

RESEARCH ARTICLE OPEN ACCESS

Bayesian Optimisation of a Two-Turbine Configuration Around a 2D Hill Using Large Eddy Simulations

Christian Jané-Ippel¹  | Nikos Bempedelis²  | Rafael Palacios¹  | Sylvain Laizet¹ 

¹Department of Aeronautics, Imperial College London, London, UK | ²School of Engineering and Materials Science, Queen Mary University of London, London, UK

Correspondence: Christian Jané-Ippel (christian.jane-ippel19@imperial.ac.uk)

Received: 1 February 2024 | **Revised:** 17 June 2024 | **Accepted:** 19 July 2024

Funding: This research was supported by the Natural Environment Research Council, Grant/Award Number: NE/S007415/1, and by the Imperial College London.

Keywords: Bayesian optimisation | complex terrain | large Eddy simulation | layout optimisation | turbulent boundary layer | wake steering

ABSTRACT

When planning a new wind farm, optimising the turbine layout to maximise power output is crucial. Traditional approaches based on analytical wake models often fail in complex terrain. In such areas, wake-to-wake interactions and terrain-induced effects need to be taken into account via high-fidelity optimisation. Building on our prior research on how best to maximise offshore wind farm power production by both micro-siting (layout optimisation) and wake steering (yaw angle optimisation), this study focuses on the optimisation of two onshore wind turbines located around a 2D hill. The proposed approach is based on high-fidelity simulations (LES) to account for the complexity of the flow and on Bayesian optimisation (BO) to systematically and efficiently find the optimal configuration for the two turbines given a set of parameters. The LES are performed with the high-order finite-difference wind farm simulator `WInC3D`, which has been modified for this study to deal with complex terrains and turbines with yaw and tilt. Following the validation of the immersed boundary method (IBM) and wall modelling used to model the complex terrain and of a modified actuator disc model (ADM) to account for tilted turbines, two BO campaigns are presented in this study, with different sets of parameters to optimise. The main flow features and some statistics for each optimum case are discussed using the flow fields with and without the optimised configuration of the upstream turbine to analyse its impact on the downstream turbine and on the overall power of the system. For the first optimisation, the design variables include the wind turbines' streamwise locations and their hub heights (4 parameters). This optimisation study benefits from an elevated upstream turbine hub, leading to a significant power increase in the downstream turbine, thanks to an acceleration of the flow in front of the hill. For the second optimisation, the design space is modified based on the data from the first optimisation, by replacing the turbine's hub height with the possibility of changing the tilt angle of the upstream turbine (4 parameters). It is found that considerable power gains can be obtained while maintaining a modest upstream turbine hub height, by introducing a positive tilt angle in the upstream turbine. Overall, it is found that by setting up the turbines in ways that exploit the speed-up of the flow around the hill, both optimised layouts achieve considerable enhancements in power density over a reference configuration.

1 | Introduction

As the global demand for sustainable and renewable energy sources intensifies, wind energy has emerged as a pivotal

solution. Its potential to mitigate greenhouse gas emissions and reduce the impact of fossil fuels on climate change has positioned it at the forefront of renewable energy strategies. A notable trend in the pursuit of increased wind energy contributions is the

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2024 The Author(s). *Wind Energy* published by John Wiley & Sons Ltd.

shift towards larger wind turbines and extensive farms. While this delivers increased energy capture, it simultaneously underscores the need for a better understanding of the dynamic interactions within the atmospheric boundary layer (ABL) and the wind farm. These interactions can significantly influence both the efficiency of power generation and the longevity of the turbines [1, 2]. In the case of onshore wind farms, the presence of hills or nearby urban areas introduces additional complexities. These can induce speed-ups, flow separation and higher levels of turbulence, all of which significantly affect power output, turbine loads and wake recovery. Hence, appropriately accounting for the effects of topography is imperative for designing wind farms in complex terrain [3].

Various methods have been devised to model the influence of topography on wind turbine wakes. A prevalent method in the industry is the superposition technique, which involves superimposing the turbine wake's velocity deficit from flat terrain onto the flow over the topography without turbines [4–7]. This technique is suitable for low to moderately complex terrains, but its applicability is challenged in landscapes of higher complexity [8, 9] as it can fail in the representation of pressure gradients, variable elevation of wake trajectories and flow separation [3]. Alternatively, steady Reynolds-Averaged Navier-Stokes (RANS) have been employed over complex terrain [8, 10, 11], offering a method to consider more complex phenomena such as blockage or pressure effects. However, RANS falls short in accurately capturing the (unsteady) wake dynamics, flow features in cases of highly complex topography (e.g., recirculation regions) and the dynamic interactions with the atmospheric flow. To address this, an approach that increasingly gains interest owing to advances in numerical methods and computing architectures is the use of large Eddy simulations (LES) [12]. In turbulence-resolving simulations, such as LES, the aforementioned issues are mitigated, as the dynamic interactions between ABL, complex terrain and turbine wakes are taken into consideration. Several studies have demonstrated good agreement between LES and experimental data [13–17], confirming its efficacy for modelling wind farm flows over complex terrain. However, this improvement in accuracy comes at a higher computational cost. Furthermore, LES have been used to study wind turbine wakes in real complex terrain sites [18–21] and coupled with the weather research and forecasting model (WRF) [22, 23].

Among the best studied examples of complex terrain are the Gaussian and cosine squared hills, which are often used as reference models for real-world topography. Liu and Stevens [15] and Zhang et al. [16] validated their LES methods that use an IBM for the representation of the topography against the experimental results of Cao and Tamura [24] for the flow of a turbulent boundary layer (TBL) over a 2D hill. Their studies provided insights into the benefits and downsides of placing wind turbines around the hill. In particular, Liu and Stevens [15] found that turbines significantly taller than the hill benefit from the accelerated flow over the terrain, with the benefits depending on their placement in relation to the hill. Moreover, they showed that turbines on the windward side matched predictions of the superimposition method while those on the leeward side did not. They also investigated the trade-offs of wind farm power when leaving a gap in front or behind the hill and concluded that it predominantly affected downstream turbines, with the

hill's recirculation zone being the overriding flow feature. On the other hand, Zhang et al. [16] examined individual wind turbine wakes over 2D hills of varying slopes, with and without flow separation. Their findings suggest that in the presence of turbine wakes, the streamwise velocity in the wake of hills rises due to mass flux conservation. They suggest placing the turbine at the hilltop as it was deemed ideal for harvesting more energy and reducing the turbine dynamic loading owing to reduced levels of turbulence. Conversely, turbines placed in the wake of steep hills experienced a decrease in performance due to reverse flow and high levels of turbulence intensity.

Wind farms sited in complex terrain suffer from power losses due to turbine-wake interactions and terrain-induced flow features. Wind farm layout optimisation attempts to mitigate these losses and maximise the farm power output. The objective function is typically the total power output, with turbine locations as design variables. Such optimisations generally employ low-fidelity flow solvers owing to the large number of function evaluations required and the long running times and high computational cost of high-fidelity solvers, such as LES. Examples include superposition methods with wake models [7, 25–27] or medium-fidelity tools such as RANS simulations [28, 29]. Recently, Bempedelis et al. [30] demonstrated the feasibility of using high-fidelity solvers for wind farm power maximisation, using a method based on LES and a Bayesian framework. They pointed out the potential benefits of using high-fidelity tools in scenarios of increased flow complexity, such as in wake steering. Given the prevalence of complex flow phenomena found in complex terrain flows, there exists an opportunity to explore layout optimisation and wake steering using high-fidelity tools. In scenarios where function evaluations are costly, such as with LES, BO has been identified as an effective and robust alternative, surpassing various other derivative-free optimisation strategies like particle swarm optimisation and explorative gradient methods in achieving optimal outcomes with fewer function evaluations [31]. Moreover, while combinations of LES and adjoint methods have been explored for wind farm control [32, 33], they often struggle with the complexities of turbulent flows. For instance, Talnikar and Wang [34] have documented issues where adjoint solutions diverge when calculating long-term averages from chaotic flow features, leading to significant errors in gradient estimations.

In this study, we aim to maximise the power of wind farms in complex terrain by combining high-fidelity LES with Bayesian optimisation (BO) to search for optimal configurations for a two-turbine set-up around a steep 2D hill. Section 2 presents a detailed validation of the LES and the methodology used for the two BOs performed in this work. In Section 3, the two optimal turbine configurations obtained are presented and compared against a conventional layout. Finally, conclusions are drawn in Section 4.

2 | Methodology

This section first outlines the key features of the LES employed in this work to evaluate the power output of the different turbine configurations. We then demonstrate the new capabilities of our wind farm simulator `Winc3D` in resolving the flow around the hill and in predicting power, by making comparisons with previously published experimental and numerical data. After the

validation, we introduce the objective function and design space for maximising power in the two BOs, followed by an overview of the employed optimisation algorithm. The section concludes with a comment on the computational cost of the optimisations.

2.1 | Large Eddy Simulation

The simulation set-up used in the optimisations has been adopted from the validation of the LES with the wind tunnel experiments of Cao and Tamura [24]. The complex terrain investigated is the constant section hill (in the spanwise direction) defined by a cosine-squared function, with a maximum height of $H = 0.04$ m and a half-width of $L = 0.1$ m. The TBL inflow in the experiments is characterised by a TBL height of $\delta = 0.25$ m, a friction velocity of $u_* = 0.1926$ m/s, and a roughness length of $y_0 = 0.004$ mm.

The computational domain has a size of $L_x \times L_y \times L_z = (5 \times 1 \times 1)\delta$. The simulations require a total non-dimensional time of $\hat{T} = 20$ to fully converge the second-order statistics after discarding $\hat{T} = 2$ of initial transient. The non-dimensional time units ($\hat{T}$) are defined as $\hat{T} = \frac{u_* t}{\delta}$, where t is the physical time in seconds. The validation results are presented for two different discretisation levels, hereafter referred to as low- and high-resolution simulations. The low-resolution grid ($n_x \times n_y \times n_z = 193 \times 97 \times 64$) has around 1M grid nodes and the non-dimensional time step used is $D\hat{T} \approx 7.7 \times 10^{-5}$. The high-resolution simulations double the resolution in each direction ($n_x \times n_y \times n_z = 385 \times 193 \times 128$), which results in around 9.5M grid nodes, and the time step used is $D\hat{T} \approx 3.85 \times 10^{-5}$. The high-resolution grid has been employed in the optimisations, due to its better agreement with the experimental data (see Section 2.1.2). Simulations with wind turbines use an ADM (Section 2.1.3), with the turbines characterised by a diameter of $D = H$ and a local thrust coefficient of $C_T = 4/3$, similar to the studies of Liu and Stevens [15] and Zhang et al. [16]. Figure 1 shows a schematic view of the computational domain with the hill and a wind turbine at the hilltop, $x/D = 0.0$, and the boundary conditions used.

A slip wall boundary condition is enforced on the top boundary, while a no-slip wall condition with a wall stress model is used on the bottom boundary. Periodic boundary conditions are applied in the spanwise planes. An inlet boundary condition is used to provide the turbulent inflow, with the inflow planes extracted

from a precursor simulation of a neutral boundary layer. The outlet boundary condition is applied with one-dimensional convection equations. The convective velocity is calculated as the time and grid-scaled average of the maximum and minimum streamwise velocities at the last interior grid point.

We perform a precursor simulation of the TBL in the experiments of Cao and Tamura [24] to model the effect of the ABL on the flow around the hill. We use the shifted periodic boundary condition of Munters et al. [35] as it allows shorter streamwise domain lengths while avoiding the locking of large-scale turbulent structures that lead to statistical inhomogeneities that can persist for long time in the simulations. A separate precursor simulation of the TBL was run for each spatial resolution with a computational domain size of $L_x \times L_y \times L_z = (3 \times 1 \times 1)\delta$. The statistics were collected for $\hat{T} = 20$. As shown in Figure 2a, the horizontally and time-averaged streamwise velocity profiles are compared against the time-averaged profiles in the experimental study of Cao and Tamura [24] and the log-law profile $\frac{\langle u_x \rangle}{u_*} = \frac{1}{\kappa} \ln\left(\frac{y}{y_0}\right)$. $\kappa = 0.4$ is the von Kármán constant. The flow is driven by a constant pressure gradient u_*^2/δ , with the total stresses in a horizontal stationary TBL following $1 - \frac{y}{L_y}$. Due to the no-slip boundary condition at the wall and the use of high-order schemes, the log-law profile is recovered after the first 3 to 5 nodes off the wall. It should be noted that it is not possible to avoid the log-layer mismatch when using Mason and Thompson modelling at the wall. The log-layer mismatch does not have a significant impact in the present work as the turbine hub-heights are outside the log-layer. Figure 2b shows the horizontally and time-averaged resolved stresses $\langle u'_x u'_y \rangle$ obtained in the precursor TBL simulations. Both low- and high-resolution simulations are in good agreement with the experimental data of Cao and Tamura [24] and follow the total theoretical stress profile almost all the way down to the wall. As expected, the high-resolution simulation resolves better the turbulence near the wall, but only marginally.

2.1.1 | Numerical Methods

The wind farm simulator WINC3D [36] is an open-source code that is part of Xcompact3D, a suite of fluid flow solvers dedicated to the study of turbulent flows on supercomputers [37]. For the present study, the governing equations are the unsteady,




FIGURE 1 | Computational domain representation of the validation set-up, featuring the constant section hill and an actuator disc at the hilltop, along with the boundary conditions.




FIGURE 2 | Horizontally and time-averaged streamwise velocity (a) and resolved stress profiles (b) of the simulated TBL compared against analytical profiles for a turbulent boundary layer and experimental data of Cao and Tamura [24].

incompressible, filtered Navier-Stokes (N-S) equations, based on an explicit LES formulation:

$$\frac{\partial \tilde{u}_i}{\partial t} + \frac{1}{2} \left(\tilde{u}_j \frac{\partial \tilde{u}_i}{\partial x_j} + \frac{\partial \tilde{u}_i \tilde{u}_j}{\partial x_j} \right) = -\frac{1}{\rho} \frac{\partial \tilde{p}}{\partial x_i} - \frac{\partial \tau_{ij}}{\partial x_j} + \nu \frac{\partial^2 \tilde{u}_i}{\partial x_j \partial x_j} + \frac{F_i}{\rho}, \quad (1)$$

$$\frac{\partial \tilde{u}_i}{\partial x_i} = 0, \quad (2)$$

where $\tilde{u}_i = (u_x, u_y, u_z)$ and $\tilde{p}$ are the spatially filtered components of velocity and pressure fields, $\rho = 1.2 \text{ kg/m}^3$ is the fluid density, τ_{ij} the subfilter-scale stresses and $\nu = 1.5 \times 10^{-5} \text{ m}^2/\text{s}$ the kinematic viscosity. F_i is the body forcing used to model the complex terrain with an IBM and also to model the wind turbines. The convective terms are written in the skew-symmetric form, as it allows for the reduction of aliasing errors while remaining energy-conserving [38]. Subfilter stresses are computed using the standard Smagorinsky model [39], defined as $\tau_{ij} = -2(C_S \Delta)^2 |\tilde{S}| \tilde{S}_{ij}$, where C_S is the so-called Smagorinsky constant, $\Delta = \sqrt[3]{\Delta x \Delta y \Delta z}$ is the geometrically averaged grid size and $\tilde{S}_{ij}$ and $|\tilde{S}|$ are the strain-rate tensor and its magnitude, respectively. The Smagorinsky constant remains constant throughout the domain but is reduced near the wall to avoid excessive dissipation at the surface [40]. To address this, we apply the wall-damping function of Mason and Thomson [41], which adjusts the Smagorinsky constant as $C_S = \left(C_0^n + \left\{ \kappa \left(\frac{y}{\Delta} + \frac{y_0}{\Delta} \right) \right\}^{-n} \right)^{-1/n}$. Here, y_0 is the boundary layer's roughness height, C_0 is the Smagorinsky constant far from the wall, with $C_0 = 0.14$, and n is a parameter determining the growth from the wall, set at $n = 3$.

The framework uses a Cartesian grid and compact finite-difference schemes for differentiating, filtering and interpolating. The spatial discretisation uses sixth-order schemes. Time advancement is performed using an explicit third-order Adams-Bashforth scheme. The high-order compact finite difference schemes employed guarantee a 'spectral-like' accuracy [42, 43] while allowing the use of non-periodic boundary conditions. The use of high-order schemes, with their quasi-spectral accuracy, significantly enhances the efficient resolution of small-scale turbulence for a given resolution, offering a beneficial trade-off between the computational cost and accuracy [42]. To avoid numerical oscillations, we apply a spatial filter along all directions to the velocity field at every time

step using a discrete filter following the tridiagonal filtering of Motheau and Abraham [44]. The discrete low-pass filter coefficients are set to satisfy the required sixth-order accuracy and a user-defined coefficient $\alpha_{\text{filter}} = 0.49$ is used to filter at close to the grid scale. The incompressibility of the velocity field is ensured by solving the pressure Poisson equation in spectral space using three-dimensional fast Fourier transforms (FFTs) and the concept of the modified wave number [43]. The pressure grid is half staggered from the velocity grid to avoid spurious pressure perturbations [42]. The simplicity of the Cartesian grid is exploited with a powerful 2D domain decomposition strategy based on standardised MPI that gives the code its excellent strong and weak scaling properties [37, 45]. The computational domain is divided into 'pencils', each handled by an MPI process. The derivatives, interpolations, and one-dimensional FFTs in the x, y and z directions are performed within the X, Y and Z pencils, allowing an efficient computation. Each operation is performed sequentially in one direction, with global transpositions enabling to switch from one pencil to another.

2.1.2 | IBM and Wall Modelling for Complex Terrain

The IBM of Gautier et al. [46] is utilised to enforce a no-slip boundary condition on the terrain surface. The IBM first identifies the nodes that fall within the solid domain, which is where the forcing term F_i is added in the N-S equations (1). The method uses Lagrangian polynomials to reconstruct a non-zero velocity field inside the solid body, while enforcing the no-slip condition at the boundaries. This approach promotes stability, eliminates discontinuities at the boundaries of the immersed object, and avoids staircase effects in the representation of the terrain surface. This IBM has been used in various studies involving both wall-resolved LES and direct numerical simulations (DNS) [47, 48], as well as in conjunction with a wall model in Jané-Ippel et al. [49]. A wall model is used due to the high-Reynolds number flow in the TBL and the relatively coarse grid near the bottom wall. The Reynolds number based on the TBL thickness and free stream velocity is $Re = \frac{U_\infty \delta}{\nu} \approx 10^5$. The wall model enforces the desired log-law profile. The current simulations use a no-slip condition at the bottom wall, and, similar to the methodology of Jaegle et al. [50], the wall model is implemented by adjusting the sub-grid stresses at the immediate off-wall grid node. The adjusted stresses, $\tau_{w,i}$, are computed assuming an instantaneous log-law

profile for a rough wall, given by $\tau_w = - \left[\frac{\kappa U_1}{\ln(y_1/y_0)} \right]^2$. Here, y_1 is the local height from the wall where the velocity $U_1 = \sqrt{u_x^2 + u_z^2}$ is sampled within the fluid domain. Good results have been obtained by modelling only the dominant stresses (τ_{xy} and τ_{yz}) while neglecting the others (τ_{xx} , τ_{xz} , τ_{yy} , and τ_{zz}), which is a suitable approximation for smooth shapes such as the hill under consideration. Both the IBM and wall modelling are key tools for representing the flow over complex terrain. Our methodology is validated against the experimental data of Cao and Tamura [24] for the same set-up using the previously defined low-resolution ($n_x \times n_y \times n_z = 193 \times 97 \times 64$) and high-resolution ($n_x \times n_y \times n_z = 385 \times 193 \times 128$) meshes.

Figure 3a shows the time-averaged streamwise velocity $\langle u_x \rangle$ colour map obtained with the high-resolution simulation. Additionally, it displays the time-averaged streamwise velocity profiles from the experimental data and the two simulations with WInc3D, all normalised by the free stream velocity U_∞ at $L_z = \delta/2$. The contour of zero velocity in the streamwise direction, highlighting the edge of reverse flow, is demarcated by the white line. Both grid resolutions align closely with the experimental data across all streamwise locations, accurately capturing the speed-up in front of the hill, separation post-crest, and reattachment at around $x/H = 5.54$. In the experimental campaign, the reattachment is reported to be around $x/H = 5.4$.

Figure 3b shows the time-averaged streamwise turbulence intensity σ_{u_x} colour map, normalised by U_∞ , from the high-resolution simulation. Profiles from both resolutions, set against the experiments and normalised by $U_\infty/5$, are also shown. Unlike

the streamwise velocity profiles, differences may be observed between the two grid resolutions. The high-resolution simulation generally concurs with the experimental data, whereas the low-resolution simulation demonstrates some discrepancies in the near wake of the hill. Nonetheless, predictions from both simulations realign downstream of $x/H \approx 5.0$.

To provide a quantitative assessment of the validation process illustrated in Figure 3, Table 1 lists the maximum percentage errors in mean flow predictions compared to experimental data for each x/H location profiled. The errors for the streamwise velocity are computed as $e_{u_x} = \frac{|u_{x,exp} - u_{x,sim}|}{U_\infty}$ and for the streamwise turbulence intensity as $e_{\sigma_{u_x}} = \frac{|\sigma_{u_x,exp} - \sigma_{u_x,sim}|}{U_\infty}$. Both resolutions display maximum errors of the same order of magnitude, approximately 5%–16% for velocity and 2%–9% for the variance. As also observed in Figure 3, there is a better match with the experimental data for the higher resolution simulations, particularly at $x/H = 0$ and $x/H = 2.5$.

Employing the IBM and a wall model with the high-resolution grid has demonstrated a robust capability in accurately predicting the flow around the hill, as evidenced by both velocity and turbulence statistics. Given this validation, the high-resolution grid is selected for use in the optimisations.

2.1.3 | Actuator Disc Model (ADM) for Yawed and Tilted Turbines

The ADM portrays the wind turbine as a non-rotating porous disc exerting a drag force within its swept area. The thrust force




FIGURE 3 | Comparison between simulation results from WInc3D and experimental data from Cao and Tamura [24] for the normalised time-averaged streamwise velocity (a) and the normalised time-averaged streamwise turbulence intensity (b). The white line in (a) shows the edge of reverse flow.

TABLE 1 | Maximum error of mean flow profiles shown in Figure 3.

Resolution / x/H location	-2.5	-1.25	0	1.25	2.5	3.75	5	6.25	7.5
e_{u_x} (%)									
Low-resolution	10.07	4.62	10.60	11.02	14.01	16.67	8.36	8.98	6.62
High-resolution	10.47	5.34	7.75	11.51	8.25	14.32	6.16	9.80	6.49
$e_{\sigma_{u_x}}$ (%)									
Low-resolution	5.41	3.05	2.45	8.80	9.00	3.25	4.20	3.35	3.26
High-resolution	5.75	1.77	1.75	6.23	5.72	4.50	5.63	3.67	3.02

F_{wt} imposed on the flow is computed based on the local velocity [40, 51] and is given by

$$F_{wt} = -\frac{1}{2}\rho A_D C'_T u_d^2, \quad (3)$$

where ρ denotes the fluid density, A_D the rotor swept area, and C'_T the local thrust coefficient as defined in Calaf et al. [40]. The disc velocity u_d is the spatially averaged velocity through the rotor, time-filtered using a one-sided exponential filter. We employ a relaxation time of $\hat{T}_{relax} = 0.27$, following the approach described by Calaf et al. [40]. The power is subsequently evaluated as $P_{wt} = -F_{wt} u_d$.

The thrust force is distributed across the rotor-swept area using super-Gaussian functions, similar to the approach of King et al. [52]. This allows modelling yawed and tilted turbines (i.e., where the rotors are not aligned with the vertical grid planes) while simultaneously ensuring a sharp representation of the disc and avoiding instabilities and singularities in the vicinity of the rotor. The super-Gaussian kernel, denoted by Γ , is expressed as

$$\Gamma(\delta_n, \delta_r) = \exp\left(-\left(\frac{\delta_n}{t_D/2}\right)^{\alpha_n} - \left(\frac{\delta_r}{D/2}\right)^{\alpha_r}\right), \quad (4)$$

where δ_n and δ_r represent the distances from the turbine centre in the normal and radial direction of the disc plane, respectively, α_n and α_r are the exponents that control the sharpness of the Gaussian profiles across the normal and radial directions, t_D is the disc thickness and D is the rotor diameter. The distributed thrust force is normalised with the volumetric integral of $\Gamma(\delta_n, \delta_r)$. The exponents used are $\alpha_n = 2$ and $\alpha_r = 8$, validated against the experimental data of Chamorro et al. [53]. The exponent α_n ensures a smooth distribution of forces along the disc thickness, while α_r ensures a sharp distribution across the disc edge (see also King et al. [52]).

The proposed ADM is validated against the LES of Heck et al. [54], where the power of a single turbine in uniform flow at different yaw angles was computed by rotating the inflow angle and fixing the disc to the Cartesian grid. The currently proposed method is tested by keeping the inflow angle fixed and yawing the turbine. Various grid resolutions have been assessed to determine the minimum grid requirements to represent a grid-misaligned disc. These are characterised by the number of grid nodes across the diameter $n_D = D/\Delta y$ and

the cell aspect ratio in the streamwise direction $AR = \Delta x/\Delta y$. Figure 4a shows the computed power for a yawed wind turbine with $C'_T = 4/3$ normalised by the power production of a turbine at zero yaw ($\gamma = 0^\circ$), denoted by $\dot{P}$, for various resolutions and the reference of Heck et al. [54]. It can be seen that WInc3D accurately predicts the power as the yaw angle is increased for all tested resolutions except when using few points across the diameter ($n_D = 6.4$) and a high aspect ratio ($AR = 2.5$) together. The predictions of two low-order models are also plotted for reference. The resolution used in the optimisations (see Section 2.1) corresponds to $n_D = 30$ and $AR = 2.5$, which is sufficient for predicting the power of yawed/tilted turbines.

The power prediction capabilities of the proposed ADM are also examined using the complex terrain simulation set-up and adding a single wind turbine at different streamwise locations with a hub height of $h = D$. Figure 4b shows the normalised power of this single wind turbine when placed at different streamwise positions, using the low and high-resolution grids, together with the corresponding 99% confidence intervals (CI). The data is in good agreement with results from the literature [15, 16]. The two levels of grid resolution give very similar results in terms of power output. Only minor discrepancies of around 3.3% and 8.4% are observed at $x/D = 1.25$ and 5, respectively.

The validation of the LES approach against experimental and numerical data, encompassing flow dynamics around the hill, power output of turbines in this complex terrain, and performance of yawed and tilted discs, sets a solid foundation for undertaking the BOs.

2.2 | Optimisation

2.2.1 | Design Variables and Objective Function

The two optimisations performed, henceforth referred to as BO1 and BO2, have four design variables each. The design variables for BO1 include the wind turbines' streamwise locations, denoted as x_1 and x_2 , and their hub heights, represented by h_1 and h_2 . These variables are normalised by the turbines' diameter. Specifically, the streamwise location x_1 spans the top of the hill ($x_1 \in [-1.25, 1.25]$), while x_2 spans the leeward side of the hill ($x_2 \in [0.75, 5.00]$). These choices allow for a thorough exploration of the hill's influence on turbine placement and its corresponding performance. Additionally, they aim to exploit




FIGURE 4 | (a) Power output of a yawed wind turbine normalised by the power output of a turbine at zero yaw angle. Comparison between different resolutions with *WInc3D* and the LES of Heck et al. [54]. The predictions of the low-order $\cos(\gamma)^2$ and $\cos(\gamma)^3$ models are also shown for reference. (b) Normalised power output of a wind turbine placed at various positions around the hill as computed by *WInc3D*, Liu and Stevens [15], and Zhang et al. [16].

the leeward region of the hill, which is not usually explored. Regarding the hub heights, ranges of $h_1, h_2 \in [0.70, 1.50]$ are selected to ensure a diverse set of turbine configurations. To ensure the turbines operate safely, a constraint is imposed in both optimisations to always maintain a minimum distance of around one turbine diameter between their centres ($x_2 - x_1 \geq 1.05$).

In BO2, the design space is modified upon the learnings from BO1 and to explore the possible benefits of wake steering. To this end, the upstream turbine's hub height is fixed to $h_1 = 1.0$, removing it from the design variables. We incorporate new design variables which can substitute the hub height optimisation with a wake steering strategy. Given the current problem setup using periodic boundary conditions in the spanwise direction, the tilt angle is considered as a means of wake steering instead of the more commonly used yaw. This approach allows us to directly influence the mid-span plane as in BO1. The tilt angle, denoted as α_1 , is optimised, allowing it to vary between -20° and $+20^\circ$ ($\alpha_1 \in [-20^\circ, +20.0^\circ]$). This range choice is similar to that used by previous optimisation studies [55, 56], which have explored tilt angles at the wind farm level on flat terrain. A negative tilt angle implies potentially aligning the turbine more closely with the inclined mean flow typically observed at the flow in the hilltop, thereby enhancing the power production of the upstream turbine. High values of negative tilt, such as -20° , are more suitable for downwind turbines, as they avoid the risk of tower strikes often associated with such configurations in conventional turbines [57]. The turbine wake is expected to steer downwards for the negative tilt. Conversely, a positive tilt angle steers the wake upwards, which may be strategically used to mitigate adverse effects on downstream turbines while optimising the airflow capture by the upstream turbine. This dual approach allows the optimisation to evaluate whether steering the wake away from downstream turbines (positive tilt) or improving alignment with the local inflow at the hilltop (negative tilt) results in better overall system performance.

A schematic of the hill and the design variables used in BO1 and BO2 is illustrated in Figure 5a,b, respectively, showcasing wind turbines at positions $x_1 = 0.0$, $h_1 = 1.0$, $\alpha_1 = 0^\circ$, $x_2 = 3.75$, and $h_2 = 1.0$.

The objective function for both BO1 and BO2 optimisations is the aggregate power output of the turbine ensemble, defined as the time-averaged power produced by the pair of turbines. The power estimates come from the LES following the methodology presented in Section 2.1. The power output of the turbines, represented as $\hat{P}$, is normalised by the average power produced by a single turbine under similar inflow conditions on a flat terrain (i.e., without the hill, in isolation). To ascertain a short but reliable time window for statistical averaging of the power, an additional investigation is performed utilising two reference simulations: one consisting of a wind turbine on top of the hill ($x/D = 0$) and the other placing the wind turbine in the hill-wake ($x/D = 5.0$). The chosen cut-off for the simulation time is dictated by achieving a variability of less than 1% in the time-averaged power for a minimum of 6 non-dimensional time units and convergence in the CI of the averaged power signal. This threshold aims to ensure statistical convergence in the context of the turbulence-induced chaotic flow behaviour.

Figure 6 presents the power signals for the two indicative scenarios, including the cumulative and total averages and 99% CI. The confidence intervals are computed following the methodology detailed in Bergmann et al. [58]. This involves utilising a correlation function to establish confidence intervals that account for the inherent temporal correlation of the Navier-Stokes equations from the LES. The plots confirm that from $\hat{T} = 10$ onward, the cumulative power adheres to the established variability criterion, and that from $\hat{T} = 16$ onward, the confidence intervals converge, showing only minimal reductions.

2.2.2 | BO

To search for the global maximum of power production for the two-turbine setup, we use a gradient-free BO strategy which makes use of surrogate models, in particular, Gaussian process (GP) models. This approach is particularly well suited for optimisation tasks with expensive objective function evaluations [30, 59, 60], providing a balance between exploration and exploitation of the design space. Our implementation of BO is based on Python's GPyOpt library [61].




FIGURE 5 | Schematic of the constant section hill with two wind turbines placed around it showing the design variables for optimisations BO1 (a) and BO2 (b).




FIGURE 6 | Power signals with total and cumulative averages and 99% CI for (a) wind turbine placed at the hilltop ($x/D = 0$) and (b) wind turbine immersed in the hill's wake ($x/D = 5.0$). The grey region indicates the initial transient discarded.

The goal of BO is to determine the optimal set of design variables, denoted $\mathbf{x}^*$, that yields the maximum value of an objective function $f(\mathbf{x})$ within a specified design space $\mathcal{X}$. Mathematically, this is expressed as

$$\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}), \quad (5)$$

where $\mathbf{x}$ is a vector representing a specific point in the design space. For both optimisations, the objective function is defined as $f(\mathbf{x}) = \hat{P}_{\text{Up}}(\mathbf{x}) + \hat{P}_{\text{Down}}(\mathbf{x})$, representing the aggregated normalised power of the upstream and downstream turbine. In BO1, the design space $\mathcal{X}_{\text{BO1}}$ comprises $\{x_1, h_1, x_2, h_2\}$ and for BO2, $\mathcal{X}_{\text{BO2}}$ consists of $\{x_1, \alpha_1, x_2, h_2\}$, with the bounds specified in Section 2.2.1. For both optimisations, the spacing constraint is expressed as $x_2 - x_1 \geq 1.05$.

BO begins with a design of experiments (DOE), which, in this work, consists of a set of initial design points near randomly generated through Latin hypercube sampling (LHS) to ensure a diverse and representative coverage of the design space. For both BO1 and BO2, it consists of 5 samples per design variable, resulting in 20 samples in total. These initial sample points are evaluated with LES, based on which the first instance of the GP surrogate model is constructed. The size of the initial sampling used has proven to be enough for convergence within 36 (BO1) or 20 (BO2) additional samples (see Section 2.2.3).

A GP model is a non-parametric model defined by its mean function $\mu(\mathbf{x})$ and its covariance function or kernel $k(\mathbf{x}, \mathbf{x}')$. The kernel measures the similarity or covariance between two points $\mathbf{x}$ and $\mathbf{x}'$ in the input space. Given the precision and control over the simulations to evaluate the objective function, and neglecting

the uncertainty introduced by some modelling approximations (e.g., ADM, tower/nacelle), we consider an observation noise-free GP formulation. Given training data $D = \{\mathbf{X}, \mathbf{y}\}$, where $\mathbf{X}$ represents the input vectors and $\mathbf{y}$ the corresponding observed values, the mean (μ) and variance (σ^2) of the GP posterior at a new point $\mathbf{x}_*$ are given by

$$\mu(\mathbf{x}_*) = \mathbf{k}_*^T \mathbf{K}^{-1} \mathbf{y} \quad (6)$$

$$\sigma^2(\mathbf{x}_*) = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^T \mathbf{K}^{-1} \mathbf{k}_* \quad (7)$$

where $\mathbf{k}_*$ is the estimated covariance vector between the training points and the unseen point $\mathbf{x}_*$, and $\mathbf{K}$ is the covariance matrix of the training data, with elements $K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$. For this work, the GP surrogate model uses the Matérn 5/2 kernel, which has proven to be effective for similar physical problems [59]. The Matérn 5/2 kernel is given by

$$k(\mathbf{x}, \mathbf{x}') = \sigma^2 \left(1 + \sqrt{5}r + \frac{5r^2}{3} \right) \exp(-\sqrt{5}r), \quad (8)$$

where

$$r = \sqrt{\sum_{i=1}^n \left(\frac{x_i - x'_i}{\ell_i} \right)^2}. \quad (9)$$

n is the number of dimensions of the input space, and ℓ_i are the characteristic length scales associated with each dimension. The kernel's hyperparameters, i.e., the variance and lengthscales, are determined by maximising the marginal likelihood using

a gradient-based method to iteratively refine these parameters. Each design variable in our model is associated with a unique lengthscale, allowing the GP to capture the influence of each variable on the output.

Once the GP model is built, the acquisition function guides the choice of subsequent evaluation points. Our acquisition function of choice is the lower confidence bound (LCB) function, which is defined as

$$\text{LCB}(\mathbf{x}) = \mu(\mathbf{x}) - w(\sigma(\mathbf{x}) + \epsilon), \quad (10)$$

where $\mu(\mathbf{x})$ and $\sigma(\mathbf{x})$ represent the GP's posterior mean and standard deviation at $\mathbf{x}$, respectively. The parameter w controls the trade-off between exploration and exploitation by weighting the emphasis placed on uncertainty, and ϵ is the jitter term. In our optimisation, these parameters are set to $w = 2$ and $\epsilon = 0.01$. A batch $\mathbf{B}_t$ of design points for simultaneous evaluation is subsequently computed as

$$\mathbf{B}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} \text{LCB}(\mathbf{x} | D_t), \quad (11)$$

where D_t encompasses all data up to iteration t . Given the considerable computational demands of the evaluations, our chosen batch size is four. We employ the local penalisation method of González et al. [62] to ensure non-overlapping and systematic exploration of the design space. The optimisation progresses iteratively, assimilating the new data in the training data set $\mathcal{D}$, updating the surrogate model and the acquisition function and selecting new points to be evaluated.

2.2.3 | Convergence and Cost

We considered a budget of 100 evaluations for each optimisation. Convergence is assessed by determining the cumulative number of simulations at which the improvement of the maximum value of power found falls below a certain threshold, in this case set to 1%, similar to the criterion selected for the end of the time-averaging period, as discussed in Section 2.2.1.

Figure 7 illustrates the progression of BO1 and BO2, showing the power deficit ΔP with respect to the global optimum found for all the simulations of the optimisation (in blue) and the optimal values found up to that point as the optimisation progressed (in black). A horizontal red line has been added marking the

-1% deficit. For BO1, the optimal power converged after 56 simulations. Similarly, BO2 reached this convergence threshold after 40 simulations. Beyond this point, the lower performance points along the blue line, explore and refine the model's understanding of less optimal regions. This exploration enhances the robustness of the GP model, particularly by confirming the boundaries of lower performance and better defining the gradient towards the optimum. Moreover, the clustering of points near the optimum shown in the blue line, increases confidence in the region around the optimum.

Each evaluation based on one LES leveraged 2048 cores from the ARCHER2 UK National Supercomputing Service, each with dual AMD EPYC 7742 64-core processors at 2.25 GHz. The wall clock time per LES is approximately 3.6 h. Across 100 evaluations, the cumulative computational cost is around 737,280 CPUh. Each optimisation requires approximately 5 days of wall clock time to complete, accounting for parallel initial sampling and the use of batches of four evaluations at each surrogate model update. Note that the computational demands and time associated with the BO algorithm are marginal compared to the LES evaluations.

3 | Results

This section presents the findings from BOs 1 and 2 (BO1 and BO2). The results compare the best designs from both optimisations with a reference scenario, henceforth referred to as the Reference case, where a wind turbine is positioned on top of the hill, with a second turbine placed at a distance of 5D downstream of the first turbine. An overview of the optimum design variables, highlighted in bold, and the corresponding values of the objective function for both BO1 and BO2, alongside those of the Reference case, are provided in Table 2.

Figure 8 illustrates the turbine placement in a flow visualisation of the two optimal scenarios, BO1 and BO2. The visualisation captures the instantaneous normalised velocity magnitude ($\|\mathbf{u}\|/U_\infty = \sqrt{u_x^2 + u_y^2 + u_z^2}/U_\infty$) across the planes and a vorticity contour ($\omega = 500\text{s}^{-1}$) in white, along with the hill and the turbine pair. This snapshot shows the interaction between the TBL and the hill, as well as the strategic positioning of the turbines and their interrelated dynamics.

In BO1, there is a significant increase in power, approximately 73% more than in the Reference case. The upstream wind turbine is






FIGURE 7 | Deficit from the optimum power as a function of the cumulative simulation number for (a) BO1 and (b) BO2. The red dashed line indicates the -1% power deficit.

TABLE 2 | Summary of optimum design variables and objective functions obtained from BO1 and BO2, alongside the Reference case.

Case	x_1	h_1	α_1 (°)	x_2	h_2	Total power
Reference case	0.00	1.00	0.0	5.00	1.00	2.70
BO1	0.12	1.50	0.0	1.17	1.02	4.66
BO2	-0.12	1.00	8.72	1.39	0.98	3.61

**FIGURE 8** | Normalised velocity magnitude flow visualisation and contour of vorticity in white of the optimum cases in BO1 (a) and BO2 (b).**FIGURE 9** | Power series with cumulative averages and 99% CI of (a) the Reference case and the optimum cases of (b) BO1 and (c) BO2.

located near the hill's centre, slightly shifted to its leeward side, and at an elevated hub height. The second turbine, placed at a distance of $1.05D$ satisfying the predefined constraint, benefits from the undisturbed wind left by the upstream turbine, as depicted in Figure 8a. BO2 was conceived in an attempt to reproduce the same behaviour through changes in the tilt angle of the upstream wind turbine instead of its hub height, thereby exploring the merits of wake steering. The best case of BO2 shows a 34% power improvement compared to the Reference case. The upstream turbine is positioned near the hill's centre, slightly towards the front, at the fixed hub height of $1D$, as seen in Figure 8b, with a tilt angle of 8.72° . This indicates that the upstream turbine is sacrificing its power generation for an overall improvement. The second turbine in BO2 is placed further from the upstream turbine compared to BO1, at a distance of $1.51D$, potentially leaving more space for the wake steering to deflect vertically the wake and have a greater effect on the downstream turbine. The difference in total power between BO1 and BO2 indicates that the limitation in tower height for the upstream wind turbine potentially limits the achievable maximum power output of BO2.

To comprehend the improved performance of the optimised configurations, we analyse the power series predicted by the ADM of the turbine pairs, as well as the mean velocity and turbulence intensity fields.

3.1 | Power Production

Figure 9 illustrates the time history of power production for each turbine, along with their cumulative averages and 99% CI for the Reference ($\hat{P}_{Ref}$), BO1 optimum ($\hat{P}_{BO1}$), and BO2 optimum ($\hat{P}_{BO2}$) cases.

In the Reference configuration, the resultant total normalised power stands at $\hat{P}_{Ref} = 2.7$. The upstream turbine generates a significant portion, $\hat{P}_{Ref_{Up}} = 2.33$, primarily due to the speed-up effect caused by the hill's presence. Conversely, the downstream turbine's contribution is only $\hat{P}_{Ref_{Down}} = 0.37$. Its reduced output is attributed to its immersion in the wake of both the hill and the upstream turbine, which adversely affects its power generation.

This power distribution is consistent with observations made when positioning the turbines individually around the hill, as illustrated in Figure 4b. Furthermore, there are notable differences in the power series of the two turbines, primarily influenced by their distinct operating environments. The upstream turbine, directly impacted by the TBL, exhibits significant variability in its power output. The TBL's dynamic and spatially inhomogeneous nature, though statistically homogeneous in the spanwise direction, is reflected in the power evolution of the upstream turbine, characterised by larger amplitude fluctuations and larger time-scale dominance. In contrast, the downstream turbine, positioned within the wake of the hill and the upstream turbine, demonstrates a signal characterised by higher frequency and lower amplitude fluctuations.

The BO1 optimal scenario achieves a total power of $\hat{P}_{BO1} = 4.66$, with the upstream and downstream turbines generating $\hat{P}_{BO1_{Up}} = 2.31$ and $\hat{P}_{BO1_{Down}} = 2.35$, respectively. Elevating the hub height of the upstream turbine and optimising the downstream turbine's placement results in a marked power surge in the latter, as evidenced by $\hat{P}_{BO1_{Down}} \gg \hat{P}_{Ref_{Down}}$. Remarkably, the power yield from the downstream turbine in BO1 is slightly above the upstream turbine power and presents similar signal characteristics to the upstream turbine. Furthermore, it is noteworthy that the power of the upstream turbine does not increase as it might be expected with the elevation of its hub height. This suggests that by elevating the upstream turbine, it potentially misses out on the benefits of the speed-up effect caused by the hill's

presence and as seen in Figure 3, the vertical gradient in the streamwise velocity is limited.

The BO2 configuration garners a total power of $\hat{P}_{BO2} = 3.61$, divided between the upstream and downstream turbines as $\hat{P}_{BO2_{Up}} = 2.14$ and $\hat{P}_{BO2_{Down}} = 1.47$, respectively. The reduction in the upstream turbine of BO2 can be explained by the use of the tilt angle of the disc. The downstream turbine continues to exhibit a substantial boost in its power generation, although lower compared to BO1 ($\hat{P}_{BO1_{Down}} > \hat{P}_{BO2_{Down}}$) due to the limitations in the upstream turbine's hub height. The time history of $\hat{P}_{BO2_{Down}}$ displays similar behaviour to that of $\hat{P}_{Ref_{Down}}$ but with higher magnitude and amplitude. This could be the imprint of the upstream turbine since both are affected by it, as seen in the following section.

3.2 | Velocity and Turbulence Intensity Fields

The difference in power generated by the downstream wind turbine across the three scenarios can be better understood by examining the time-averaged velocity fields. Figure 10a,c,e depicts the normalised time-averaged streamwise velocity field in the mid-plane for each scenario, with streamlines represented by black lines. The contour of zero streamwise velocity, highlighting the edge of reverse flow, is demarcated by the white line. For clarity, Figure 10 includes the tower of the turbines; however, these are not modelled in the simulations.




FIGURE 10 | Normalised time-averaged streamwise velocity fields (left column) and streamwise turbulence intensity fields (right column) at the symmetry plane for the Reference case (top row) and both optimised scenarios, BO1 (middle row) and BO2 (bottom row). The black lines are streamlines of the velocity field, while the white lines indicate the contours of zero streamwise velocity.

In the Reference scenario, the downstream turbine is adversely influenced by the low velocity imparted by the wake of both the hill and the upstream turbine, explaining its diminished power output. Conversely, the BO1 optimum configuration positions the downstream turbine to circumvent both the hill and turbine wake, attempting to exploit the local flow acceleration between the hill and the upstream wind turbine. In the BO2 optimum case, due to the limitations in hub height, the downstream turbine cannot avoid the upstream wakes, and therefore, it overlaps with the upstream turbine wake. The wake steering performed by the upstream turbine allows an increase in the kinetic energy available downstream by deflecting the wake vertically at the cost of producing less power upstream. As shown in the next section, this trade-off is not trivial. In both optimum cases, power increases are due to the second turbine placed very close to the first one. This is owed to the blockage of the upstream turbine (see Section 3.3). Although this turbine acts as a sink of momentum overall, it diverts some of the flow momentum to the region below it, which can be harvested by the second turbine, but only if placed close enough. Also, it is clear that in all cases the presence of the turbines influences the recirculation region of the hill. Both effects cannot be fully captured by the superposition technique, showcasing the merits of using a LES-based optimisation.

Further insights can be obtained from the second-order statistics. Figure 10b,d,f displays the normalised time-averaged streamwise turbulence intensity in the mid-plane for each scenario. In the Reference case, the downstream turbine experiences a highly turbulent flow, triggered by the flow separation at the hill, combined with the turbulence generated by the upstream turbine. In contrast, the BO1 optimum operates under relatively clean conditions. The BO2 optimum case is affected marginally by the wakes of the hill and the upstream turbine, albeit less compared to the Reference case. High levels of turbulence intensity appear when the upstream and downstream turbine wakes merge in BO2. In both the Reference case and BO2, the downstream turbine would be negatively impacted in terms of dynamic loadings due to turbulence. Overall, the optimally

configured turbines reduce the turbulent kinetic energy of the hill wake when compared to the Reference case.

3.3 | Influence of the Upstream Wind Turbine

The optimisation results indicate that the wind turbines should be placed close to each other, prompting a closer look at the influence of the upstream turbine in terms of blockage and wake steering.

In BO1 we compare the flow fields from simulations with and without the upstream turbine to analyse this influence. The presence of the upstream turbine leads to a 5% increase in the power of the downstream turbine, a result of the flow being re-directed by the upstream turbine's blockage. Figure 11a illustrates the velocity deficit evaluated in a streamwise plane 0.25D upstream of the downstream turbine ($x/D = 0.92$). The velocity deficit in BO1 is defined as $\langle u_{D-BO1} \rangle = \langle u_{nt} \rangle - \langle u_t \rangle$, where $\langle u_{nt} \rangle$ is the time-averaged streamwise velocity field in the absence of the upstream turbine and $\langle u_t \rangle$ is the mean streamwise velocity with the presence of the upstream turbine, both normalised by U_∞ . The horizontal line at $y/D = 0.7$ marks the hill's edge at this longitudinal location, and the dashed lines show the contours of velocity deficit at -0.02 and -0.04 . These contours indicate regions where the flow speed is increased by 2% and 4% of U_∞ , respectively. The turbines' centres and the extent of their rotor discs are projected onto the plane for scale. The influence of the upstream turbine is visible as a positive velocity deficit, slightly covering the downstream turbine's projected area. The contours reveal areas of faster flow close to the upstream turbine and along the hill edge. This flow acceleration spreads across a large area impacting the downstream turbine, which explains the observed performance boost. The acceleration is particularly strong near the hill, indicating that the upstream turbine's presence has altered the hill's recirculation zone across the span of the hill.

In BO2, the upstream turbine has a positive tilt angle, which, despite reducing its power output by 6% if compared to the




FIGURE 11 | Velocity deficit caused by the upstream turbine in (a) BO1 and (b) BO2 evaluated 0.25D upstream of the downstream turbine. Contours are shown with dashed lines and the grey area indicates the hill's section.

no-tilt case, leads to a substantial increase of 23.5% for the downstream turbine's power. Figure 11b illustrates the alterations in the flow pattern due to the upstream turbine's tilt, applying the velocity deficit concept as defined previously. The velocity deficit in BO2 (u_{D-BO2}) is computed by contrasting the flow with and without the upstream turbine's tilt at a plane $0.25D$ upstream of the downstream turbine ($x/D = 1.139$). The horizontal line at $y/D = 0.57$ designates the hill's edge, and the dashed lines delineate the contours of velocity deficit at -0.12 , -0.02 , 0.02 , and 0.12 . Tilting diverts the wake upwards, creating the kidney-shaped accelerated flow in the lower region of the wake's trajectory, which benefits the downstream turbine over a considerable area. Concurrently, as also seen in Figure 11a, the tilt induces modifications in the hill's recirculation, notably a horizontal deceleration over the hill edge, which has a marginal adverse effect on the downstream turbine's performance.

To further understand the impact of the tilt angle in BO2 on the total power, we turned to the insights provided by the GP surrogate model constructed at the end of BO2. The 4-dimensional surrogate model can be sliced and analysed in a simplified manner by fixing three of its design variables at the BO2 optimum, x_1 , x_2 and h_2 , and observing variations solely due to the tilt angle, α_1 . The result is a 1-dimensional function, which presents both the mean value predicted by the surrogate model and the accompanying standard deviation, which captures the associated uncertainty. For the purpose of this section, an additional training point has been considered in the surrogate model, corresponding to the BO2 optimum with zero tilt. As depicted in Figure 12, introducing a tilt angle of approximately $\alpha_1 \approx 9^\circ$ yields a 4.3% increase in total power in comparison to configurations without tilt. This result is especially noteworthy given that the two turbines are positioned in close proximity to each other, underscoring the tilt angle's significant potential to enhance total power generation for this case. As indicated by the surrogate model, using negative tilt angles reduces the performance of the turbine pair. This indicates the importance of allowing the downstream turbine to improve its conditions rather than harvesting more power upstream with an alignment to the inclined flow in front of the hill.




FIGURE 12 | Mean and standard deviation power prediction of the GP surrogate model against the tilt angle by fixing x_1 , x_2 and h_2 to the BO2 optimum values. The GP surrogate model is built with the 100 evaluations of BO2, plus the BO2 optimum without tilt.

4 | Conclusions

This numerical study focused on the optimisation of wind turbine layouts in complex terrains to maximise power output. As a proof-of-concept, the focus is on a simplified set-up with two wind turbines around a 2D hill with a limited number of 4 parameters to optimise. High-fidelity simulations (LES) were used to model the complex flow dynamics around the 2D hill and the wake interactions between two wind turbines. A BO framework was employed to systematically and efficiently find the optimal turbine layout for a given set of parameters. The two optimisations presented in this work considered parameters such as turbine locations, hub heights, and tilt angle of the upstream turbines. The first one focused on the streamwise locations and hub heights of the turbines. The second incorporated tilt angle for the upstream turbine to explore wake steering effects.

The wind farm simulator WInc3D, modified for this study, was validated against reference data to demonstrate the suitability of the IBM and wall modelling to handle complex terrains. The ADM implemented in WInc3D was extended to model yawed and tilted discs, also showing excellent agreement with results from the literature.

It was found that for the first optimisation, the best configuration positioned the upstream turbine near the hill's centre with an elevated hub height, significantly boosting the downstream turbine's power due to flow acceleration in front of the hill. For the second optimisation, the design space was modified based on the findings of the first optimisation, replacing the hub height with tilt angle for the upstream turbine as a parameter to optimise. Introducing a positive tilt angle for the upstream turbine resulted in a considerable power gain while maintaining a modest hub height.

This study demonstrated that optimising wind turbine layouts using high-fidelity LES combined with BO can lead to substantial power output improvements in complex terrains, by exploiting terrain-induced flow features and strategic turbine placement to enhance overall wind farm efficiency. The findings of this work are in line with the conclusions of our previous study [30] focusing on offshore wind farms without terrain-induced flow features, where we found that it is possible to mitigate power losses caused by wake effects through two distinct strategies: layout optimisation (micro-siting) and wake steering through yaw angle optimisation.

Future work will aim to study more complex terrains and larger wind farms with more parameters to optimise, while also improving the quality of the optimisation by adding unsteady loading and fatigue constraints into the optimisation process.

Acknowledgements

This work used the ARCHER2 UK National Supercomputing Service (<https://www.archer2.ac.uk>) via a Pioneer project and the UK Turbulence Consortium (EP/R029326/1). The first author thanks NERC and SSE Renewables Services (UK) Limited for funding a PhD studentship in the Department of Aeronautics at Imperial College London via the Science and Solutions for a Changing Planet DTP (NE/S007415/1).

Data Availability Statement

The WInc3D version used in this work is available at <https://github.com/cjaneippel/Incompact3d/tree/2d-hill>. The data that support this study are available upon reasonable request.

References

1. P. Veers, K. Dykes, E. Lantz, et al., "Grand Challenges in the Science of Wind Energy," *Science* 366, no. 6464 (2019): eaau2027.
2. P. Veers, K. Dykes, S. Basu, et al., "Grand Challenges: Wind Energy Research Needs for a Global Energy Transition," *Wind Energy Science* 7 (2022): 2491–2496.
3. F. Porté-Agel, M. Bastankhah, and S. Shamsoddin, "Wind-Turbine and Wind-Farm Flows: A Review," *Boundary-Layer Meteorology* 174 (2020): 1–59.
4. A. Crespo and J. Hernández, "A Numerical Model of Wind Turbine Wakes and Wind Farms," in *European Wind Energy Conference EWEC*, eds. W. Palz and E. Sesto, Vol. 2 (Rome, Italy, 1986), 111–115.
5. A. Crespo, F. Manuel, J. Grau, and J. Hernández, "Modelization of Wind Farms in Complex Terrain. Application to the Monteahumada Wind Farm," in *Proceedings of the 1993 European Community Wind Energy Conference* (Travemünde, Germany, 1993), 436–439.
6. A. Hyvärinen and A. Segalini, "Qualitative Analysis of Wind-Turbine Wakes Over Hilly Terrain," *Journal of Physics: Conference Series* 854 (2017): 12023.
7. R. Brogna, J. Feng, J. N. Sørensen, W. Z. Shen, and F. Porté-Agel, "A New Wake Model and Comparison of Eight Algorithms for Layout Optimization of Wind Farms in Complex Terrain," *Applied Energy* 259 (2020): 114189, <https://www.sciencedirect.com/science/article/pii/S0306261919318768>.
8. E. S. Politis, J. Prospathopoulos, D. Cabezon, K. S. Hansen, P. K. Chaviaropoulos, and R. J. Barthelmie, "Modeling Wake Effects in Large Wind Farms in Complex Terrain: The Problem, the Methods and the Issues," *Wind Energy* 15, no. 1 (2012): 161–182, <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.481>.
9. A. Segalini, "Linearized Simulation of Flow Over Wind Farms and Complex Terrains," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 375, no. 2091 (2017): 20160099, <https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2016.0099>.
10. F. Castellani, D. Astolfi, M. Mana, et al., "Investigation of Terrain and Wake Effects on the Performance of Wind Farms in Complex Terrain Using Numerical and Experimental Data," *Wind Energy* 20 (2017): 1277–1289.
11. W. Z. Shen, W. J. Zhu, E. Barlas, and Y. Li, "Advanced Flow and Noise Simulation Method for Wind Farm Assessment in Complex Terrain," *Renewable Energy* 143 (2019): 1812–1825.
12. S.-P. Breton, J. Sumner, J. N. Sørensen, K. S. Hansen, S. Sarmast, and S. Ivanell, "A Survey of Modelling Methods for High-Fidelity Wind Farm Simulations Using Large Eddy Simulation," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 375, no. 2091 (2017): 20160097.
13. X. Yang, K. B. Howard, M. Guala, and F. Sotiropoulos, "Effects of a Three-Dimensional Hill on the Wake Characteristics of a Model Wind Turbine," *Physics of Fluids* 27 (2015): 25103.
14. S. Shamsoddin and F. Porté-Agel, "Large-Eddy Simulation of Atmospheric Boundary-Layer Flow Through a Wind Farm Sited on Topography," *Boundary-Layer Meteorology* 163 (2017): 1–17.
15. L. Liu and R. J. A. M. Stevens, "Effects of Two-Dimensional Steep Hills on the Performance of Wind Turbines and Wind Farms," *Boundary-Layer Meteorology* 176 (2020): 251–269.
16. Z. Zhang, P. Huang, G. Bitsuamlak, and S. Cao, "Large-Eddy Simulation of Wind-Turbine Wakes Over Two-Dimensional Hills," *Physics of Fluids* 34 (2022): 65123.
17. D. Wang, D. Feng, H. Peng, et al., "Implications of Steep Hilly Terrain for Modeling Wind-Turbine Wakes," *Journal of Cleaner Production* 398 (2023): 136614, <https://www.sciencedirect.com/science/article/pii/S0959652623007722>.
18. J. Berg, N. Troldborg, N. N. Sørensen, E. G. Patton, and P. P. Sullivan, "Large-Eddy Simulation of Turbine Wake in Complex Terrain," *Journal of Physics: Conference Series* 854 (2017): 12003.
19. X. Yang, M. Pakula, and F. Sotiropoulos, "Large-Eddy Simulation of a Utility-Scale Wind Farm in Complex Terrain," *Applied Energy* 229 (2018): 767–777.
20. A. S. Dar, J. Berg, N. Troldborg, and E. G. Patton, "On the Self-Similarity of Wind Turbine Wakes in a Complex Terrain Using Large Eddy Simulation," *Wind Energy Science* 4 (2019): 633–644.
21. F. Bernardoni, M. Guzman, and S. Leonardi, "Implications of Complex Terrain Topography on the Performance of a Real Wind Farm," *Journal of physics: Conference series* 2505 (2023): 12052, <https://doi.org/10.1088/1742-6596/2505/1/012052>.
22. J. Wagner, T. Gerz, N. Wildmann, and K. Gramitzky, "Long-Term Simulation of the Boundary Layer Flow Over the Double-Ridge Site During the Perdigão 2017 Field Campaign," *Atmospheric Chemistry and Physics* 19 (2019): 1129–1146.
23. B. Kale, S. Buckingham, J. van Beeck, and A. Cuerva-Tejero, "Multi-Scale Modeling of a Wind Turbine Wake in Complex Terrain," *Journal of Physics: Conference Series* 2505, no. 1 (2023): 12012, <https://doi.org/10.1088/1742-6596/2505/1/012012>.
24. S. Cao and T. Tamura, "Experimental Study on Roughness Effects on Turbulent Boundary Layer Flow Over a Two-Dimensional Steep Hill," *Journal of Wind Engineering and Industrial Aerodynamics* 94 (2006): 1–19.
25. J. Feng and W. Z. Shen, "Wind Farm Layout Optimization in Complex Terrain: A Preliminary Study on a Gaussian Hill," *Journal of Physics: Conference Series* 524, no. 1 (2014): 12146, <https://doi.org/10.1088/1742-6596/524/1/012146>.
26. J. Feng, W. Z. Shen, and Y. Li, "An Optimization Framework for Wind Farm Design in Complex Terrain," *Applied Sciences* 8, no. 11 (2018): 2053, <https://www.mdpi.com/2076-3417/8/11/2053>.
27. S. R. Reddy, "An Efficient Method for Modeling Terrain and Complex Terrain Boundaries in Constrained Wind Farm Layout Optimization," *Renewable Energy* 165 (2021): 162–173, <https://www.sciencedirect.com/science/article/pii/S0960148120316426>.
28. L. E. Boni Cruz and B. Carmo, "Wind Farm Layout Optimization Based on CFD Simulations," *Journal of the Brazilian Society of Mechanical Sciences and Engineering* 42 (2020): 433.
29. J. Allen, R. King, and G. Barter, "Wind Farm Simulation and Layout Optimization in Complex Terrain," *Journal of Physics: Conference Series* 1452, no. 1 (2020): 12066, <https://doi.org/10.1088/1742-6596/1452/1/012066>.
30. N. Bempedelis, F. Gori, A. Wynn, S. Laizet, and L. Magri, "Data-Driven Optimisation of Wind Farm Layout and Wake Steering With Large-Eddy Simulations," *Wind Energy Science* 9, no. 4 (2024): 869–882, <https://wes.copernicus.org/articles/9/869/2024/>.
31. A. Larroque, M. F. de Pando, and L. Lafuente, "Cylinder Drag Minimization Through Wall Actuation: A Bayesian Optimization Approach," *Computers & Fluids* 240 (2022): 105370, <https://www.sciencedirect.com/science/article/pii/S0045793022000500>.
32. J. P. Goit and J. Meyers, "Optimal Control of Energy Extraction in Wind-Farm Boundary Layers," *Journal of Fluid Mechanics* 768 (2015): 5–50.

33. W. Munters and J. Meyers, "Dynamic Strategies for Yaw and Induction Control of Wind Farms Based on Large-Eddy Simulation and Optimization," *Energies* 11, no. 1 (2018): 177, <https://www.mdpi.com/1996-1073/11/1/177>.
34. C. Talnikar and Q. Wang, "Adjoint-Based Trailing Edge Shape Optimization of a Transonic Turbine Vane Using Large Eddy Simulations," (2020), arXiv preprint arXiv:2011.06744.
35. W. Munters, C. Meneveau, and J. Meyers, "Shifted Periodic Boundary Conditions for Simulations of Wall-Bounded Turbulent Flows," *Physics of Fluids* 28, no. 2 (2016): 25112, <https://doi.org/10.1063/1.4941912>.
36. G. Deskos, S. Laizet, and R. Palacios, "WInc3D: A Novel Framework for Turbulence-Resolving Simulations of Wind Farm Wake Interactions," *Wind Energy* 23 (2020): 779–794.
37. P. Bartholomew, G. Deskos, R. A. S. Frantz, F. N. Schuch, E. Lamballais, and S. Laizet, "Xcompact3d: An Open-Source Framework for Solving Turbulence Problems on a Cartesian Mesh," *SoftwareX* 12 (2020): 100550.
38. A. G. Kravchenko and P. Moin, "On the Effect of Numerical Errors in Large Eddy Simulations of Turbulent Flows," *Journal of Computational Physics* 131, no. 2 (1997): 310–322, <https://www.osti.gov/biblio/494295>.
39. J. Smagorinsky, "General Circulation Experiments With the Primitive Equations: I. The Basic Experiment," *Monthly Weather Review* 91, no. 3 (1963): 99–164, https://journals.ametsoc.org/view/journals/mwre/91/3/1520-0493_1963_091_0099_gcewtp_2_3_co_2.xml.
40. M. Calaf, C. Meneveau, and J. Meyers, "Large Eddy Simulation Study of Fully Developed Wind-Turbine Array Boundary Layers," *Physics of Fluids* 22 (2010): 15110.
41. P. J. Mason and D. J. Thomson, "Stochastic Backscatter in Large-Eddy Simulations of Boundary Layers," *Journal of Fluid Mechanics* 242 (1992): 51–78.
42. S. Laizet and E. Lamballais, "High-Order Compact Schemes for Incompressible Flows: A Simple and Efficient Method With Quasi-Spectral Accuracy," *Journal of Computational Physics* 228 (2009): 5989–6015.
43. S. K. Lele, "Compact Finite Difference Schemes With Spectral-Like Resolution," *Journal of Computational Physics* 103 (1992): 16–42.
44. E. Motheau and J. Abraham, "A High-Order Numerical Algorithm for DNS of Low-Mach-Number Reactive Flows With Detailed Chemistry and Quasi-Spectral Accuracy," *Journal of Computational Physics* 313 (2016): 430–454, <https://www.sciencedirect.com/science/article/pii/S0021999116001315>.
45. S. Laizet and N. Li, "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 (2011): 1735–1757.
46. R. Gautier, S. Laizet, and E. Lamballais, "A DNS Study of Jet Control With Microjets Using an Immersed Boundary Method," *International Journal of Computational Fluid Dynamics* 28, no. 6-10 (2014): 393–410.
47. T. Dairay, E. Lamballais, and S. Benhamadouche, "Mesh Node Distribution in Terms of Wall Distance for Large-Eddy Simulation of Wall-Bounded Flows," *Flow, Turbulence and Combustion* 100 (2018): 617–626.
48. E. Lamballais, R. Vicente Cruz, and R. Perrin, "Viscous and Hyperviscous Filtering for Direct and Large-Eddy Simulation," *Journal of Computational Physics* 431 (2021): 110115, <https://www.sciencedirect.com/science/article/pii/S0021999121000073>.
49. C. Jané-Ippel, N. Bempedelis, R. Palacios, and S. Laizet, "High-Fidelity Simulations of Wake-to-Wake Interaction in an Atmospheric Boundary Layer Over a Complex Terrain," *Journal of Physics: Conference Series* 2505, no. 1 (2023): 12033, <https://doi.org/10.1088/1742-6596/2505/1/012033>.
50. F. Jaegle, O. Cabrit, S. Mendez, and T. Poinso, "Implementation Methods of Wall Functions in Cell-Vertex Numerical Solvers," *Flow, turbulence and combustion* 85 (2010): 245–272, <https://link.springer.com/article/10.1007/s10494-010-9276-1>.
51. N. Bempedelis, S. Laizet, and G. Deskos, "Turbulent Entrainment in Finite-Length Wind Farms," *Journal of Fluid Mechanics* 955 (2023): A12.
52. R. N. King, K. Dykes, P. Graf, and P. E. Hamlington, "Optimization of Wind Plant Layouts Using an Adjoint Approach," *Wind Energy Science* 2, no. 1 (2017): 115–131, <https://wes.copernicus.org/articles/2/115/2017/>.
53. L. P. Chamorro and F. Porté-Agel, "Effects of Thermal Stability and Incoming Boundary-Layer Flow Characteristics on Wind-Turbine Wakes: A Wind-Tunnel Study," *Boundary-layer meteorology* 136 (2010): 515–533, <https://link.springer.com/article/10.1007/s10546-010-9512-1>.
54. K. S. Heck, H. M. Johlas, and M. F. Howland, "Modelling the Induction, Thrust and Power of a Yaw-Misaligned Actuator Disk," *Journal of Fluid Mechanics* 959 (2023): A9.
55. J. J. Cutler, A. P. Stanley, A. Ning, and J. J. Thomas, "Optimization of Turbine Tilt in a Wind Farm," in *AIAA Scitech 2021 Forum* (2021), 1180, <https://arc.aiaa.org/doi/abs/10.2514/6.2021-1180>.
56. J. Annoni, A. Scholbrock, M. Churchfield, and P. Fleming, "Evaluating Tilt for Wind Plants," in *2017 American Control Conference (ACC)* (IEEE, 2017), 717–722.
57. I. Janajreh, R. Qudaih, I. Talab, and C. Ghenai, "Aerodynamic Flow Simulation of Wind Turbine: Downwind Versus Upwind Configuration," *Energy Conversion and Management* 51, no. 8 (2010): 1656–1663, <https://www.sciencedirect.com/science/article/pii/S0196890409005196>.
58. M. Bergmann, C. Morsbach, G. Ashcroft, and E. Kügeler, "Statistical Error Estimation Methods for Engineering-Relevant Quantities From Scale-Resolving Simulations," *Journal of Turbomachinery* 144, no. 3 (2021): 31005, <https://doi.org/10.1115/1.4052402>.
59. M. Diessner, J. O'Connor, A. Wynn, et al., "Investigating Bayesian Optimization for Expensive-to-Evaluate Black Box Functions: Application in Fluid Dynamics," *Frontiers in Applied Mathematics and Statistics* 8 (2022): 1076296, <https://www.frontiersin.org/articles/10.3389/fams.2022.1076296>.
60. N. Bempedelis and L. Magri, "Bayesian Optimization of the Layout of Wind Farms With a High-Fidelity Surrogate Model," (2023), arXiv preprint arXiv:2302.01071.
61. The GPyOpt authors, "GPyOpt: A Bayesian Optimization Framework in Python," (2016), <http://github.com/SheffieldML/GPyOpt>.
62. J. González, Z. Dai, P. Hennig, and N. Lawrence, "Batch Bayesian Optimization via Local Penalization," *Artificial intelligence and statistics* (PMLR, 2016), 648–657.