero. The overarching goal of this paper is to propose a computational method to maximise the power production of wind farms with two practical design strategies. First, we propose a gradient-free method to optimise the wind farm power production with high-fidelity surrogate models based on large-eddy simulations and a Bayesian framework. Second, we apply the proposed method to maximise wind farm power production by both micro-siting (layout optimisation) and wake steering (yaw angle optimisation). Third, we compare the optimisation results with the optimisation achieved with low-fidelity wake models. Finally, we propose a simple multi-fidelity strategy by combining the inexpensive wake models with the high-fidelity framework. The proposed gradient-free method can effectively maximise wind farm power production. Performance improvements relative to wake-model optimisation strategies can be attained, particularly in scenarios of increased flow complexity, such as in the wake steering problem, in which some of the assumptions in the simplified flow models become less accurate. The optimisation with high-fidelity methods takes into account nonlinear and unsteady fluid mechanical phenomena, which are leveraged by the proposed framework to increase the farm output. This paper opens up opportunities for wind farm optimisation with high-fidelity methods and without adjoint solvers.

## 1 Introduction

The transition towards a net-zero world is intertwined with a continually growing demand for renewable energy sources. Wind power has emerged as the leading contributor to renewable electricity in numerous countries (Our World in Data, 2022). Governments worldwide have recognised the significance of this resource and have committed to substantial capacity expansions as part of their energy strategies. As wind energy takes on an increasingly central role as an energy source, it is critical for it to mature in terms of efficiency and resilience. This entails addressing a number of challenges, which range from turbine aerodynamics, their interaction with the atmosphere, and plant-level control (Veers et al., 2019, 2022).

One key issue is that as wind farms grow in size, they become increasingly less efficient. This degradation in efficiency is primarily attributed to wake losses (Barthelmie et al., 2009), which arise when a significant portion of the turbines operate within the wake fields generated by neighbouring turbines within the same farm. A wind turbine operating within a wake field



is a negative operability issue for two reasons. First, the reduction of its power output due to the decelerated incoming wind,  
25 and, second, the increase of fatigue loading due to higher levels of turbulence.

The long-established strategy for mitigating wake losses is to optimally position the wind turbines across the available land. This approach is most commonly referred to as *layout optimisation* or *micro-siting*. Building upon the work by Mosetti et al. (1994), numerous studies have been published in the literature on maximising the production (or minimising the cost of energy production) of a wind farm by optimising turbine placement (see, for instance, Grady et al. (2005); Kusiak and Song (2010);  
30 Chowdhury et al. (2012); Stanley and Ning (2019), and the review articles by Herbert-Acero et al. (2014) and Shakoor et al. (2016)).

More recently, an approach called *wake steering* was proposed as an alternative strategy to mitigate wake losses. In wake steering, upstream turbines are yawed to deflect their wakes away from downstream turbines. Although this results in a reduction in power production from the upstream turbines, the downstream turbines generate more power, leading to an overall  
35 increase in the farm's power output. Demonstrations of performance gains in a number of computational (Fleming et al., 2015; Gebraad et al., 2016), experimental (Adaramola and Krogstad, 2011; Campagnolo et al., 2016; Bastankhah and Porté-Agel, 2019) and field (Fleming et al., 2017; Howland et al., 2019; Fleming et al., 2019; Simley et al., 2021) studies led to the recent implementation of wake steering in commercial wind plants (see, for instance, a press release by WindESCo (2023)).

In the vast majority of relevant studies published in the literature, the optimisation of the farm layout or the turbine yaw angles  
40 is carried out with low-fidelity flow solvers, commonly known as *wake models*, as cost function evaluators. Wake models are tools that estimate the velocity deficit downstream of a wind turbine and, by means of superposition, the overall flow in a wind farm (see, for example, Jensen (1983); Frandsen et al. (2006); Bastankhah and Porté-Agel (2014); Bempedelis and Steiros (2022), and the review article by Porté-Agel et al. (2020)). Wake models are built on simplified assumptions for the turbulent wake of a porous disk and do not account for various phenomena such as unsteadiness, non-linear interactions, or blockage,  
45 among others. As a consequence, optimisation based on wake models misses opportunities to improve farm performance by manipulating and potentially exploiting these unaccounted phenomena. Furthermore, wake models can produce predictions which are non-smooth or discontinuous as functions of the design variables (e.g., turbine coordinates or yaw angles). This renders their use within gradient-based optimisation algorithms challenging (Gori et al., 2023). Nevertheless, wake models are almost invariably used in optimisation studies due to their low computational cost, and the expertise gained from decades of  
50 development and application (Asmuth et al., 2023).

On the other hand, higher fidelity flow models have been employed in a limited number of layout optimisation studies (King et al., 2017; Antonini et al., 2018; Allen et al., 2020; Antonini et al., 2020). In these works, wind farm layouts were optimised with an adjoint gradient-based approach and steady-state Reynolds-averaged Navier-Stokes (RANS) simulations. The use of steady RANS enables the consideration of a number of the aforementioned phenomena, such as blockage or  
55 pressure effects. However, the unsteady wake dynamics and the wake-to-wake and atmosphere-to-wake interactions, which play a critical role in large wind farm flows, cannot be appropriately accounted for. While higher-fidelity flow models, such as large-eddy simulations (LES), aptly take these phenomena into account, the computation of the gradient of a long-term<sup>1</sup>

---

<sup>1</sup>By "long-term" we mean after many Lyapunov times.



time average of a chaotic turbulent flow is ill posed (Wang, 2013; Blonigan and Wang, 2018; Huhn and Magri, 2022), which is a limit for adjoint methods. Despite these challenges, large-eddy simulations and adjoint methods have been combined in the context of wind farm control (Goit and Meyers, 2015; Munters and Meyers, 2018), showcasing the potential of dynamic control strategies. Alternatively, LES have been utilised in hybrid or multi-fidelity approaches (Bokharaie et al., 2016; Kirby et al.). For instance, in Bokharaie et al. (2016), LES was employed to calibrate the parameters of wake models, which, in turn, were used to find optimal wind farm layouts.

In this work, we propose a data-driven framework for optimising the output of a wind farm based on a gradient-free Bayesian approach and high-fidelity large-eddy simulations of the wind farm flow. The structure of the paper is as follows. The proposed optimisation strategy is outlined in Sect. 2. Section 3 discusses its application to a sixteen-turbine layout optimisation problem. In Sect. 4, the framework is applied to a ten-turbine yaw angle optimisation (wake steering) problem. Finally, conclusions end the paper in Sect. 5.

## 2 Methodology

The methodology is presented in two parts. First, we introduce the optimisation framework in Sect. 2.1. Second, in Sect. 2.2, we describe the flow solver employed to predict the wind farm flow and power output.

### 2.1 Bayesian optimisation

Bayesian optimisation (BO) is a gradient-free global optimisation strategy, which is particularly attractive for optimising complex functions (as in large codes for which the adjoint algorithm is cumbersome). Its effectiveness has been demonstrated in various fluid mechanical applications (e.g., Mahfoze et al. (2019); Huhn and Magri (2022); O'Connor et al. (2023)). It seeks for the extrema of the objective function by constructing a probabilistic surrogate model of it, typically a Gaussian process (GP), and by exploring the parameter space with an acquisition function (AF).

A Gaussian Process (GP) model is a non-parametric probabilistic model that is defined by a prior mean function,  $\mu_0(\mathbf{x})$ ,  $\mathbf{x} \in \mathbb{R}^d$ , which is usually assumed zero, and a covariance function (also known as the kernel),  $k \in \mathbb{R}(\mathbf{x}, \mathbf{x}')$ . Given a set  $\mathcal{D}$  of  $m$  observations  $\mathbf{Y} \in \mathbb{R}^{m \times 1} = f(\mathbf{X}) = [y_1, \dots, y_m]^T$  at the input locations  $\mathbf{X} \in \mathbb{R}^{m \times d} = [\mathbf{x}_1, \dots, \mathbf{x}_m]^T$ , as training data, and assuming a zero prior mean, we can compute the posterior mean,  $\boldsymbol{\mu}_* \in \mathbb{R}^{n \times 1}$ , and variance,  $\boldsymbol{\sigma}_*^2 \in \mathbb{R}^{n \times 1}$ , for a set of  $n$  test points  $\mathbf{X}_* \in \mathbb{R}^{n \times d}$  as

$$\boldsymbol{\mu}_* = \mathbf{K}_*^T (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{Y} \quad (1)$$

$$\boldsymbol{\sigma}_*^2 = \text{diag} \left( \mathbf{K}_{**} - \mathbf{K}_*^T (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{K}_* \right) \quad (2)$$

where  $\mathbf{K} \in \mathbb{R}^{m \times m} = k(\mathbf{X}, \mathbf{X})$ ,  $\mathbf{K}_* \in \mathbb{R}^{m \times n} = k(\mathbf{X}, \mathbf{X}_*)$ ,  $\mathbf{K}_{**} \in \mathbb{R}^{n \times n} = k(\mathbf{X}_*, \mathbf{X}_*)$ , and  $\sigma^2 \in \mathbb{R}$  is the observational noise. These expressions provide an estimate of the underlying objective function and the corresponding uncertainty for every point in the parameter space. *Training* a GP model involves determining the values of the kernel hyper-parameters that



maximise the marginal likelihood. For a more detailed discussion on GPs, the reader is referred to Rasmussen and Williams (2006).

90 The acquisition function  $AF(\mathbf{x})$ , which is built on the surrogate model whose first two statistical moments are given by equations (1)-(2), determines the next point(s) to be evaluated, by taking into account both the mean and uncertainty for any point  $\mathbf{x}$  of the parameter space, and providing the probability, or amount, by which  $\mathbf{x}$  can improve the current optimum. Exploration, which directs the search towards unexplored areas of the parameter space, and exploitation, which focuses the search in the vicinity of known promising solutions, are balanced based on the trade-off between uncertainty reduction and  
95 potential improvement. In brief, Bayesian optimisation is described in Algorithm 1.

---

**Algorithm 1** Bayesian Optimisation

---

```
1: Acquire a set of initial objective function observations  $\mathcal{D}$ 
2: Initialise the variables holding the best solution,  $\mathbf{x}_{\text{opt}}$  and  $y_{\text{opt}}$ 
3: for  $i = 1$  to stopping criterion do
4:   // Update the surrogate model using the available data
5:   Train the GP model given  $\mathcal{D}$ 
6:   // Find the next point to evaluate using the AF
7:    $\mathbf{x}_{\text{next}} \leftarrow \arg \max AF(\mathbf{x})$ 
8:   // Evaluate the objective function at  $\mathbf{x}_{\text{next}}$ 
9:    $y_{\text{next}} \leftarrow f(\mathbf{x}_{\text{next}})$ 
10:  // Update the dataset with the new evaluation
11:  Append  $\mathbf{x}_{\text{next}}$  and  $y_{\text{next}}$  to  $\mathcal{D}$ 
12:  // Update the best solution found so far
13:  if  $y_{\text{next}} > y_{\text{opt}}$  then
14:     $\mathbf{x}_{\text{opt}} \leftarrow \mathbf{x}_{\text{next}}$ 
15:     $y_{\text{opt}} \leftarrow y_{\text{next}}$ 
16:  end if
17: end for
18: return  $\mathbf{x}_{\text{opt}}, y_{\text{opt}}$ 
```

---

In the current work, Bayesian optimisation is implemented with the `GPyOpt` library (The GPyOpt authors, 2016). As discussed in more detail in Sect. 3.2, we consider GP priors characterised by the Rational Quadratic kernel (Rasmussen and Williams, 2006). We further assume that our LES observations are high-fidelity, and, as such, noise-free (i.e.,  $\sigma^2$  is negligible). To select the next observation point(s), we use the Lower Confidence Bound acquisition function,  $AF(\mathbf{x}) = \mu_*(\mathbf{x}) - \lambda\sigma_*(\mathbf{x})$ ,  
100 which allows for straightforward control of the exploration-to-exploitation ratio through the tunable parameter  $\lambda$ , which weighs the mean/exploitation  $\mu_*$  and uncertainty/exploration  $\sigma_*$  terms of the GP model, with larger values of  $\lambda$  favouring exploration. Other AFs (e.g., Probability of Improvement, Expected Improvement) or combination of them may be used (e.g., Huhn and Magri (2022)). The kernel hyper-parameter and AF optimisations are carried out with the L-BFGS algorithm (Liu and Nocedal,



1989). Finally, we compute design points to be explored in batches (batch BO), following the local penalisation method  
105 proposed in González et al. (2016). This allows us to efficiently utilise our high-performance computing resources by running  
several LES at the same time.

## 2.2 Large-eddy simulations

To predict the wind farm flow and power output, we employ the wind farm simulator `winc3d` (Deskos et al., 2020) of the  
open-source finite-difference framework `xcompact3d` (Bartholomew et al., 2020). `xcompact3d` solves the incompressible  
110 filtered Navier-Stokes equations on Cartesian meshes using sixth-order compact schemes and a third-order Adams-Bashforth  
method for time integration (Laizet and Lamballais, 2009). Efficient scaling up to hundreds of thousands of CPU cores is made  
possible by utilising the `2Decomp` & `FFT` library for parallelisation (Laizet and Li, 2011). In this work, the Smagorinsky sub-  
filter scale model (Smagorinsky, 1963) is used to account for the effects of the unresolved fluid motions. The wind turbines  
are modelled with the actuator disk method, which is described in more detail, together with a related validation study, in  
115 Bempedelis et al. (2023). `winc3d` has been validated against experimental and field data in several wind turbine and wind  
farm flow problems (Deskos et al., 2019, 2020; Bempedelis and Steiros, 2022; Steiros et al., 2022; Jané-Ippel et al., 2023).

In order to realistically simulate the interaction of the wind farm with the turbulent atmospheric flow, we perform precursor  
simulations of fully-developed neutral atmospheric boundary layers. The precursor atmospheric flow simulations are run  
until statistical convergence is reached, after which we begin storing two-dimensional planes of the flow field, normal to the  
120 streamwise direction, at every timestep. These planes are subsequently used as inflow conditions in the wind farm simulations.

## 3 Layout optimisation (micro-siting)

In layout optimisation, the objective is to maximise the power production of a wind farm with  $N$  wind turbines by adjusting  
their position  $\mathbf{c} = [\mathbf{x}, \mathbf{y}]^T$ , with  $\mathbf{x} = (x_1, \dots, x_N)$  and  $\mathbf{y} = (y_1, \dots, y_N)$ , within a given space  $\mathbf{X}$ . The space  $\mathbf{X}$  corresponds to  
the available land where the wind turbines may be installed. The farm power production is typically a weighted sum of all the  
125 different wind states  $K$  (wind direction, speed, turbulence) which the farm experiences over a period of time. The turbines can  
align themselves with the mean direction of incoming wind. The optimisation problem is therefore expressed as

$$\arg \max_{\mathbf{c}} \sum_{k=1}^K a_k \sum_{n=1}^N P_{n,k}(\mathbf{c}) \quad (3a)$$

$$\text{s.t.} \quad \mathbf{c} \in \mathbf{X}, \quad (3b)$$

$$\|\mathbf{c}_i - \mathbf{c}_j\| > D, \text{ for } i, j = 1, \dots, N \text{ and } i \neq j \quad (3c)$$

130 where  $P_{n,k}$  is the power output of the  $n$ -th turbine given wind state  $k$ , and  $a_k$  is the weight (i.e. probability) of each wind state,  
as dictated by the local meteorological data. Equation (3b) constrains the turbines within the available land. To avoid turbine  
overlapping, a second constraint is added to ensure that the centres of the turbines are spaced at least one turbine diameter  $D$   
apart (Eq. (3c)).



In a preliminary proof-of-concept study (Bempedelis and Magri, 2023), an early version of the proposed framework was  
 135 applied to a simple layout optimisation problem (five wind turbines in a  $6D \times 6D$  square, operating under a single wind state).  
 It was shown that the method is capable of finding optimal designs in a few iterations whilst leveraging flow phenomena that  
 are unaccounted for in low-fidelity solvers (e.g., local speed-ups) to deliver increased wind farm performance.

However, modern wind farms are composed of several more turbines. The layout optimisation problem therefore becomes  
 high-dimensional, with the design space being  $2N$ -dimensional. Bayesian optimisation in low dimensions is well established,  
 140 but its application to high-dimensional problems poses a significant challenge due to the increase in complexity in both GP  
 regression and the search for the next designs (AF optimisation) (Binois and Wycoff, 2022).

Here, we propose a way to reduce the dimensionality of the layout optimisation problem. It is based on the assumption that,  
 in most cases, the turbines in a wind farm are identical. In that case, the cost function is invariant under permutations of the  
 turbine labelling, and there are  $N!$  different combinations of the design variables that produce the same output. Such behaviour  
 145 is a clear barrier to efficient optimisation. To remedy this, we could augment the training dataset with all  $N!$  permutations of the  
 evaluated designs. However, this quickly gives rise to high complexity, as the computational cost for GP regression scales with  
 $\mathcal{O}(m^3)$ , where  $m$  is the number of points in the training dataset. Alternatively, without loss of generality, we propose working  
 on a reduced design space by enforcing a set of  $N - 1$  additional constraints (Eq. (4d)). The constraints enforce a particular  
 ordering of the turbines, in one direction only, thereby constraining the design variables to that particular subspace of the full  
 150 parameter space. The subspace size is smaller by a factor  $1/N!$  relative to the original one, indicating the equivalence of the  
 two approaches. (A third way, not explored in this work, is to encode the turbine ordering invariance property of the objective  
 function in the kernel; see Duvenaud (2014) for additional details). The layout optimisation problem is thus reformulated as

$$\arg \max_{\mathbf{c}} \sum_{k=1}^K a_k \sum_{n=1}^N P_{n,k}(\mathbf{c}) \quad (4a)$$

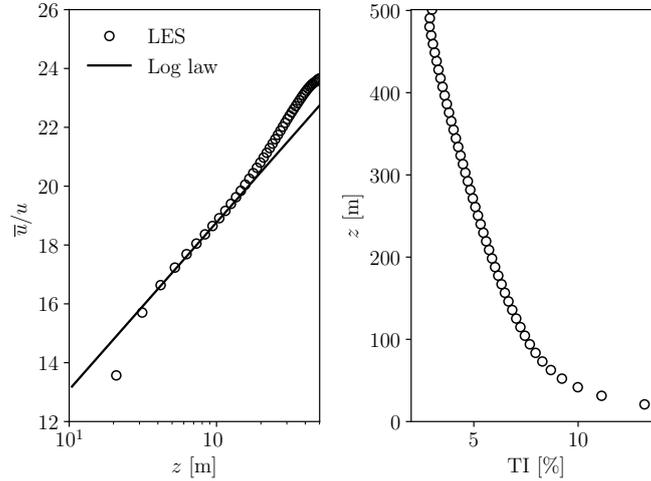
$$\text{s.t.} \quad \mathbf{c} \in \mathbf{X}, \quad (4b)$$

$$155 \quad \|\mathbf{c}_i - \mathbf{c}_j\| > D, \text{ for } i, j = 1, \dots, N \text{ and } i \neq j, \quad (4c)$$

$$x_i \geq x_{i-1}, \text{ for } i = 2, \dots, N \quad (4d)$$

### 3.1 Layout optimisation: Problem set-up

We consider a wind farm of  $N = 16$  turbines, each with  $D = 100$  m rotor diameter and  $h = 100$  m hub height. The wind  
 turbines operate at a constant thrust coefficient,  $C_T = 0.75$ . The available land is a square of size  $18D \times 18D$ . We consider  
 160 an evenly-weighted six-directional wind rose. This is a set-up close to those examined in King et al. (2017) and Antonini  
 et al. (2018). The atmospheric boundary layer is characterised by friction velocity  $u^* = 0.442 \text{ ms}^{-1}$ , height  $\delta = 501$  m and  
 roughness length  $z_0 = 0.05$  m, which correspond to conditions in the North Sea (Wu and Porté-Agel, 2015). The velocity at  
 the hub height of the turbines is  $U_h = 8.3 \text{ ms}^{-1}$ , and the turbulence intensity at the same level is  $\text{TI}_h = 7.4\%$  (see Fig. 1 for  
 the simulated atmospheric wind profiles).



**Figure 1.** Mean streamwise velocity (left) and turbulence intensity (right) of the simulated atmospheric boundary layer.

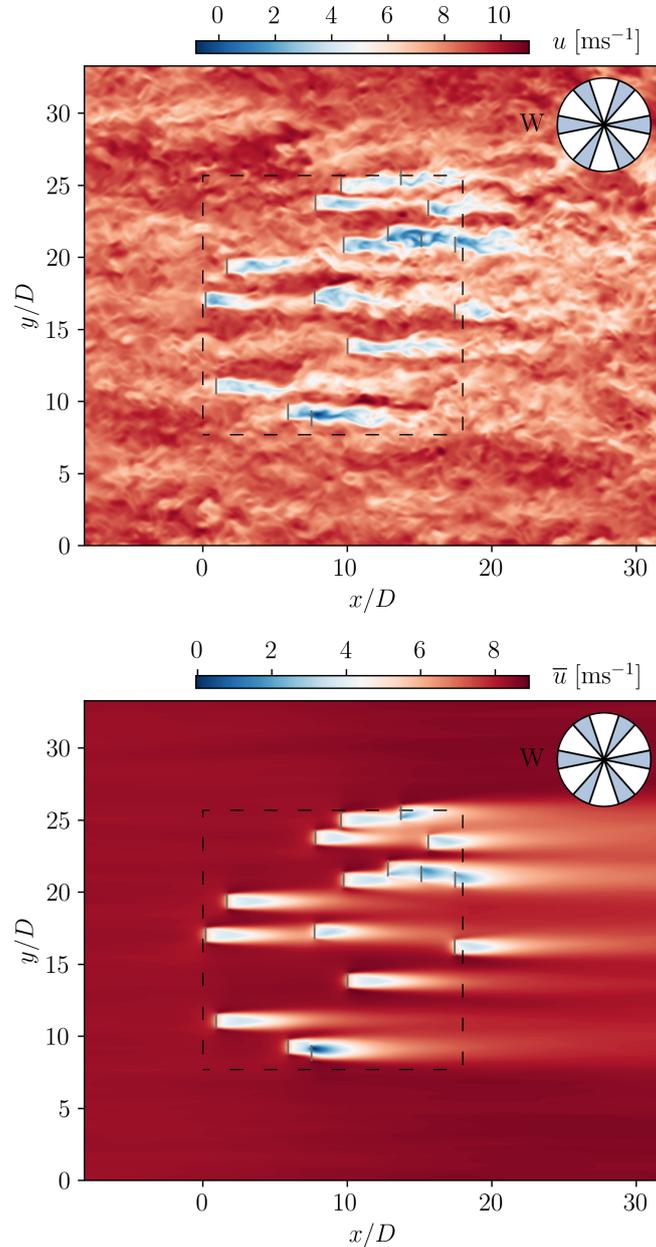
165 The wind farm is embedded in a computational domain of size  $4008 \times 3340 \times 501$  m, which is discretised with a uniform  
 mesh with  $\Delta x = \Delta y = \Delta z = 10.4375$  m (corresponding to  $\approx 10$  points across the turbine rotor). As shown in Bempedelis  
 et al. (2023), this resolution is sufficient to predict the power production of a large wind farm. This was also confirmed for  
 the considered set-up (result not shown for brevity). The timestep is chosen such that the maximum Courant (CFL) number  
 remains below 0.2. Data are averaged over a 2.5-hour period of farm operation, following the time required for the flow to  
 170 develop.

Figure 2 shows the contours of the instantaneous and mean streamwise flow at the turbine hub level for a randomly sampled  
 farm layout under westerly wind, together with details on the described set-up (wind rose and available land). It also reveals  
 the highly complex nature of wind farm flows, with turbine wakes interacting with one another and with the atmospheric  
 turbulence.

### 175 3.2 Layout optimisation: Results

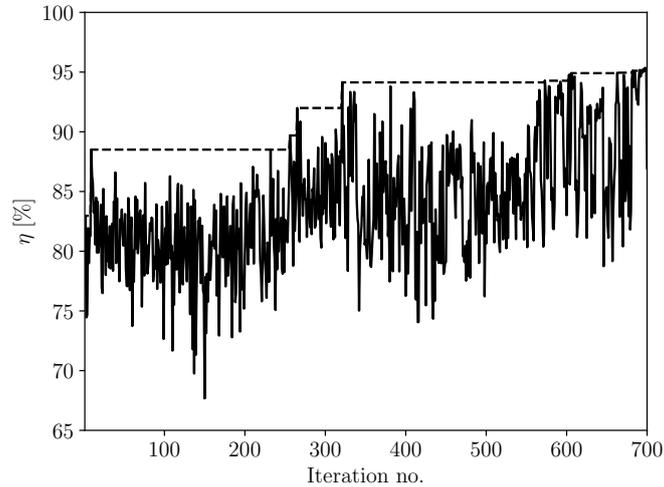
The optimisation starts by evaluating the farm power output of a large set of layouts sampled from a Latin hypercube, along  
 with two user-designed layouts. These are a uniform  $4 \times 4$  layout with 6D spacing between the turbines and a layout where the  
 turbines are equidistantly distributed on the land border only. The optimisation then progresses iteratively, as described in Sect.  
 2, and terminates after 700 iterations (for a total of 4200 wind farm LES). This was decided in consideration of the available  
 180 computational resources (a comment on the cost of the simulations can be found in Appendix A) and the optimisation's  
 progression. The optimisation history is shown in Fig. 3. To facilitate comparison, we introduce the metric  $\eta$ , defined as the  
 ratio between the overall farm power output and a reference ideal output based on the output of a single wind turbine

$$\eta = \frac{\sum_{k=1}^K a_k \sum_{n=1}^N P_{n,k}}{N \times P_{\text{single}}} \quad (5)$$



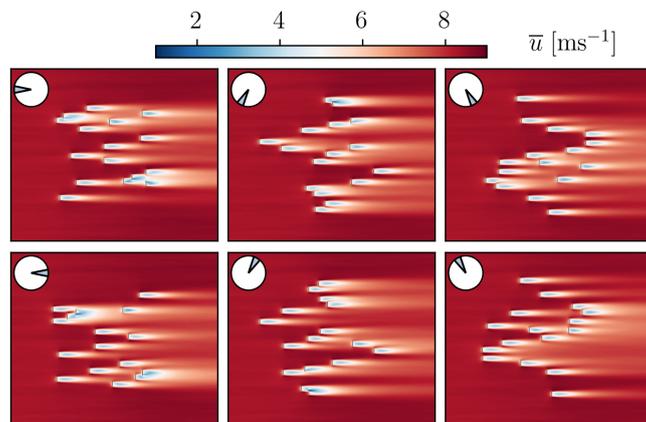
**Figure 2.** (Top) Instantaneous and (bottom) mean streamwise velocity at the turbine hub height for a randomly sampled layout, with the wind blowing from the west. The full six-directional wind rose is shown at the top right. The borders of the available land are indicated with dashed black lines. The wind turbines are denoted with grey solid lines.

with  $P_{\text{single}}$  being computed in a separate LES, where a single turbine was placed at the center of the available land. Basically,  
185  $\eta$  is a normalised form of the objective function defined in Eq. (4a), based on the production of  $N$  similar stand-alone turbines. It may therefore be interpreted as a farm efficiency coefficient primarily describing wake losses.



**Figure 3.** Efficiency of designs evaluated during the layout optimisation process. The dashed line shows the evolution of the best-performing design.

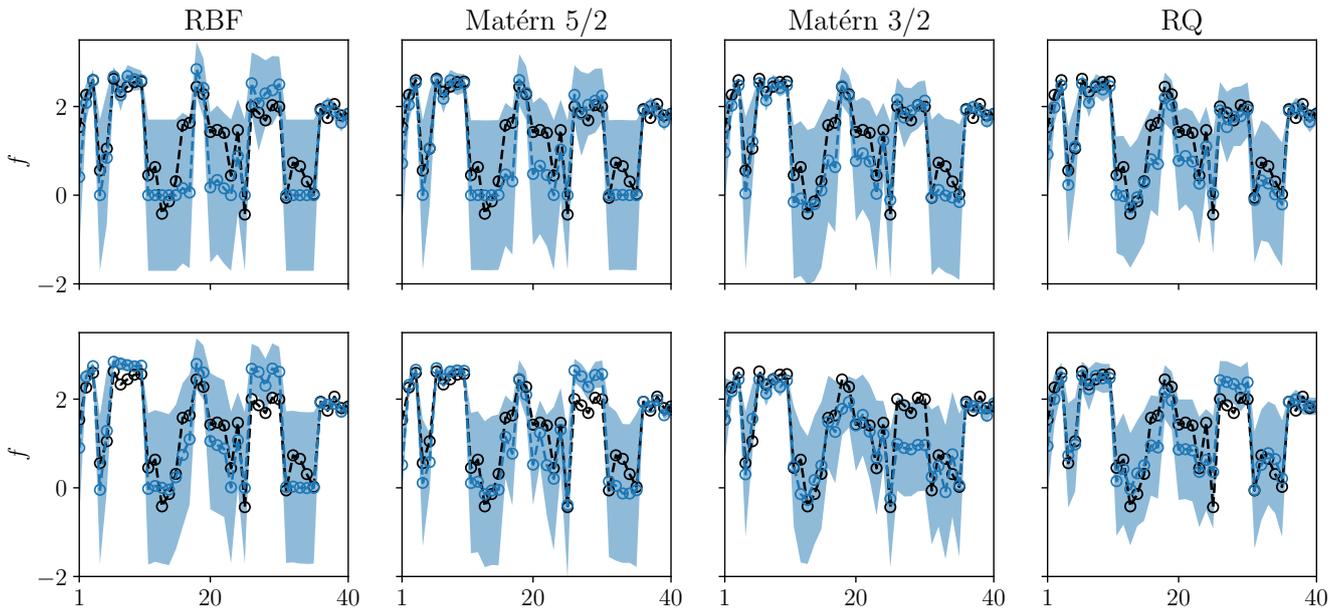
At the end of the optimisation, the best-performing design has  $\eta = 95.3\%$  efficiency, with the corresponding mean flow fields shown in Fig. 4 for every wind direction. We observe limited turbine-wake interference, especially in the diagonal wind directions, which is associated with increased power production by the wind farm in those directions. For reference, the  $4 \times 4$  aligned layout and the layout shown in Fig. 2 have efficiencies  $\eta = 88.5\%$  and  $\eta = 80.4\%$ , respectively.



**Figure 4.** Mean flow fields for the best-performing design proposed by the large-eddy simulation Bayesian optimisation framework (LES-BO). In all cases, the wind blows from left to right, with the wind farms rotated to match the wind direction shown on the inset wind roses. The reference non-rotated layout is that of the westerly wind case (top left).



In parallel with the training of the surrogate model, we can also learn what constitutes good modelling practices. By means of validation on subsequent test points, we observe that the rational quadratic (RQ) kernel outperforms both the widely used Matérn and the squared exponential (RBF) kernels. Figure 5 shows the predictions of the regressed GPs (with a training dataset of 600 points) for the 40 points tested next. The test points include designs suggested with different levels of balance between exploration and exploitation (by tuning the  $\lambda$  parameter in the AF, see Sect. 2.1). All kernels perform adequately in regions near points in the training dataset. However, the RQ kernel is shown to extrapolate more accurately. Similar conclusions hold for both anisotropic (ARD) kernels (in which each design variable is associated with its own lengthscales) and isotropic kernels. In Bempedelis and Magri (2023), it was shown that it is possible to find optimal designs with isotropic kernels. Here, we find that anisotropic kernels show slightly improved accuracy. However, this comes at the price of larger computational cost and complexity. In particular, care needs to be taken when using anisotropic kernels in high-dimensional problems, to ensure that sufficient data are available in the training set to achieve robustness in the GP hyper-parameter optimisation.

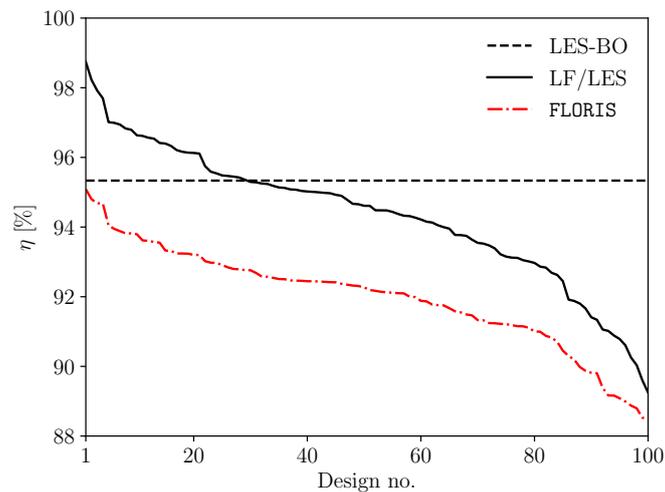


**Figure 5.** GP regression with different covariance functions on a standardised training dataset of 600 points. Predictions of the regressed models for the 40 points computed next in the optimisation chain. The GP mean predictions are denoted with blue markers. The shaded areas represent 95% confidence intervals. Black markers denote the LES predictions. (Top row): Isotropic kernels. (Bottom row): Anisotropic kernels.

In order to assess the quality of the best-performing design obtained with the proposed framework (hereafter referred to as LES-BO), we perform a series of optimisations with the FLOW Redirection and Induction in Steady State (FLORIS) software (NREL, 2023). FLORIS is an open-source wind farm simulator that incorporates several widely used wake models. In this study, the Gauss-Curl Hybrid model (King et al., 2021) is employed for predicting the wind farm flow and power output,

while the optimisation algorithm used is the gradient-based Sequential Least Squares Programming (SLSQP) method (Kraft, 1988). For more details on the FLORIS simulations, please see Appendix B. In consideration of the multi-modal nature of the layout optimisation problem, and to enable a more comprehensive comparison with the LES-BO method, we perform 100 independent FLORIS layout optimisations, initialised with 99 random layouts and a uniform  $4 \times 4$  layout with  $6D$  spacing  
210 between the turbines.

The efficiencies of the 100 optimal designs outputted by FLORIS are then evaluated using Winc3d (the designs suggested by the low-fidelity model and evaluated with LES will be referred to as LF/LES). The results, ordered from best to worst, are presented in Fig. 6, together with the predictions of FLORIS and the best-performing design obtained with LES-BO.



**Figure 6.** Efficiency of optimal designs outputted by FLORIS, evaluated with both FLORIS and LES (LF/LES). The dashed line shows the optimal design suggested by the LES-BO framework.

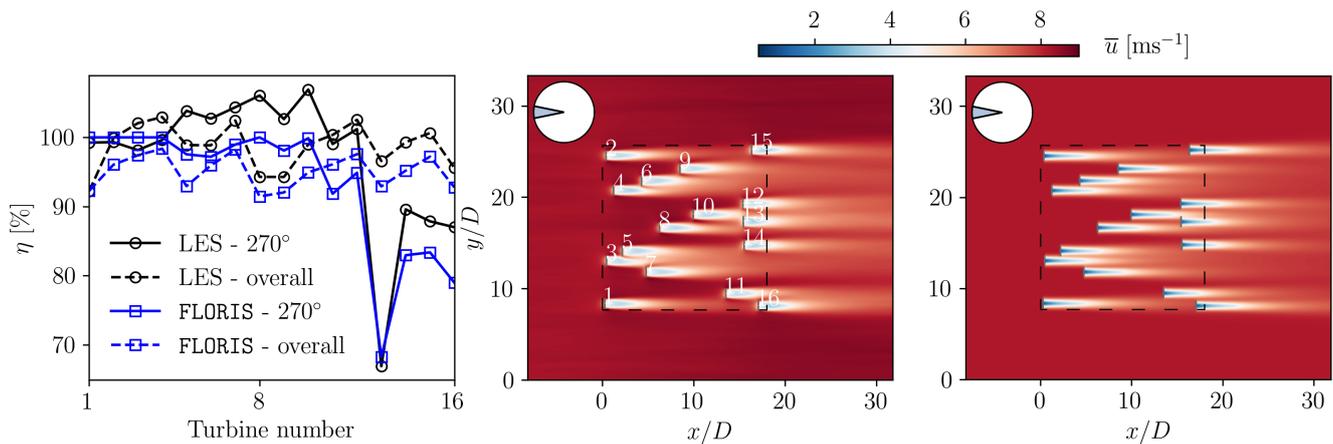
Figure 6 allows us to draw a number of conclusions. First, the proposed LES-BO framework is capable of finding a design  
215 that produces more power than  $\approx 70\%$  of the optimal designs suggested by FLORIS. Taking into account the relatively small amount of iterations to find this design, the option of using more advanced BO strategies (e.g., Eriksson et al. (2019)), and the fact that our focus is on feasibility rather than performance, we conclude that optimising wind farms using a high-fidelity surrogate modelling approach is achievable.

However, several layouts suggested by the wake model-based optimisation outperform those suggested by LES-BO. This  
220 finding supports the practice of using wake models to design wind farms in the wind energy industry, as they demonstrate excellent performance at only a fraction of the cost (see also Thomas et al. (2019)).

Nevertheless, the efficiency predictions of FLORIS are lower compared to those of LES. As shown in Fig. 6, the trends of both predictions are similar, but the offset between the two is not constant, with LES predicting a wider range of efficiencies. These differences are attributed to the flow phenomena that are unaccounted for in FLORIS. To identify these, Fig. 7 shows the  
225 efficiency of individual turbines of the FLORIS-suggested best design, evaluated by both FLORIS and LES, for a single wind

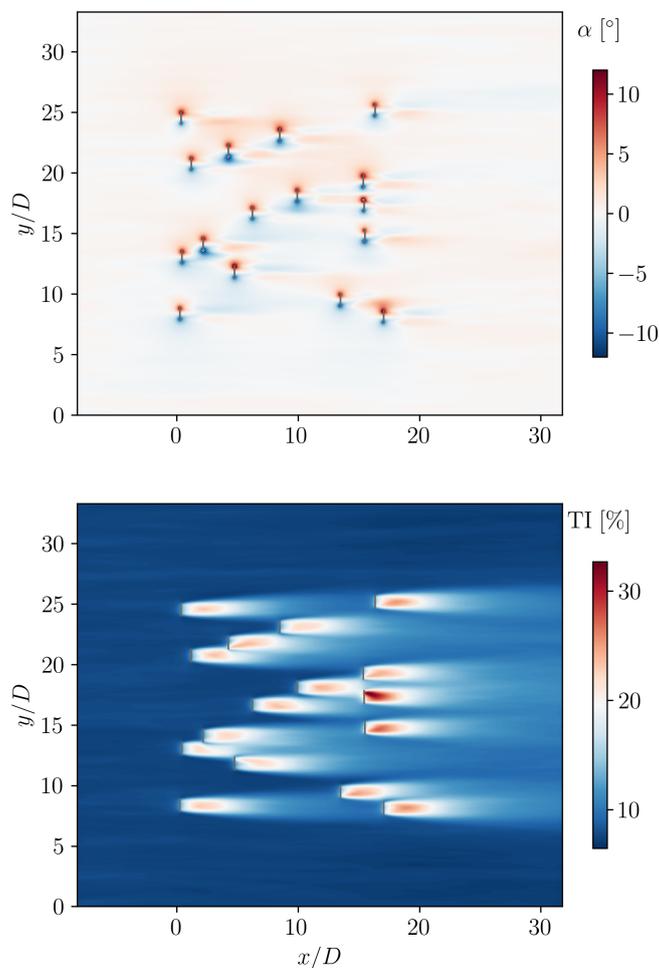


direction ( $270^\circ$ ) and the full wind rose. In the case of the westerly wind, we observe that the increase in power production comes from downstream turbines. According to LES, the first turbines are producing slightly less power, indicating the negative impact of global (farm-level) blockage that is not taken into account in the wake models. However, the production of downstream turbines is increased, to levels above those of an isolated turbine, due to local flow speed-ups. (Here, additional simulations with three times larger spanwise extent were performed to evaluate the impact of domain blockage. This was confirmed to be small, with differences in efficiency  $\approx 1\%$ ). Similar benefits owing to local blockage were also reported in King et al. (2017) and Antonini et al. (2018). In the case of the full wind rose, all turbines operate at an efficiency above 90%, with some turbines exceeding, on average, the production of an isolated turbine. Overall, the individual turbine efficiency predictions of the two flow models for the full wind rose (aggregate power) are qualitatively similar, with varying levels of differences.



**Figure 7.** FLORIS-suggested best design. (Left): Individual turbine efficiencies predicted by FLORIS and LES. (Middle) Mean streamwise flow for westerly wind, LES. (Right) Mean streamwise flow for westerly wind, FLORIS.

Figure 7 also shows the mean streamwise flow fields predicted by FLORIS and LES for the westerly wind case. Together with Fig. 8, which shows the LES predictions for the in-plane velocity angle and the turbulence intensity, we observe several phenomena, in addition to the local speed-ups, which are taken into account only in LES (see also figures 2 and 4). These include cross-stream flow, wake deflections and curvature (e.g., at the sides of the farm), pressure gradients, and wake merging (e.g., turbines 11 and 16). Another difference between the predictions of the two flow models is in the near wake of the turbines. Wake models assume that self-similarity is established immediately downstream of a turbine, which is known to be inaccurate (see e.g., Steiros et al. (2023)). Although this may not be a significant concern in the relatively sparse arrangement considered in this work, it is expected to play a more significant role in farms with higher turbine density. Finally, Fig. 9 shows a smoke-type visualisation of the turbine wakes for the FLORIS-suggested best farm layout. To achieve this, a passive scalar is emitted from the turbine rotors and is transported by the flow. This enables us to observe the meandering of the wakes, and possibly some indications of weakly—owing to the stabilisation effects of base bleed (Steiros et al., 2020, 2021)—coupled dynamics (e.g., turbines 12,13 and 14) (Peschard and Le Gal, 1996).

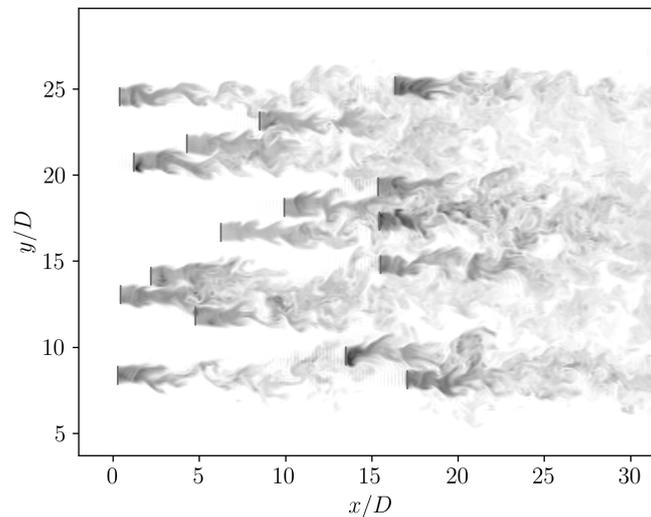


**Figure 8.** FLORIS-suggested best design. LES predictions for (top) angle of horizontal plane velocity vector, (bottom) turbulence intensity.

### 3.3 Combining wake models with LES: A multi-fidelity approach

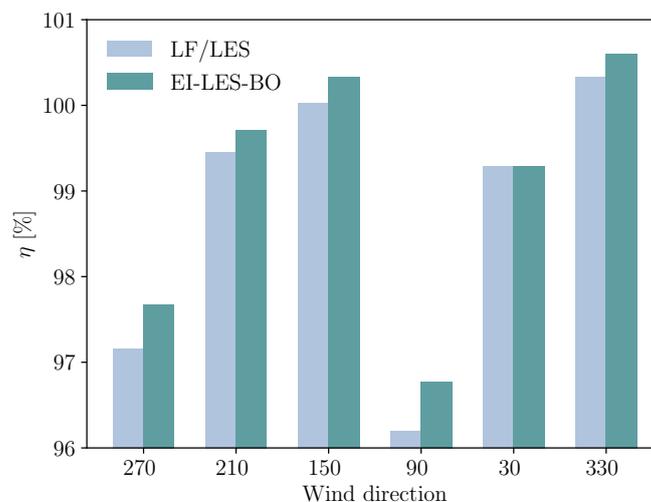
Section 3.2 demonstrated the feasibility of optimising the layout of wind farms with the proposed LES-BO framework, along with the role of the mechanisms that are neglected by low-fidelity flow solvers. However, for the number of optimisation iterations performed, LES-BO was outperformed by the multi-start gradient-based optimisation with low-fidelity models.

In order to show the benefits of adopting the LES-BO approach, we propose a simple multi-fidelity strategy. We train a surrogate model with a dataset composed of only the first 300 samples from our original investigation (see Fig. 3) and the 100 designs suggested by the low-order model (see Fig. 6). Following initial training, the optimisation proceeds as discussed in Sect. 2, and similar to Sect. 3.2. However, our surrogate is now informed with well-performing designs obtained from the low-fidelity optimisation. We refer to this strategy as the extrema-informed LES-BO framework (EI-LES-BO).



**Figure 9.** Wake visualisation in the FLORIS-suggested best design case by means of transport of a passive scalar shed from the turbine rotors.

Within just thirty iterations, EI-LES-BO manages to improve upon the FLORIS-suggested best design by  $\delta\eta = 0.32\%$ . This is achieved by leveraging the flow phenomena that are unaccounted for in FLORIS. Figure 10 shows the efficiency of the two designs for each direction in the wind rose. EI-LES-BO delivers increased performance for most wind directions.



**Figure 10.** Efficiency of best-performing LF/LES and EI-LES-BO layouts for each direction in the wind rose.



#### 4 Wake steering (yaw angle optimisation)

260 In wake steering, the objective is to maximise the power output of a wind farm by adjusting the angle  $\gamma$  at which turbines face the wind (a yaw angle  $\gamma = 0^\circ$  corresponds to the turbines being aligned with the mean wind direction). The position and properties of the wind turbines are considered known. As a result, optimisation can be carried out independently for each wind state. The (static) yaw angle optimisation problem for a single wind state can then be expressed as

$$\arg \max_{\gamma} \sum_{n=1}^N P_n(\gamma) \quad (6a)$$

265 s.t.  $\gamma \in \Gamma$  (6b)

where  $\gamma = (\gamma_1, \dots, \gamma_N)$  denotes the turbine yaw angles and  $\Gamma$  is the range of admissible misalignment angles.

##### 4.1 Wake steering optimisation: Problem set-up

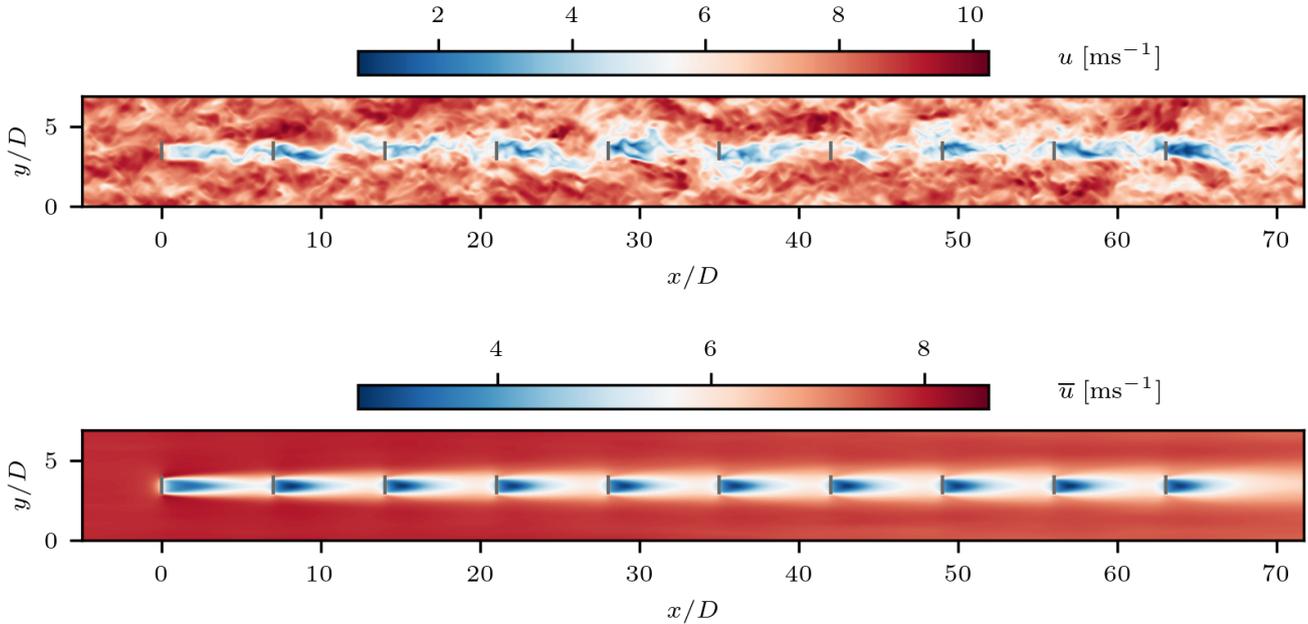
The problem we consider draws on the Horns Rev I wind farm, which is located in the North Sea, and consists of eighty turbines with diameter  $D = 80$  m and hub height  $h = 70$  m, arranged in an  $8 \times 10$  grid with  $7D$  spacing between the turbines  
270 (Barthelmie et al., 2009). Data for the atmospheric boundary layer are taken from Wu and Porté-Agel (2015), with  $u^* = 0.442$   $\text{ms}^{-1}$ , height  $\delta = 500$  m and roughness length  $z_0 = 0.05$  m (similar to the conditions considered in Sect. 3). Under these conditions, the velocity and turbulence intensity at hub height are  $U_h = 7.9$   $\text{ms}^{-1}$  and  $\text{TI}_h = 8.0\%$ , respectively.

We assume that the turbines operate with a constant thrust coefficient,  $C_T = 0.75$ , and may yaw up to  $\pm 30^\circ$  ( $\Gamma \in [-30, 30]$ ). We consider a single wind direction ( $270^\circ$ ), which is among the most prevalent ones, and the wind is aligned with the wind  
275 farm rows. This case is selected because wake steering optimisation studies based on low-fidelity models have yielded wildly varying optimum yaw settings, which range from zero to the maximum admissible yaw angle (see e.g., Zong and Porté-Agel (2021); Gori et al. (2023)). For the purposes of the present optimisation study, we simulate only one row of the wind farm and use periodic conditions on the lateral boundaries. This follows a number of studies that have reported row independence for the considered problem (Zong and Porté-Agel, 2021; Gori et al., 2023). For instance, Zong and Porté-Agel (2021) reported an  
280 estimated asymptotic wake deflection of  $0.8D$  under a yaw condition of  $\gamma = 30^\circ$ , effectively demonstrating that row-to-row interactions have negligible influence on wake steering effectiveness for the investigated turbine spacings.

The domain of size  $6144 \times 560 \times 1024$  km is discretised with a uniform mesh with  $\Delta x = \Delta y = \Delta z = 8$  m (corresponding to 10 points across the turbine rotor, a resolution equivalent to the one used in Sect. 3). Data are gathered over a 3-hour period of farm operation, following the time required for the flow to develop. Figure 11 shows contours of the instantaneous and mean  
285 streamwise flow at the turbine hub level, for the case of non-yawed conditions, ( $\gamma = 0^\circ$ ).

##### 4.2 Wake steering optimisation: Results

The initial training dataset comprises fifty designs. These include yaw angle combinations sampled with the Latin Hypercube method, the non-yawed condition shown in Fig. 11, and a design in which the yaw angles decrease linearly, from the maximum



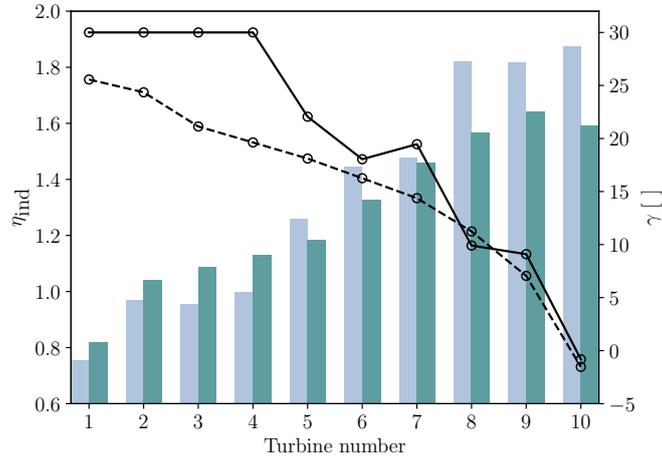
**Figure 11.** (Top) Instantaneous and (bottom) mean streamwise velocity at the turbine hub height for non-yawed conditions.

admissible yaw angle for the most upstream turbine to zero yaw angle for the most downstream one (Zong and Porté-Agel, 2021). Significant improvements in power output are observed within just ten iterations of the LES-BO optimisation. To facilitate comparison between all considered conditions, we define the metric  $\eta$  as the ratio of improvement over the non-yawed case, namely when  $\gamma = 0^\circ$

$$\eta = \frac{\sum_{n=1}^N P_n(\gamma)}{\sum_{n=1}^N P_n(\gamma = 0^\circ)} \quad (7)$$

The above metric is a measure of power output improvements when adopting wake steering. Similar to Sect. 3, the best performing LES-BO design is compared with the FLORIS-suggested optimal design (Fig. 12). In particular, Fig. 12 shows the optimal yaw angles suggested by each framework, along with the corresponding LES predictions for individual turbine efficiency. The latter is defined as  $\eta_{\text{ind},i} = P_i/P_i(\gamma = 0^\circ)$ .

Qualitatively, both frameworks suggest a decreasing trend in yaw angles as we move downstream in the farm, similar to results reported in Bastankhah and Porté-Agel (2019). However, differences in magnitude can be observed, with the LES-BO framework favouring larger yaw angles in the upstream turbines. This results in reduced power output for the first four turbines. For downstream turbines, both frameworks suggest yaw angles of similar magnitudes. However, because they are less affected by the wakes of upstream turbines, downstream turbines in the LES-BO design produce more power. Overall, the LES-BO framework is able to find a design that is 4% more efficient compared to the design suggested by FLORIS, with efficiencies

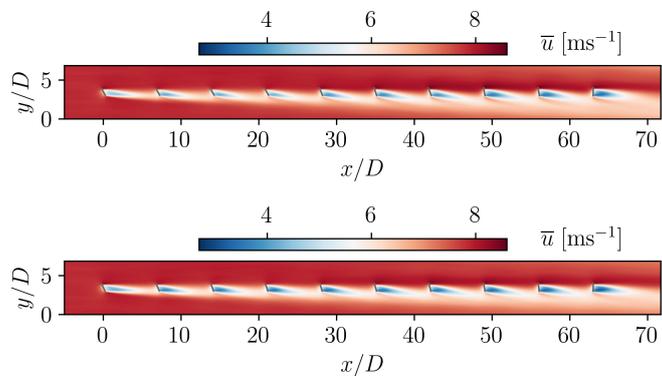


**Figure 12.** Optimal yaw angles suggested by LES-BO (solid line) and FLORIS (dashed line). Individual turbine efficiencies are shown as bars with solid fill (LES-BO) and bars with hatched pattern (LF/LES).

305  $\eta = 1.28$  and  $\eta = 1.24$ , respectively. Figure 13 shows the mean streamwise velocity field for the two cases (LES-BO best and LF/LES).

0.7% out of the 4% increase in efficiency is provided by leveraging the effects of confinement (blockage). This is measured by testing the designs in LES with five times the spanwise extent of the domain, effectively simulating them as individual rows of turbines. The difference in efficiency between the designs in the unconfined case reduced to 3.3%. Besides affecting the flow both locally and globally (see discussion in Sect. 3.2), confinement also affects wake recovery through increased levels of shear.

310



**Figure 13.** Mean streamwise flow at the turbine hub height. (Top) LES-BO, (bottom) LF/LES.



## 5 Conclusions

This work introduces an optimisation framework aimed at enhancing the efficiency of power production in wind farms. The proposed method follows a Bayesian approach and utilises surrogate models based on high-fidelity large-eddy simulations of wind farm flows. As part of an extensive computational campaign involving around 5000 large-eddy simulations, the framework was effectively used to mitigate losses caused by wake effects through two distinct strategies: layout optimisation (also known as micro-siting), and wake steering through yaw angle optimisation. The achieved optimisation outcomes were also compared with results obtained from low-fidelity wake model-based optimisation. Finally, a simple strategy was proposed to combine both large-eddy simulations and wake models in a multi-fidelity approach.

In the layout optimisation problem, the performance of wake models was notably confirmed to be highly effective, with the best layouts found having  $\approx 10\%$  increased efficiency compared with the efficiency of a standard aligned layout. In the case of wake steering, the framework relying on high-fidelity simulations outperformed the wake model-based optimisation by a considerable margin (4% increased efficiency, reaching an overall 28% increase compared with the non-yawed conditions). The increased complexity of the flow in this scenario rendered the utilisation of high-fidelity flow models advantageous. In future work, this framework will be extended to other flow cases of increased complexity, such as wind farms on complex terrain or under stratified conditions.

### Appendix A: Computational cost

Simulations were performed using 128 or 256 CPU cores on ARCHER2 (<https://www.archer2.ac.uk/>) and CSD3 (<https://www.csd3.cam.ac.uk/>). The cost of each simulation ranged from  $\approx 300$  to  $\approx 900$  CPUh, mostly depending on the type and amount of outputted data. These included instantaneous and time-averaged fields, and time-resolved probe and turbine data.

### Appendix B: Low-order modelling and optimisation

Low-order farm modelling is conducted using version 3.4 of FLORIS (NREL, 2023), an open-source wind farm simulator developed by the National Renewable Energy Laboratory (NREL). The framework incorporates several widely-used steady-state analytical wake models to predict the wind farm flow and power output. In this study, we employ the Gauss-Curl Hybrid wake model (King et al., 2021), with the deflection model of Bastankhah and Porté-Agel (2016), and the sum of squares freestream superposition (SOSFS) model of Katic et al. (1987). For all models, we use the FLORIS default parameters. The atmospheric conditions specified in sections 3.1 and 4.1 are matched by specifying a wind profile with a power law relationship (with a best fitting exponent 0.133) and turbulence intensity at the hub height as computed in the LES.

In terms of optimisation, the algorithm of choice is the gradient-based Sequential Least Squares Programming (SLSQP) method (Kraft, 1988). In the farm layout optimisation, 100 different optimisations are independently conducted, each starting from different initial conditions. The SLSQP parameters  $ftol = 10^{-9}$  and  $eps = 0.01$  are set following FLORIS recommendations (NREL, 2023). In the wake steering optimisation problem, the optimal yaw angles obtained in Gori et al. (2023) are set



as initial conditions to avoid initialisation sensitivity issues. Furthermore, the SLSQP parameters  $ftol = 10^{-12}$  and  $eps = 0.05$  are defined as suggested in Gori et al. (2023) for wake steering application on the Horns Rev wind farm.

*Data availability.* The data that support this study are available upon reasonable request.

345 *Acknowledgements.* NB, SL and LM are supported by EPSRC, Grant No. EP/W026686/1. SL and LM also acknowledge financial support from EPSRC, Grant No. EP/Y005619/1. LM also acknowledges financial support from the ERC Starting Grant PhyCo No. 949388. The authors would like to thank the UK Turbulence Consortium (EP/R029326/1 and EP/X035484/1) for providing access to ARCHER2. We also acknowledge use of the Cambridge Service for Data Driven Discovery (CSD3) system. FG would like to acknowledge the Department of Aeronautics, Imperial College London, for supporting this work with a doctoral studentship.

350 *Author contributions.* Nikolaos Bempedelis: Conceptualisation, Methodology, Investigation, Writing - Original Draft. Filippo Gori: Investigation, Writing - Review & Editing. Andrew Wynn: Writing - Review & Editing. Sylvain Laizet: Writing - Review & Editing, Funding acquisition. Luca Magri: Conceptualisation, Writing - Review & Editing, Funding acquisition.

*Competing interests.* The authors declare that they have no conflict of interest.



## References

- 355 Adaramola, M. and Krogstad, P.-Å.: Experimental investigation of wake effects on wind turbine performance, *Renewable Energy*, 36, 2078–2086, 2011.
- Allen, J., King, R., and Barter, G.: Wind farm simulation and layout optimization in complex terrain, in: *Journal of Physics: Conference Series*, vol. 1452, p. 012066, IOP Publishing, 2020.
- Antonini, E. G., Romero, D. A., and Amon, C. H.: Optimal design of wind farms in complex terrains using computational fluid dynamics and adjoint methods, *Applied Energy*, 261, 114 426, 2020.
- 360 Antonini, E. G. A., Romero, D. A., and Amon, C. H.: Continuous adjoint formulation for wind farm layout optimization: A 2D implementation, *Applied Energy*, 228, 2333–2345, 2018.
- Asmuth, H., Korb, H., and Ivanell, S.: How Fast is Fast Enough? Industry Perspectives on the Use of Large-eddy Simulation in Wind Energy, in: *Journal of Physics: Conference Series*, vol. 2505, p. 012001, IOP Publishing, 2023.
- 365 Barthelmie, R. J., Hansen, K., Frandsen, S. T., Rathmann, O., Schepers, J. G., Schlez, W., Phillips, J., Rados, K., Zervos, A., Politis, E. S., and Chaviaropoulos, P. K.: Modelling and measuring flow and wind turbine wakes in large wind farms offshore, *Wind Energy*, 12, 431–444, 2009.
- Bartholomew, P., Deskos, G., Frantz, R. A. S., Schuch, F. N., Lamballais, E., and Laizet, S.: Xcompact3D: An open-source framework for solving turbulence problems on a Cartesian mesh, *SoftwareX*, 12, 100 550, 2020.
- 370 Bastankhah, M. and Porté-Agel, F.: A new analytical model for wind-turbine wakes, *Renewable Energy*, 70, 116–123, 2014.
- Bastankhah, M. and Porté-Agel, F.: Experimental and theoretical study of wind turbine wakes in yawed conditions, *Journal of Fluid Mechanics*, 806, 506–541, 2016.
- Bastankhah, M. and Porté-Agel, F.: Wind farm power optimization via yaw angle control: A wind tunnel study, *Journal of Renewable and Sustainable Energy*, 11, 2019.
- 375 Bempedelis, N. and Magri, L.: Bayesian optimization of the layout of wind farms with a high-fidelity surrogate model, arXiv preprint arXiv:2302.01071, 2023.
- Bempedelis, N. and Steiros, K.: Analytical all-induction state model for wind turbine wakes, *Physical Review Fluids*, 7, 034 605, 2022.
- Bempedelis, N., Laizet, S., and Deskos, G.: Turbulent entrainment in finite-length wind farms, *Journal of Fluid Mechanics*, 955, A12, 2023.
- Binois, M. and Wycoff, N.: A survey on high-dimensional Gaussian process modeling with application to Bayesian optimization, *ACM Transactions on Evolutionary Learning and Optimization*, 2, 1–26, 2022.
- 380 Blonigan, P. J. and Wang, Q.: Multiple shooting shadowing for sensitivity analysis of chaotic dynamical systems, *Journal of Computational Physics*, 354, 447–475, 2018.
- Bokharaie, V. S., Bauweraerts, P., and Meyers, J.: Wind-farm layout optimisation using a hybrid Jensen–LES approach, *Wind Energy Science*, 1, 311–325, 2016.
- 385 Campagnolo, F., Petrović, V., Bottasso, C. L., and Croce, A.: Wind tunnel testing of wake control strategies, in: 2016 American Control Conference (ACC), pp. 513–518, IEEE, 2016.
- Chowdhury, S., Zhang, J., Messac, A., and Castillo, L.: Unrestricted wind farm layout optimization (UWFLO): Investigating key factors influencing the maximum power generation, *Renewable Energy*, 38, 16–30, 2012.
- Deskos, G., Laizet, S., and Piggott, M. D.: Turbulence-resolving simulations of wind turbine wakes, *Renewable Energy*, 134, 989–1002, 390 2019.



- Deskos, G., Laizet, S., and Palacios, R.: WInc3D: A novel framework for turbulence-resolving simulations of wind farm wake interactions, *Wind Energy*, 23, 779–794, 2020.
- Duvenaud, D.: Automatic model construction with Gaussian processes, Ph.D. thesis, University of Cambridge, 2014.
- Eriksson, D., Pearce, M., Gardner, J., Turner, R. D., and Poloczek, M.: Scalable global optimization via local Bayesian optimization, *Advances in Neural Information Processing Systems*, 32, 2019.
- 395 Fleming, P., Gebraad, P. M., Lee, S., van Wingerden, J.-W., Johnson, K., Churchfield, M., Michalakes, J., Spalart, P., and Moriarty, P.: Simulation comparison of wake mitigation control strategies for a two-turbine case, *Wind Energy*, 18, 2135–2143, 2015.
- Fleming, P., Annoni, J., Shah, J. J., Wang, L., Ananthan, S., Zhang, Z., Hutchings, K., Wang, P., Chen, W., and Chen, L.: Field test of wake steering at an offshore wind farm, *Wind Energy Science*, 2, 229–239, 2017.
- 400 Fleming, P., King, J., Dykes, K., Simley, E., Roadman, J., Scholbrock, A., Murphy, P., Lundquist, J. K., Moriarty, P., Fleming, K., et al.: Initial results from a field campaign of wake steering applied at a commercial wind farm—Part 1, *Wind Energy Science*, 4, 273–285, 2019.
- Frandsen, S., Barthelmie, R., Pryor, S., Rathmann, O., Larsen, S., Højstrup, J., and Thøgersen, M.: Analytical modelling of wind speed deficit in large offshore wind farms, *Wind Energy*, 9, 39–53, 2006.
- Gebraad, P. M., Teeuwisse, F. W., Van Wingerden, J., Fleming, P. A., Ruben, S. D., Marden, J. R., and Pao, L. Y.: Wind plant power optimization through yaw control using a parametric model for wake effects—a CFD simulation study, *Wind Energy*, 19, 95–114, 2016.
- 405 Goit, J. P. and Meyers, J.: Optimal control of energy extraction in wind-farm boundary layers, *Journal of Fluid Mechanics*, 768, 5–50, 2015.
- González, J., Dai, Z., Hennig, P., and Lawrence, N.: Batch Bayesian optimization via local penalization, in: *Artificial Intelligence and Statistics*, pp. 648–657, PMLR, 2016.
- Gori, F., Laizet, S., and Wynn, A.: Sensitivity analysis of wake steering optimisation for wind farm power maximisation, *Wind Energy Science Discussions*, 2023, 1–33, 2023.
- 410 Grady, S., Hussaini, M., and Abdullah, M. M.: Placement of wind turbines using genetic algorithms, *Renewable Energy*, 30, 259–270, 2005.
- Herbert-Acero, J. F., Probst, O., Réthoré, P.-E., Larsen, G. C., and Castillo-Villar, K. K.: A review of methodological approaches for the design and optimization of wind farms, *Energies*, 7, 6930–7016, 2014.
- Howland, M. F., Lele, S. K., and Dabiri, J. O.: Wind farm power optimization through wake steering, *Proceedings of the National Academy of Sciences*, 116, 14 495–14 500, 2019.
- 415 Huhn, F. and Magri, L.: Gradient-free optimization of chaotic acoustics with reservoir computing, *Physical Review Fluids*, 7, 014 402, 2022.
- Jané-Ippel, C., Bempedelis, N., Palacios, R., and Laizet, S.: High-fidelity simulations of wake-to-wake interaction in an atmospheric boundary layer over a complex terrain, in: *Journal of Physics: Conference Series*, vol. 2505, p. 012033, IOP Publishing, 2023.
- Jensen, N.: A note on wind generator interaction, Tech. Rep. RISO-M-2411, Risoe National Laboratory, Roskilde (Denmark), 1983.
- 420 Katic, I., Højstrup, J., and Jensen, N.: A Simple Model for Cluster Efficiency, *EWEC’86 Proceedings*, 1, 407–410, 1987.
- King, J., Fleming, P., King, R., Martínez-Tossas, L. A., Bay, C. J., Mudafort, R., and Simley, E.: Control-oriented model for secondary effects of wake steering, *Wind Energy Science*, 6, 701–714, 2021.
- King, R. N., Dykes, K., Graf, P., and Hamlington, P. E.: Optimization of wind plant layouts using an adjoint approach, *Wind Energy Science*, 2, 115–131, 2017.
- 425 Kirby, A., Briol, F.-X., Dunstan, T. D., and Nishino, T.: Data-driven modelling of turbine wake interactions and flow resistance in large wind farms, *Wind Energy*, n/a, <https://doi.org/https://doi.org/10.1002/we.2851>.
- Kraft, D.: A software package for sequential quadratic programming, vol. 88 of *Deutsche Forschungs- und Versuchsanstalt für Luft- und Raumfahrt Köln: Forschungsbericht*, Wiss. Berichtswesen d. DFVLR, 28 edn., DFVLR-FB88-28, 1988.



- Kusiak, A. and Song, Z.: Design of wind farm layout for maximum wind energy capture, *Renewable Energy*, 35, 685–694, 2010.
- 430 Laizet, S. and Lamballais, E.: High-order compact schemes for incompressible flows: A simple and efficient method with quasi-spectral accuracy, *Journal of Computational Physics*, 228, 5989 – 6015, 2009.
- Laizet, S. and Li, N.: Incompact3d: A powerful tool to tackle turbulence problems with up to  $\mathcal{O}(10^5)$  computational cores, *International Journal for Numerical Methods in Fluids*, 67, 1735–1757, 2011.
- Liu, D. C. and Nocedal, J.: On the limited memory BFGS method for large scale optimization, *Mathematical Programming*, 45, 503–528, 435 1989.
- Mahfoze, O., Moody, A., Wynn, A., Whalley, R., and Laizet, S.: Reducing the skin-friction drag of a turbulent boundary-layer flow with low-amplitude wall-normal blowing within a Bayesian optimization framework, *Physical Review Fluids*, 4, 094 601, 2019.
- Mosetti, G., Poloni, C., and Diviacco, B.: Optimization of wind turbine positioning in large windfarms by means of a genetic algorithm, *Journal of Wind Engineering and Industrial Aerodynamics*, 51, 105–116, 1994.
- 440 Munters, W. and Meyers, J.: Dynamic strategies for yaw and induction control of wind farms based on large-eddy simulation and optimization, *Energies*, 11, 177, 2018.
- NREL: FLORIS. Version 3.4, GitHub, <https://doi.org/10.5281/zenodo.7942258>, 2023.
- Our World in Data: Share of electricity production from wind, 2022, <https://ourworldindata.org/grapher/share-electricity-wind>, accessed 18 August, 2022.
- 445 O’Connor, J., Diessner, M., Wilson, K., Whalley, R. D., Wynn, A., and Laizet, S.: Optimisation and analysis of streamwise-varying wall-normal blowing in a turbulent boundary layer, *Flow, Turbulence and Combustion*, 110, 993–1021, 2023.
- Peschard, I. and Le Gal, P.: Coupled wakes of cylinders, *Physical Review Letters*, 77, 3122, 1996.
- Porté-Agel, F., Bastankhah, M., and Shamsoddin, S.: Wind-turbine and wind-farm flows: a review, *Boundary-Layer Meteorology*, 174, 1–59, 2020.
- 450 Rasmussen, C. E. and Williams, C. K.: *Gaussian processes for machine learning*, MIT Press, 2006.
- Shakoor, R., Hassan, M. Y., Raheem, A., and Wu, Y.-K.: Wake effect modeling: A review of wind farm layout optimization using Jensen’s model, *Renewable and Sustainable Energy Reviews*, 58, 1048–1059, 2016.
- Simley, E., Fleming, P., Girard, N., Alloin, L., Godefroy, E., and Duc, T.: Results from a wake-steering experiment at a commercial wind plant: investigating the wind speed dependence of wake-steering performance, *Wind Energy Science*, 6, 1427–1453, 2021.
- 455 Smagorinsky, J.: General circulation experiments with the primitive equations: I. The basic experiment, *Monthly Weather Review*, 91, 99–164, 1963.
- Stanley, A. P. and Ning, A.: Massive simplification of the wind farm layout optimization problem, *Wind Energy Science*, 4, 663–676, 2019.
- Steiros, K., Kokmanian, K., Bempedelis, N., and Hultmark, M.: The effect of porosity on the drag of cylinders, *Journal of Fluid Mechanics*, 901, R2, 2020.
- 460 Steiros, K., Bempedelis, N., and Ding, L.: Recirculation regions in wakes with base bleed, *Physical Review Fluids*, 6, 034 608, 2021.
- Steiros, K., Bempedelis, N., and Cicolin, M.: An analytical blockage correction model for high-solidity turbines, *Journal of Fluid Mechanics*, 948, A57, 2022.
- Steiros, K., Obligado, M., and Vassilicos, J. C.: The influence of the vortex shedding instability on the far wake of bluff bodies, Under review, 2023.
- 465 The GPyOpt authors: GPyOpt: A Bayesian Optimization framework in Python, <http://github.com/SheffieldML/GPyOpt>, 2016.



- Thomas, J. J., Annoni, J., Fleming, P. A., and Ning, A.: Comparison of wind farm layout optimization results using a simple wake model and gradient-based optimization to large eddy simulations, in: AIAA Scitech 2019 Forum, p. 0538, 2019.
- Veers, P., Dykes, K., Lantz, E., Barth, S., Bottasso, C. L., Carlson, O., Clifton, A., Green, J., Green, P., Holttinen, H., et al.: Grand challenges in the science of wind energy, *Science*, 366, eaau2027, 2019.
- 470 Veers, P., Dykes, K., Basu, S., Bianchini, A., Clifton, A., Green, P., Holttinen, H., Kitzing, L., Kosovic, B., Lundquist, J. K., et al.: Grand Challenges: wind energy research needs for a global energy transition, *Wind Energy Science*, 7, 2491–2496, 2022.
- Wang, Q.: Forward and adjoint sensitivity computation of chaotic dynamical systems, *Journal of Computational Physics*, 235, 1–13, 2013.
- WindESCo: WindESCo delivers wind industry’s first major wake steering installation, <https://www.windesco.com/blog/windesco-delivers-wind-industrys-first-major-wake-steering-installation>, accessed 23 July, 2023.
- 475 Wu, Y.-T. and Porté-Agel, F.: Modeling turbine wakes and power losses within a wind farm using LES: An application to the Horns Rev offshore wind farm, *Renewable Energy*, 75, 945–955, 2015.
- Zong, H. and Porté-Agel, F.: Experimental investigation and analytical modelling of active yaw control for wind farm power optimization, *Renewable Energy*, 170, 1228–1244, 2021.