



Sensitivity analysis of wake steering optimisation for wind farm power maximisation

Filippo Gori , Sylvain Laizet , and Andrew Wynn

Department of Aeronautics, Imperial College London, SW7 2AZ London, UK

Correspondence: Filippo Gori (f.gori21@imperial.ac.uk)

Abstract. Modern large-scale wind farms consist of multiple turbines clustered together, usually in well-structured formations. Clustering has a number of drawbacks during a wind farm's operation, as some of the downstream turbines will inevitably operate in the wake of those upstream, with a significant reduction in power output and an increase in fatigue loads. Wake steering, a control strategy in which upstream wind turbines are misaligned with the wind to redirect their wakes away from downstream turbines, is a promising strategy to mitigate power losses. The purpose of this work is to investigate the sensitivity of open-loop wake steering optimisation in which an internal predictive wake model is used to determine the farm power output as a function of the turbine yaw angles. Three different layouts are investigated with increasing levels of complexity. A simple 2×1 farm layout in aligned conditions is first considered, allowing for a careful investigation of sensitivity to wake models and operational set-points. A medium-complexity case of a generic 5×5 farm layout in aligned conditions is examined, to enable the study of a more complex design space. The final layout investigated is the Horns Rev wind farm (80 turbines), for which there has been very little study of the performance or sensitivity of wake steering optimisation. Overall, the results indicate a strong sensitivity of wake steering strategies to both analytical wake model choice, and to the particular implementation of algorithms used for optimisation. Significant variability can be observed in both farm power improvement and optimal yaw settings, depending on the optimisation set-up. Through a statistical analysis of the impact of optimiser initialisation and a study of the multi-modal and discontinuous nature of the underlying farm power objective functions, this study shows that the uncovered sensitivities represent a fundamental challenge to robustly identifying globally optimal solutions for the high-dimensional optimisation problems arising from realistic wind farm layouts. This paper proposes a simple strategy for sensitivity mitigation by introducing additional optimisation constraints, leading to higher farm power improvements and more consistent, coherent, and practicable optimal yaw angle settings.

1 Introduction

Wind energy now plays a central role in meeting world energy requirements, driven by the urgent need to mitigate climate change and a significant recent reduction in its levelised cost. New and ambitious international renewables targets, such as the European Commission's proposed installation of up to 450 GW offshore wind capacity by 2050 (Pryor et al., 2021), will require further significant increases in wind energy capacity. Large wind arrays with closely spaced turbines will be used to meet this demand. These have logistical advantages over sparse layouts in terms of lower infrastructure and maintenance costs



and increased energy density. However, the performance of large wind arrays is degraded by the aerodynamic interactions between turbines. Upstream-generated turbine wakes exhibit lower wind speeds and increased turbulence intensity, causing typical annual power losses of 10 % to 30 % (Nygaard, 2014; Barthelmie et al., 2009, 2010) in addition to increased fatigue loading.

30 Wake steering, in which upstream turbines are yawed to deflect their wakes away from downstream machines, is a leading control technique used to mitigate the effects of wake-turbine interactions. Farm power increases ranging from 5 % to 47 % have been demonstrated in high-fidelity simulations (Fleming et al., 2015, 2018), wind tunnel campaigns (Campagnolo et al., 2016; Bastankhah et al., 2019), and field tests (Howland et al., 2019). To successfully implement wake steering, one must navigate a natural trade-off: while wake steering can increase wind speed and power production at downstream turbines, it also
35 reduces both the effective swept area and power production of any yawed wake-generating turbine. For large arrays of turbines, in which there are multiple competing trade-offs arising from each interacting pair of turbines, the question of maximising farm power via wake steering becomes both a high-dimensional and non-convex optimisation problem. Such problems are challenging to solve, and their solutions can exhibit high parametric sensitivity, which may partially explain the wide range of reported farm power improvements in the literature.

40 With a view to enabling robust and predictable wake steering optimisation for large wind farms, this paper investigates the *sensitivity* of wake steering optimisation. We focus on open-loop approaches in which internal predictive wake models are used to determine the farm power output (the cost to be maximised) as a function of the turbine yaw angles (the decision variables). A number of predictive models will be considered, including the Jensen (Jensen, 1983), Multi-zone (Gebraad et al., 2014), Gaussian (Bastankhah and Porté-Agel, 2014) and Gauss-Curl-Hybrid (King et al., 2021) models, each of which is
45 suitably adapted using the approaches of Jiménez et al. (2009) and Bastankhah et al. (2022) to capture yaw-induced wake modifications. These are chosen to be representative of common models used in wake-steering optimisation (Kheirabadi and Nagamune, 2019).

By optimisation sensitivity, we refer to the variability of wake steering optimisation to (i) the choice of the underlying predictive model; and (ii) the choice of optimisation algorithm and its particular parametric implementation; and (iii) to the
50 given operating set-point, such as farm layout or atmospheric conditions. Furthermore, we view sensitivity from two different perspectives. The first is to consider the sensitivity, or variance, of the optimized farm power to the underlying model, algorithm, or operational condition. The second, and arguably the most important, is to study the sensitivity of the optimal decision variables themselves, i.e., the optimal yaw angles obtained from wake steering optimisation. Of particular interest is to identify situations of high sensitivity: that is, when either the predicted maximised power output, or the optimal yaw angles, are seen to
55 vary significantly with small changes to the wake model, algorithm, or operating set-point. When high sensitivity is identified, we will also study potential algorithmic changes to reduce it.

Understanding, quantifying, and reducing cases of high sensitivity is of fundamental importance for robust wake steering. To explain this, it should be noted that to implement open-loop wake steering, tuning, or identification, of wake model parameters from field data or high-fidelity simulations is required. As with any tuned model, there will be a mismatch between the model's
60 prediction and the true farm behaviour. A first problem arises if the predicted optimised farm power output is highly sensitive



to the wake model parameters or algorithmic implementation. In this case, small modelling mismatches may cause large deviations between expected and true farm performance, possibly nullifying any predicted power increase when applied in the field. A second problem occurs in the case of high optimal yaw angle sensitivity. Here, a small parametric change (e.g., to model re-tuning, or to atmospheric conditions) may require large yaw angle changes, and it is practicably undesirable to make large control input changes in response to only small operational perturbations. As will be discussed in Sect. 4.3, sensitivity of optimal yaw angles can be observed in recent open-loop wake steering studies for large-scale farms such as the Horns Rev installation (Dou et al., 2020). More generally, we will show that high decision variable sensitivity also presents a fundamental challenge which must be overcome when one seeks to globally optimise wake steering strategies for the high-dimensional optimisation problems arising from realistic farm layouts.

In this paper, we do not consider active wake control or closed-loop wake control in which online sensor information is used to dynamically adapt wake steering strategies (see Kheirabadi and Nagamune (2019) and Houck (2021) for comprehensive reviews). While recent, and promising, feedback control strategies have been proposed (Doekemeijer et al., 2019, 2020; Howland et al., 2020, 2022), closed-loop approaches still typically employ predictive wake models in order to calculate their time-varying control inputs. Consequently, the question of predictive sensitivity is still important in the more complex closed-loop setting. By first studying open-loop wake optimisation strategies in this paper, we will therefore also provide insight into the behaviour of more sophisticated closed-loop approaches to wake steering optimisation.

Due to the large number of underlying modelling, algorithmic, and parametric choices, in this paper we propose a hierarchy of test-cases to facilitate the understanding of the different sensitivities present in wake steering optimisation. At the simplest level, a minimal optimisation problem consisting of a 2×1 farm layout in aligned conditions is first considered. In this case, optimisation can be performed by a simple parametric sweep, removing any sensitivity to algorithmic choice or implementation. The 2×1 case, therefore, allows a careful investigation of sensitivity to only the underlying wake model and operational set-points. Next, a medium-complexity case of a 5×5 farm layout in aligned conditions is considered. This may be viewed as several non-trivially interacting 2×1 blocks, leading to a more complex design space. Sensitivity, therefore, depends on both the algorithmic implementation and underlying wake model, but quantification of the 2×1 case facilitates the decoupling of these two dependencies. We present a statistical analysis of optimiser sensitivity in this case, in addition to comparing the performance of global and gradient-based optimisation algorithms. Finally, a representative layout corresponding to the Horns Rev wind farm is considered. The investigation of this more complex optimisation problem will be used to confirm whether the wake steering sensitivities uncovered in the low and medium complexity cases transfer to a realistic setup. Finally, we propose a simple strategy for sensitivity mitigation, based on appropriately chosen optimisation constraints, which is effective even for this realistic and complex wake steering optimisation problem. To the best of the authors' knowledge, previous parametric investigations into wake optimisation sensitivity (Rak and Santos Pereira, 2022; Göçmen et al., 2022) have not studied the effect of optimiser class or analysed sensitivity from a statistical perspective, and have not presented possible strategies for sensitivity mitigation.

The remainder of the paper is organised as follows. Section 2 gives an overview of the FLOW Redirection and Induction in Steady State (FLORIS) modelling framework and introduces the wake and deflection models. In Sect. 3, a description of the



optimisation setup is provided, including wind farm layouts, atmospheric conditions, and optimisation algorithms. Sensitivity results and discussion are presented in Sect. 4 for the 2×1 (Sect. 4.1), 5×5 (Sect. 4.2), and Horns Rev (Sect. 4.3) cases. Finally, conclusions are given in Sect. 5. Appendices provide additional details on wake and deflection model formulations (Appendix-A) and a validation of the optimisation framework (Appendix-B).

100 2 Wind farm modelling

Wind farm modelling in the present work is conducted with version 2.4 of the open-source FLORIS framework (NREL, 2021) by the National Renewable Energy Laboratory (NREL). The wake models considered in this paper are Jensen, Multi-zone, Gaussian, and Gauss-Curl Hybrid (GCH) and their respective deflection models. The four selected models, fully described in Appendix-A, have different complexities, physical modelling capabilities, and empirical parameter dependencies and are selected to reflect the models commonly used by the wind energy community in optimisation and control studies (Kheirabadi and Nagamune, 2019). We emphasise that the purpose of this study is not to identify which model is best to use to study a particular wind farm. Rather, we seek to identify situations in which wake steering optimisation is highly sensitive either to the choice of underlying model, or to the interaction between the model choice and the particular parametric implementation of commonly-used optimisers (see Section 3 for a full description of optimisation algorithms used), since high sensitivity presents a fundamental barrier to robust optimisation performance, even using models tuned to high-fidelity simulation data or field data. As a result, the wake and deflection model parameters for this study simply use the FLORIS recommended standard values, as reported in Table 1.

For completeness, we now give a brief overview of FLORIS's wind farm modelling structure. First, an initial condition is defined by specifying atmospheric inflow, wind farm layout, turbine geometry, and operational conditions. Next, a wake model calculates each turbine's steady streamwise velocity deficit. The computation is sequential to allow additional considerations on added turbulence. For yawed turbine cases, a deflection model is employed to determine and apply a cross-stream shift in the streamwise velocity deficit field. The GCH model includes a lateral and vertical velocity component estimation for an additional streamwise velocity correction. Finally, streamwise velocity deficits for each turbine are combined with a superposition model and are applied to the initial flow field. In the current study, all wake models use the sum of squares freestream superposition (SOSFS) model developed by Katic et al. (1987). Based on the conservation of the mean kinetic energy deficit during wake interaction, the combined wake velocity U_w , dependant on the three spatial dimensions (x, y, z) , is defined as

$$U_w(x, y, z) = U_\infty - \sqrt{\sum_i (U_\infty - u_w^i(x, y, z))^2}, \quad (1)$$

where U_∞ is the mean free stream velocity, and u_w^i is the wake streamwise velocity induced by turbine i in stand-alone conditions.

125 Farm power calculations are conducted as follows. Turbine operational profiles consist of lookup tables for power and thrust coefficients (C_P and C_T , respectively) generated by the FAST (Jonkman, 2021) (Fatigue, Aerodynamics, Structures, and Turbulence) aeroelastic simulator developed by NREL. Given a resolved wind farm flow field $U_w(x, y, z)$, the streamwise

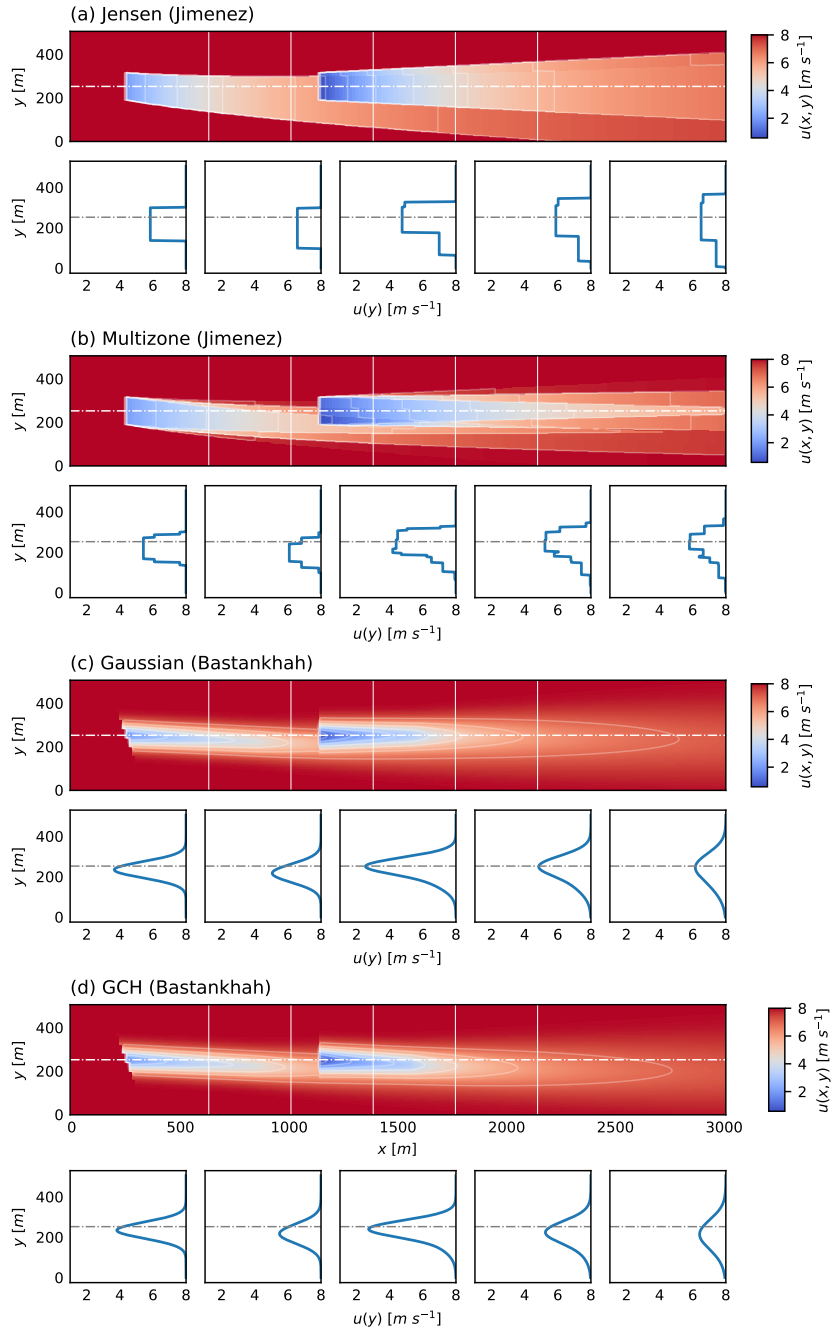


Figure 1. Streamwise velocity at hub height of a 2×1 farm layout for (a) Jensen (Jimenez), (b) Multi-zone (Jimenez), (c) Gaussian (Bastankhah) and (d) GCH (Bastankhah) wake and deflection models at an upstream rotor yaw angle of 20° . White vertical lines represent the streamwise locations for the velocity profile section plots, ordered by increasing downstream distance, while the horizontal lines indicate the upstream rotor centerline.



velocities at the turbine rotor grid points $U_{\text{rotor}}(x, y, z)$ are defined by interpolation. The averaged streamwise velocity \bar{U} of the spatial extent of the un-yawed turbine is expressed by

$$130 \quad \bar{U} = \sqrt[3]{U_{\text{rotor}}^3}. \quad (2)$$

Finally, turbine power is calculated as

$$P = 0.5\rho AC_P(\bar{U}(\cos\gamma)^w)^3, \quad (3)$$

where ρ is the air density, A is the un-yawed rotor swept area, C_P is the interpolated power coefficient based on the rotor averaged streamwise velocity \bar{U} , γ is the turbine yaw angle to the inflow wind direction, and w is a tuneable parameter
135 matching the power loss due to yaw misalignment observed in high-fidelity simulations (following (Fleming et al., 2017), $w = 0.627$).

Figure 1 presents a visual overview of the main wake modelling capabilities of the models employed in this study, including streamwise velocity deficit distributions, expansion rates, and wake deflections due to yaw. This figure will be referred to subsequently to help understand the sensitivity of wind farm wake steering optimisation to the choice of underlying wake
140 models. Further details on wake and deflection model formulations are provided in Appendix-A.

3 Optimisation and wind farm setups

Open-loop wake steering optimisation for power maximisation is conducted on three different farm layouts. The first two are a 2×1 farm layout and a 5×5 farm layout, both in fully aligned wind conditions, with NREL 5 MW turbines (Jonkman et al., 2009) with rotor diameter $D = 126$ m and hub height $H = 90$ m. In the base condition, turbine spacing is set to 7 rotor
145 diameters in the streamwise direction and 5 rotor diameters in the cross-stream direction. The third farm layout considered is representative of the offshore Horns Rev wind farm (Jensen et al., 2004). It consists of 80 Vestas V-80 2 MW turbines with rotor diameter $D = 80$ m and hub height $H = 70$ m, arranged in a rhomboid shape. For the considered wind direction ($wd = 270^\circ$), the farm exhibits an aligned layout with ten rows and eight columns and a distance between two consecutive turbines of $7D$ for both the streamwise and cross-stream directions.

150 For all the optimisation test-cases, the incoming flow is fully aligned with the farm layout with a wind speed of 8 m s^{-1} in 5 % ambient turbulence intensity. Wind shear and veer are not taken into account. The sensitivity of wake steering strategies to operating set-points is only investigated for the minimal 2×1 case, where variations in streamwise spacing and ambient turbulence intensity are performed. Validation for the employed optimisation framework is conducted by comparing results with the findings by Rak and Santos Pereira (2022) and is provided in Appendix-B.

155 The optimisation problem's objective function in all cases is chosen to be the normalised farm power production

$$G(\boldsymbol{\gamma}) := \frac{P(\boldsymbol{\gamma})}{P(\mathbf{0})}, \quad (4)$$

where $\boldsymbol{\gamma} = \{\gamma_i\}_{i=1}^N \in \mathbb{R}^N$ are the yaw angles of the N turbines considered in each configuration, $P(\mathbf{0})$ is the farm power without wake steering, and $P(\boldsymbol{\gamma})$ is the farm power achieved for a particular choice of yaw angles $\boldsymbol{\gamma}$. In all cases, we solve the



constrained maximisation problem

160
$$\max_{\gamma} G(\gamma) \tag{5}$$

subject to $-25^\circ \leq \gamma_i \leq 25^\circ, \quad 1 \leq i \leq N,$

which imposes typical operational bounds of $\pm 25^\circ$ on the yaw misalignment of upstream turbines. In section 4.3, additional constraints corresponding to modified bounds and yaw angle monotonicity will be added to the optimisation problem (5) to mitigate optimisation sensitivity. These constraints are imposed by adding further linear inequality constraints to (5) and consequently do not result in any intrinsic increase in complexity of the optimisation problem.

165 **3.1 Optimisation algorithms and parameter settings**

The optimisation algorithms used are the Sequential Least Squares Programming (SLSQP) method developed by Kraft (1988) and the Trust Region Bayesian Optimisation (TuRBO) approach by Eriksson et al. (2019). The gradient-based optimiser SLSQP is the standard algorithm used in the FLORIS framework. TuRBO represents a novel probabilistic approach for the global optimisation of large-scale high-dimensional problems. In this method, after an initial space sampling with a Latin Hypercube
170 Sampling (LHS) method, multiple Gaussian process surrogate models are fitted in evolving promising areas of the design space, referred to as trust regions. Using Thompson Sampling (TS) as an acquisition function, a principled global allocation of samples across the models is achieved via an implicit bandit approach. The number of initial evaluations generated by the LHS method is set to double the number of optimisation variables N (here the number of turbines). Moreover, a single trust region is used in all cases to enable a fair comparison with the SLSQP algorithm. The parameters used for SLSQP and TuRBO
175 optimisers are shown in Table 2.

As will be discussed subsequently for the 5×5 case (see, e.g., Figure 6), the considered objective functions $G(\gamma)$ typically exhibit multi-modality, flat regions and discontinuities. These features present a fundamental challenge when seeking to robustly identify global maxima. In the case of SLSQP, a sensible choice of optimisation parameters is required to address this complexity and to essentially regularise the optimisation problem. The parameter to consider is the precision goal for the value
180 of the objective function in the stopping criteria, called $ftol$. It corresponds to the achievement of equal objective function values in the evaluations performed for the numerical approximation of the gradient. Large values of $ftol$ may cause the optimiser to terminate in flat regions, significantly increasing the variability of the objective function and the optimal decision variables to the choice of optimisation initialisation. However, a small $ftol$ typically increases the number of objective function evaluations required for convergence.

185 In this study, $ftol = 1e-16$ for the 5×5 case, and $ftol = 1e-12$ for the Horns Rev case are identified as a suitable trade-off between these two behaviours, based on preliminary investigations. The different choice of parameter between the two cases is motivated by the higher dimensionality of the Horns Rev optimisation problem.

The second SLSQP parameter to consider is $eps \in [0, 1]$, corresponding to the step-size (in this case equal to $eps \times [25^\circ - (-25^\circ)]$) used for the finite difference gradient approximation. Large eps values lead to inaccurate gradient approximations.
190 However, as shown in Fig. 6, the objective functions considered in this study may have either discontinuities or discontinuous



Table 1. Wake and deflection model parameters.

Model	Parameter	Parameter value
Jensen (Jimenez)	k (on/offshore)	0.05 or 0.04
	k_d	0.05
Multi-zone (Jimenez)	k_e	0.05
	$m_{e,q}$	[-0.5, 0.3, 1.0]
	$M_{u,q}$	[0.5, 1.0, 5.5]
	a_u	12.00
	b_u	1.30
	k_d	0.05
Gaussian - GCH (Bastankhah)	k_a	0.380
	k_b	0.004
	α	0.580
	β	0.077
	ϵ	0.2D
	ϕ	2.000

Table 2. SLSQP and TuRBO parameters.

Optimiser	Parameter	Parameter value
SLSQP (5×5)	$ftol$	1e-16
	eps [$-25^\circ, +25^\circ$]	5e-2
SLSQP (Horns Rev)	$ftol$	1e-12
	eps [$-25^\circ, +25^\circ$]	5e-2
	eps [$0^\circ, +25^\circ$]	1e-1
TuRBO	τ_{succ}	3.0
	τ_{fail}	25.0
	L_{min}	0.5 ⁷
	L_{max}	1.6
	L_{init}	0.8

derivatives, especially when using Jensen and Multi-zone models. This ‘rough’ objective function behaviour, in conjunction with small eps values, may cause the SLSQP optimiser to identify optimal decision variables in local maxima. A 2.5° step-size ($eps = 0.05$) is identified to be a suitable value since this relatively large value of eps has the effect of artificially smoothing the rough objective function while limiting the loss of information about the local gradient.

195 4 Results

4.1 Two-turbine case - NREL 5 MW

We consider the problem of farm power maximisation for a 2×1 wind farm layout with NREL 5 MW turbines. The only optimisation variable is the yaw angle $\gamma_1 \in [-25^\circ, 25^\circ]$ of the upstream turbine since the optimal yaw angle of the downstream turbine is necessarily zero. Optimisation is therefore performed via a simple parameter sweep of γ_1 . Different streamwise turbine spacings and ambient turbulence intensity are considered, in order to provide insights into the sensitivity of both optimised farm power and the optimal yaw angle, denoted γ_1^* , to operational conditions. Observations made for this minimal study will be used to explain and analyse the sensitivity and performance of wake steering optimisation for more complex farm layouts in Sections 4.2 and 4.3.

Figures 2-a,b show the extracted power $P_{WF}(\gamma_1)$ of the entire 2×1 farm and the extracted power $P_{T_2}(\gamma_1)$ of the downstream turbine only. Power production of the upstream machine is the same for all models. Clear sensitivity to model choice is observed




Figure 2. Power function as upstream turbine yaw angle γ_1 ranges $\pm 25^\circ$. Top plots: wind farm power P_{WF} (a) and downstream turbine power P_{T_2} (b) for 7 D streamwise turbine spacing and ambient turbulence intensity $I_0 = 5\%$. Middle plots: wind farm power P_{WF} in $I_0 = 5\%$ with variable streamwise spacing: 4 D (c) and 10 D (d). Bottom plots: wind farm power P_{WF} for 7 D downstream spacing and variable turbulence intensity: $I_0 = 10\%$ (e) and $I_0 = 15\%$ (f).

in Figure 2-a with the Jensen model's optimal yaw angle at $\gamma_1^* = 0^\circ$ while, in contrast, all other models have optimal yaw angles close to constraint boundaries with $\gamma_1^* > 20^\circ$. Model sensitivity is caused by the flatness of the Jensen model's downstream power curve apparent in Figure 2-b, that is, $\frac{d}{d\gamma_1} P_{T_2}(\gamma_1)$ is significantly smaller for Jensen than for other models. Flatness arises due to the uniform, or “top-hat”, profile of the Jensen distribution (see Figure 1), which results in a lack of sensitivity of streamwise velocity deficit to moderate yaw perturbations. In contrast, it can also be seen in Figure 1 that all other models predict a streamwise velocity deficit which is both larger at the centreline and decreases more rapidly in the transverse direction.



Consequently, Gaussian, GCH and Multi-zone models give downstream power curves $P_{T_2}(\gamma_1)$, which are more sensitive to changes in γ_1 than the Jensen model, explaining the strong model-dependent sensitivity of the wake steering optimisation problem for the 2×1 farm.

215 Model sensitivity can be eliminated by increasing either turbine spacing or turbulence intensity I_0 . Figures 2-c,d show that the optimal yaw angle for the Gaussian, GCH and Multi-zone models decreases with increased turbine spacing. Although not shown, $\gamma_1^* \approx 0^\circ$ for all models for spacing greater than $32 D$. Similar behaviour can be observed in Figures 2-e,f as I_0 increases, with $\gamma_1^* \approx 0^\circ$ whenever $I_0 \geq 15\%$ for all models which incorporate this parameter (the Multi-zone model does not depend on I_0). From an optimisation perspective, increased spacing and turbulence intensity have the analogous effect of
220 increasing wake recovery and effective wake diameter at the downstream turbine. This reduces the sensitivity of $P_{T_2}(\gamma_1)$ to γ_1 and renders the wake steering optimisation problem trivial. Despite this observation, it is clear from Figure 2 that for any practically-relevant turbine spacing of up to $10D$, and for turbulence intensities of $I_0 \leq 10\%$ (at a typical $7 D$ spacing), it is still the case that significant model sensitivity is present in the 2×1 farm power maximisation problem.

With a view towards understanding the sensitivity of wake steering optimisation for more complex farms, it is important to
225 highlight the *multi-modality* of the farm power curves $P_{WF}(\gamma_1)$ for the Gaussian, GCH and Multi-zone models, that is, the presence of multiple and distinct maxima. Multi-modality can be seen for all considered spacings in Figures 2-a,c,d, and for $I_0 \leq 10\%$ in Figure 2-e. For the Gaussian model, this arises simply due to the symmetry of $P_{T_2}(\gamma_1)$ for the aligned inflow conditions considered here. Multi-zone and GCH models are weakly asymmetric due to the incorporation of wake rotation effects, with $P_{T_2}(\gamma_1) > P_{T_2}(-\gamma_1)$, for $\gamma_1 > 0$, in the case considered here of a clockwise spinning rotor. However, the farm
230 power $P_{WF}(\gamma_1)$ curves for both models are still clearly multi-modal.

In summary, this minimal 2×1 case demonstrates that wake steering optimisation can exhibit model sensitivity due to the flatness of the Jensen model, while the Gaussian, GCH and Multi-zone models exhibit multi-modal objective functions. Both observations are of importance to understanding the behaviour and sensitivity of wake steering optimisation for the following cases of more complex farm layouts.

235 4.2 Multiple-turbine optimisation: 5×5 array of NREL 5 MW turbines

We consider wake steering optimisation for farm power maximisation on a 5×5 wind farm layout with NREL 5 MW turbines. The objective of this medium-complexity case is to understand the sensitivity of wake steering performance to both the choice of wake model, and to the choice and particular implementation of an optimisation algorithm. Due to the high-dimensional and multi-modal nature of the underlying optimisation problem in this case, statistical analysis is required to quantify optimi-
240 sation sensitivity (see Section 4.2.3). However, we first study several indicative cases, which will render the statistical results transparent.

4.2.1 Initial yaw angle sensitivity in gradient-based optimisation

Figure 3 shows the comparison between wake models for an SLSQP optimisation initialised with a yaw angle of 12° for all turbines. Significant discrepancies between models are observed for both predicted wind farm power improvement, and for the

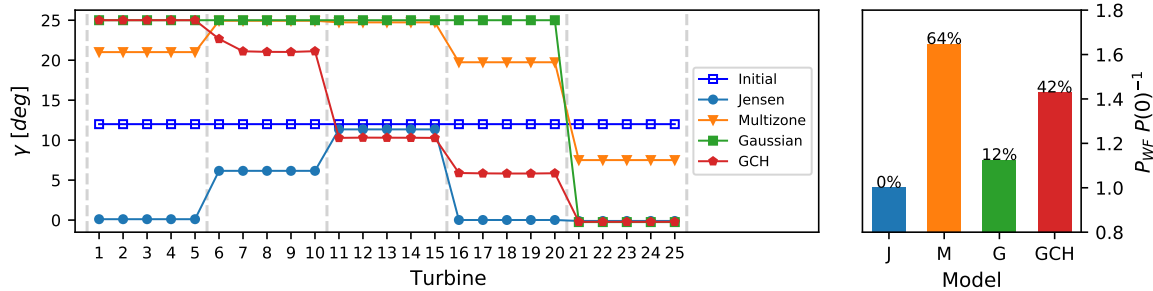


Figure 3. Model comparison for a 5×5 wind farm SLSQP optimisation. Left plot: optimal set-points per turbine, with 12° initial yaw angles illustrated in blue. Dashed lines delimit wind farm rows. Right plot: resulting objective function (normalised farm power) for Jensen (J), Multi-zone (M), Gaussian (G) and Gauss-Curl Hybrid (GCH). Turbine naming convention: turbine “1” at bottom-left, turbine “5” at top-left, turbine “21” at bottom-right, and turbine “25” at top-right.

245 corresponding optimal yaw angles. Farm power variations of up to 64 % occur between models. Even the Gaussian and the GCH wake models have an optimised farm power difference of 30 % despite the similarity in these models’ velocity deficit distribution (see Figure 1).

Arguably more important, from an implementation perspective, is the significant inter-model variation in optimal yaw angles. For the Jensen model, Figure 3 indicates non-zero optimal yaw angles for only the second (at 6°) and third rows (at 9°) of turbines. This counter-intuitive solution should be viewed in the context of a 0 % farm-power improvement, and the power-curve flatness described in Section 4.1. The Jensen model’s lack of sensitivity implies that moderate yaw angle variations are essentially indistinguishable to the optimiser and, consequently, a gradient-based algorithm such as SLSQP can terminate at any point in a large set of possible yaw angles: those obtained in Figure 3 are just one such near-optimal solution. Optimal yaw angle differences are also observed between the Gaussian model, which saturates the upper bound constraints at 25° for all but the final turbine row, and the GCH model, whose optimal angles are monotonically decreasing with turbine row. This is attributed to the GCH model’s ability to capture secondary steering effects. Finally, the Multi-zone model’s optimal yaw angles are non-monotonic with row, but also identify the non-physical setting of 8° for the final row of turbines (0° is optimal). Similar to the results for the Jensen model, this indicates optimisation sensitivity to initial yaw angle.

To further explore the sensitivity to initial yaw angles between models, Figure 4 shows the results of 10 SLSQP optimisation runs, each initialised by sampling the 25 initial yaw angles from a uniform distribution on $[-25^\circ, +25^\circ]$, with these initial conditions then used for all four models. All models exhibit strong sensitivity to initial yaw angles. Interestingly, for Jensen and Gaussian models, while there is almost no variation in normalised farm power improvement, there is high variability of the optimal decision variables. As discussed in Section 4.1, for Jensen, this arises due to cost-function flatness, and for the Gaussian model, due to the multi-modality of its objective function. A further consequence of multi-modality, which can be seen in Figure 4-c, is that yaw constraints may be saturated at either boundary (i.e. at $+25^\circ$ or at -25°), rather than uniformly at $+25^\circ$ as in the single case shown in Figure 1. Further, this column-wise switching of optimal yaw angle signs has a negligible

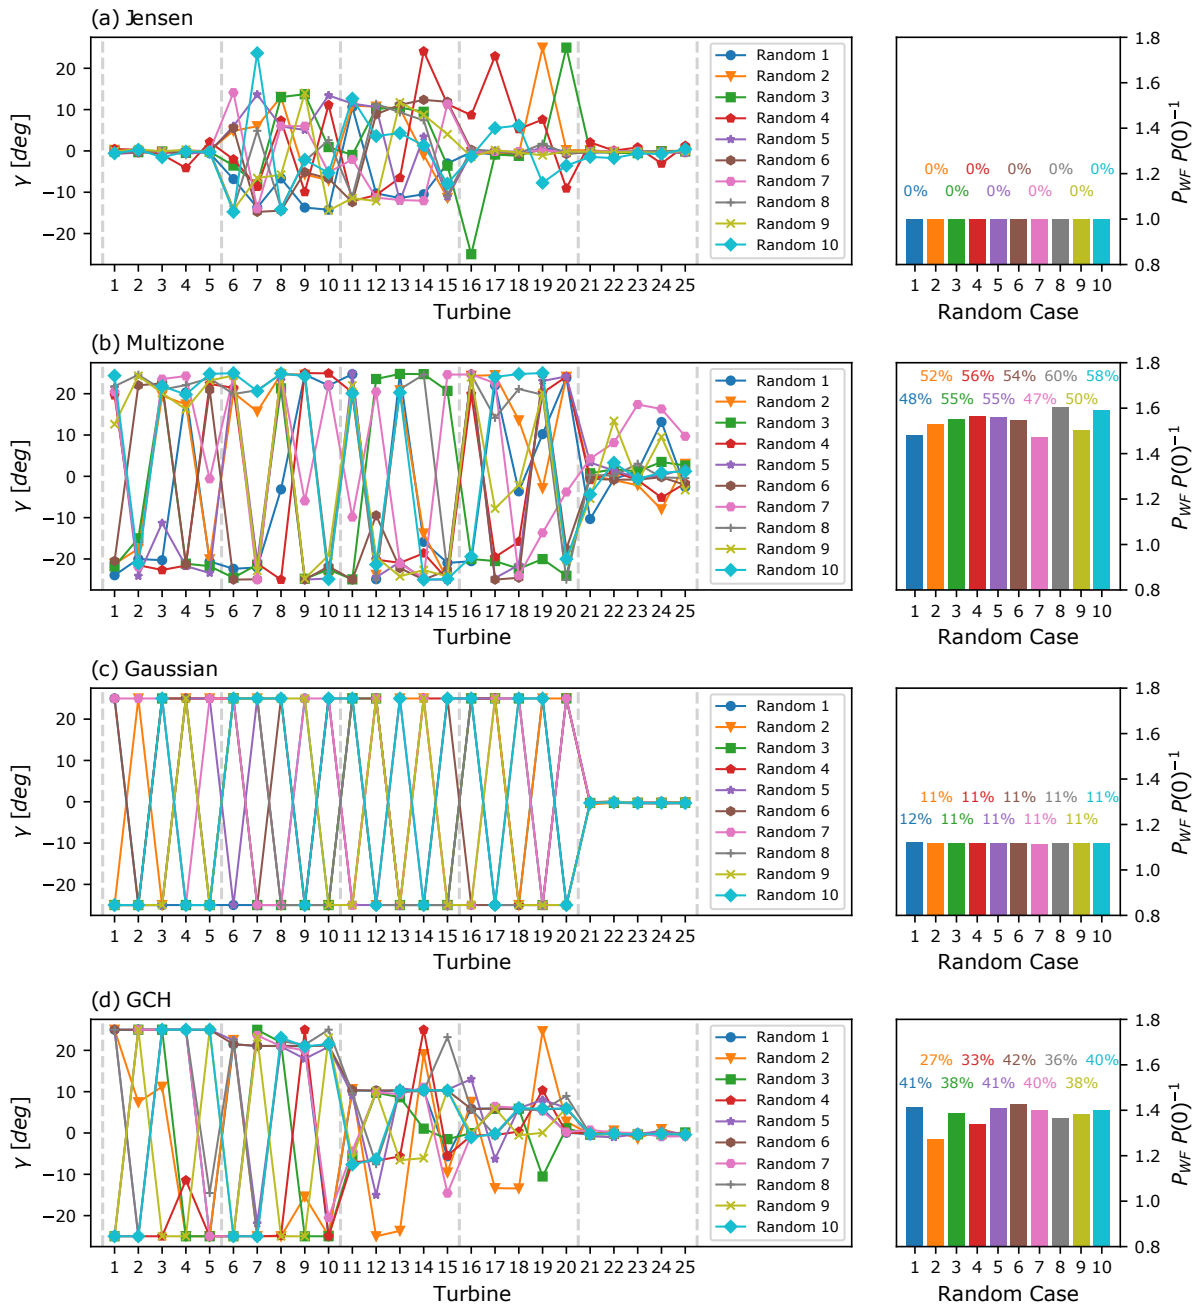


Figure 4. Initialisation sensitivity of a 5×5 wind farm SLSQP optimisation for (a) Jensen, (b) Multi-zone, (c) Gaussian, and (d) GCH wake models. Left plot: optimal set-points per turbine for the random sets “Random” of initial conditions. Dashed lines delimit wind farm rows. Right plot: resulting objective function (normalised farm power) per initialisation case. Turbine naming convention: turbine “1” at bottom-left, turbine “5” at top-left, turbine “21” at bottom-right, and turbine “25” at top-right.



effect on the predicted farm power improvement. Such sign inconsistency is clearly undesirable for practical implementation of wake steering.

For the Multi-zone and GCH models, sensitivity of both the normalised farm power (of up to 13 % and 15 %, respectively) and of the optimal yaw angles can be observed in Figure 4. The Multi-zone model (Figure 4-b) is highly sensitive to initialisation, caused by its highly non-linear and multi-modal farm power function—discussed in more detail in Section 4.2.2—which arises due to this model’s definition of splitting the wake into various zones, with sharp interfaces between each zone (see Figure 1-b). The GCH model (Fig. 4-d) has distinct local maxima, which is clear from the large differences observed in the normalised power increase across the 10 optimisation runs. Similar to the Gaussian model, multi-modality causes column-wise variation in the sign of the optimal decision variables in a number of runs (e.g., run “Random 10”). However, due to the fact that the GCH model also captures secondary-steering and wake rotation effects, such sign switches now also correspond to sub-optimal *local* maxima, with an observed range of 15 % in normalised farm power for even the small number of cases considered here.

To investigate initialisation sensitivity in more detail, we study one additional SLSQP optimisation run—denoted Test Case 1—for each of the considered models, again initialised by sampling 25 turbine yaw angles from a uniform distribution on $[-25^\circ, +25^\circ]$. The resulting optimal yaw angles are shown in Figure 5-left, while the streamwise velocity at hub height for the optimal yaw configuration is shown in Figure 5-right. It is clear that the initial yaw angles (blue markers) directly influence the optimal yaw angles (red markers) for the Jensen, Multi-zone and Gaussian models, both in their sign and absolute value. For the Gaussian model, for example, the sign of the initial angle for each turbine directly determines the sign of its optimised angle, which either saturates the constraints or is zero if the turbine is in the final row. Figure 5-right shows the corresponding, and practically undesirable, within-column yaw angle oscillations and their effect on the streamwise velocity fields.

For the GCH model in Figure 5-d, the sign influence of initial yaw angles is still apparent, although less obvious than for the other models. As discussed above, the resulting within-column oscillations of optimal yaw angles correspond to local, not global, maxima. To quantify this observation, consider the five columns of turbines in Figure 5-d-right for the GCH-based optimisation. The upper column, at $y \approx 2,500$ m, exhibits no within-column sign fluctuations of optimal yaw angle. This is in contrast to the other four columns, which extract, respectively, 13 %, 2 %, 16 %, and 7 % less power than the upper, sign consistent, column. Understanding the origin of these local maxima, and why the gradient-based optimiser SLSQP identifies different maxima for different initial yaw values requires a more careful analysis of the underlying objective function geometries.

4.2.2 Objective function geometry

Figure 6 shows, for each model, two planar slices through the objective functions $G(\gamma)$ for Test Case 1. Each is obtained by varying the yaw angle of a single chosen turbine (as indicated in each subfigure), with each presented slice also containing the optimal yaw angles obtained for each model. All four models exhibit local extrema, confirming the multi-modal nature of all farm power functions. Even the Gaussian model (Figures 6-c,g), which has the smoothest objective function, has local maxima due to the optimisation constraints. That the obtained optimum in this case (shown by red markers) is *not* global is

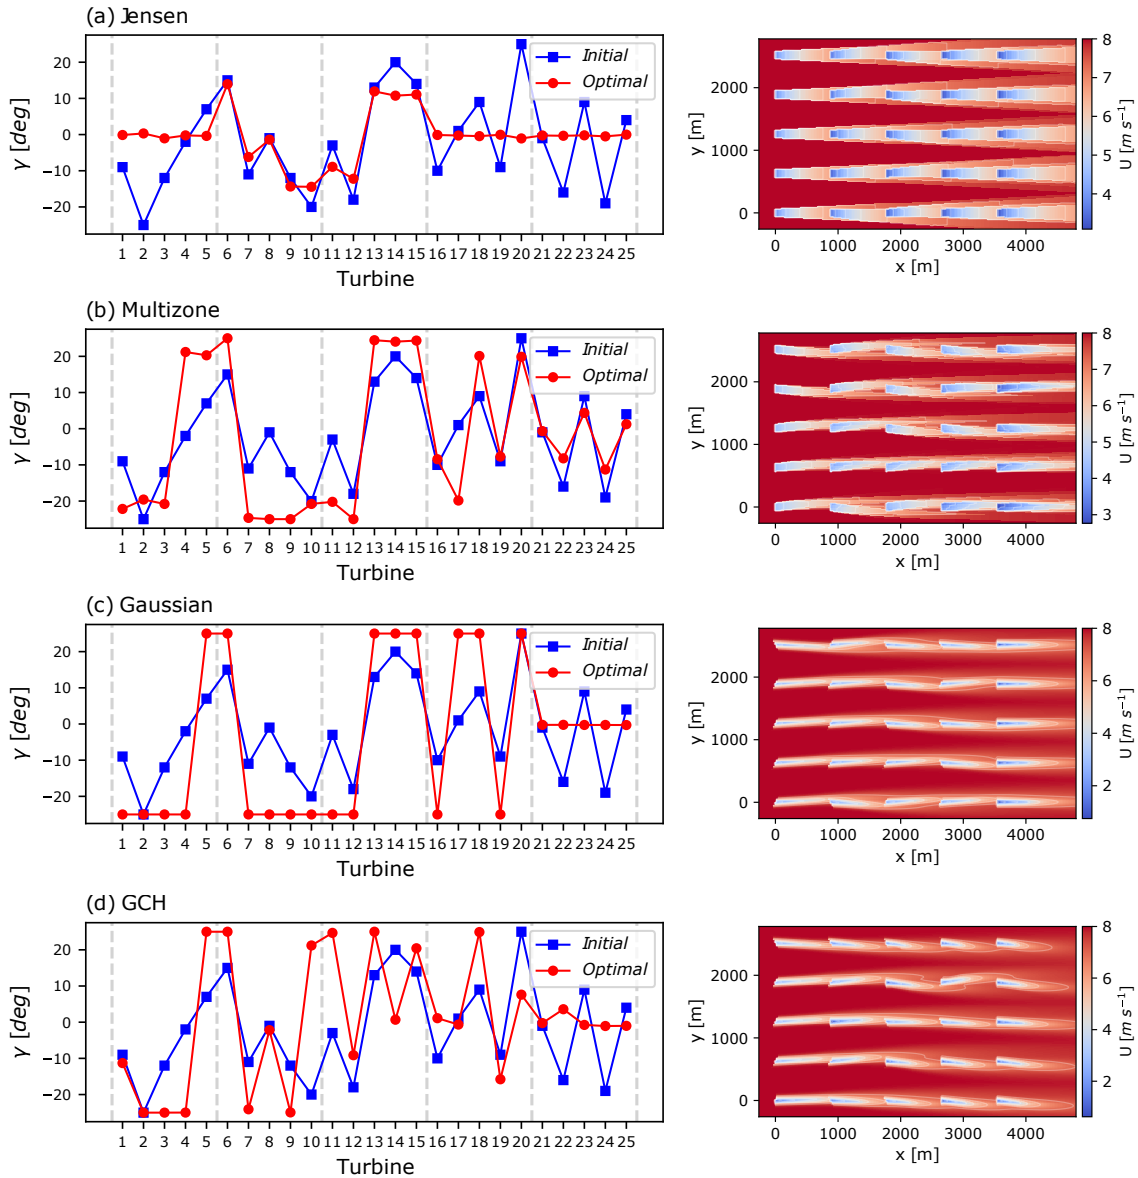


Figure 5. “Test Case 1” SLSQP optimisation results for (a) Jensen, (b) Multi-zone, (c) Gaussian, and (d) GCH wake models. Left plot: optimised yaw angle per turbine given “Test Case 1” random set of initial points. Dashed lines delimit wind farm rows. Right plot: streamwise velocity at hub height for optimised yaw set-points. Turbine naming convention: turbine “1” at bottom-left, turbine “5” at top-left, turbine “21” at bottom-right, and turbine “25” at top-right.

now transparently explained by the fact that the initial yaw angle of turbine T6 is positive, while that of T10 is negative (see Figure 5-c).

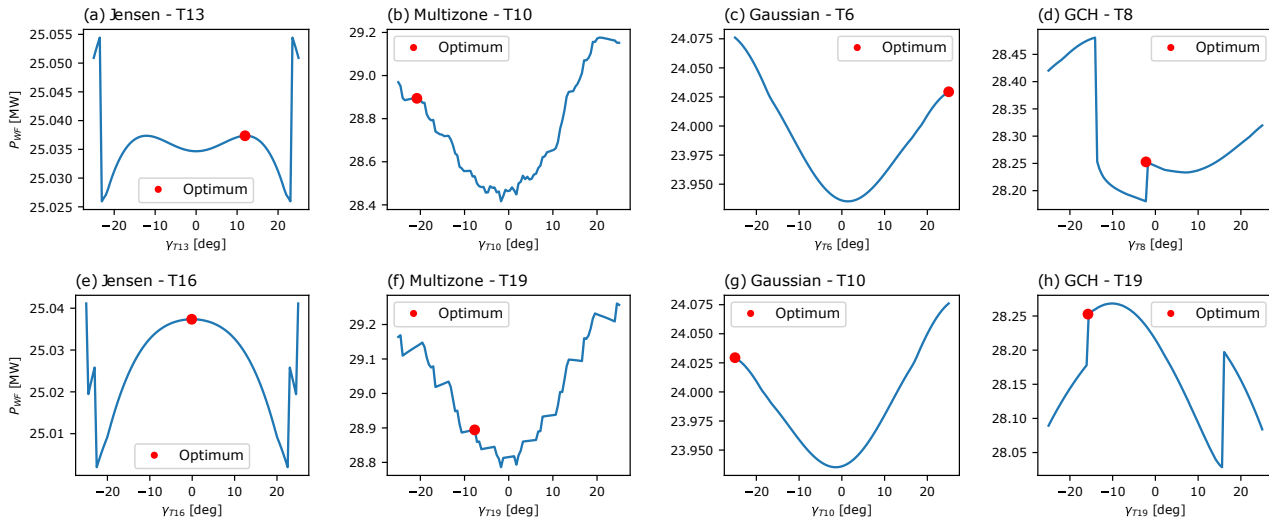


Figure 6. Wind farm power for (a,e) Jensen, (b,f) Multi-zone, (c,g) Gaussian, and (d,h) GCH wake models as a function of a single turbine yaw angle. Turbines yaw settings range $\pm 25^\circ$, with SLSQP optimisation solutions as starting points (illustrated in red for the swept turbine “T”). Turbine naming convention: turbine “1” at bottom-left, turbine “5” at top-left, turbine “21” at bottom-right, and turbine “25” at top-right.

The remaining three models have objective functions with either discontinuities or discontinuous derivatives. This is particularly prominent for the Jensen and Multi-zone models in Figures 6-a,e and 6-b,f, respectively. This property is further exacerbated by the multiple wake interactions present for the 5×5 farm. As discussed in Section 3.1, a careful choice of optimisation parameter ϵ can help mitigate the effect of cost-function discontinuities. Since ϵ determines the step size used for finite-difference gradient approximations, using a larger value can numerically smooth the objective function. However, in this case, it can be seen that even a relatively large, and standard, choice of ϵ corresponding to a 2.5° step-size cannot overcome the observed cost-function roughness for Jensen, Multi-zone or GCH models (only local maxima are obtained). This suggests that an optimal value of ϵ may be highly problem dependent for non-smooth underlying models, and thus hard to identify a priori.

Finally, we comment on the discontinuities present in the GCH model. These seem counter-intuitive given the GCH model’s use of smooth velocity distributions, as shown in Figure 1-d. However, the sharp discontinuities apparent in Figures 6-d,h arise due to secondary-steering effects, where the wake of a downstream turbine in the GCH model can experience a sudden change in deflection due to the overlapping with the influence region of the vortices generated by one or more upstream turbines. Figure 7 confirms this behaviour by showing the flow field on either side of one jump discontinuity present in the cost-function slice of T8 in Figure 6-d. Here, a small 1° change in yaw angle of a single machine gives a 1 % loss in total farm power production.




Figure 7. Wind farm power function discontinuity for the GCH wake model. Left plot (a): wind farm power as a function of turbine “8” yaw angle, with the optimised solution as a starting point. Right plots: streamwise velocity at hub height of the third column of turbines for $\gamma_{T8} = -14.5^\circ$ (b) and $\gamma_{T8} = -13.5^\circ$ (c). Turbine naming convention starting from $x = 0$ and moving downstream: turbine “3”, turbine “8”, turbine “13”, turbine “18”, and turbine “23”.

The intrinsically multi-modal, and often rough, nature of farm power objective functions for the 5×5 farm presents a challenge for gradient-based optimisation. In section 4.3 we provide a solution to this problem for gradient-based optimisers. However, we first consider the possible advantage of using a global optimisation strategy and perform a statistical comparison of optimisation sensitivity between gradient-based and global optimisation approaches for the 5×5 farm power maximisation problem.

4.2.3 Sensitivity comparison of gradient-based and global optimisation algorithms for wake steering

We compare the performance of the SLSQP and TuRBO optimisers described in Section 3.1 in terms of the statistics of two properties: the normalised farm power increase, and the total number of wind farm evaluations required. Statistical analysis is required since SLSQP is highly sensitive to initial conditions (see Section 4.2.1), and TuRBO is stochastic by nature. For each of the four models, wake steering optimisation is repeated 50 times per optimiser, each initialised with an independent set of random initial conditions sampled from a uniform distribution on $[-25^\circ, +25^\circ]$. To ensure a fair comparison, SLSQP and TuRBO algorithms share the same sets of initial yaw angles. The maximum number of evaluations for TuRBO is set to 500, corresponding to the observed average number of objective function evaluations used by SLSQP. Finally, to limit the bias of extreme cases in the function evaluation statistics, each algorithm is assumed to be converged when 95 % of the actual wind farm power improvement is reached.

Figure 8’s upper row shows the comparison between SLSQP and TuRBO with the mean, standard deviation, minimum and maximum values of normalised farm power increase shown. Higher mean values are obtained for all models by the global TuRBO algorithm. Although the balance between exploration and exploitation can represent a challenge when using global optimisers, especially in highly dimensional search spaces, the multi-modal nature of the farm power function appears to give an advantage to TuRBO’s global approach for this example. Moreover, for all models except Jensen, a lower standard




Figure 8. Comparison of the objective function (a-d) and the number of evaluations (e-h) statistics between SLSQP and TuRBO optimisation algorithms for (a,e) Jensen, (b,f) Multi-zone, (c,g) Gaussian, and (d,h) GCH wake models. Mean values are illustrated with error bars for one standard deviation (red or green) and minimum and maximum values (black).

340 deviation in the obtained optimal farm power is found when using TURBO, indicating that it is less sensitive to yaw angle
 340 initialisation.

Statistics (mean, standard deviation, minimum and maximum values) of the number of evaluations required for each optimiser are shown in the lower row of Figure 8. For all the wake models, the TuRBO algorithm requires roughly the same or lower evaluations than SLSQP to achieve, on average, higher wind farm power improvements. Moreover, TuRBO exhibits lower or similar standard deviations than SLSQP, indicating a lower impact of initialisation on computational requirements.

345 Figure 9 presents three indicative cases to understand the comparison of SLSQP and TuRBO optimisers in the case of the GCH model. These are selected to be: the TURBO run achieving farm power production closest to the mean power increase, and the two SLSQP runs corresponding to the minimum and the maximum farm power increases achieved. Figure 9-a shows the objective function $G(\gamma)$ whenever a farm evaluation is evaluated by the respective optimiser. It is interesting that the best-performing SLSQP run (blue markers) converges after fewer iterations, and to a higher farm power, than the indicative
 350 TuRBO case presented and, in fact, outperforms all TuRBO runs (see Figure 8). This suggests that, if initialised in the region of attraction of a global (or near-global) maxima, a gradient-based optimiser may exhibit faster convergence than a generic global strategy.

The streamwise velocity flow fields at hub height for the three considered cases are also shown in Figure 9. Inter-column switches of optimal yaw angles are again present in the worst performing SLSQP case (green markers in Figure 9-a), which
 355 requires 88 % more objective function evaluations and has a 10 % drop in farm power compared to the best-performing

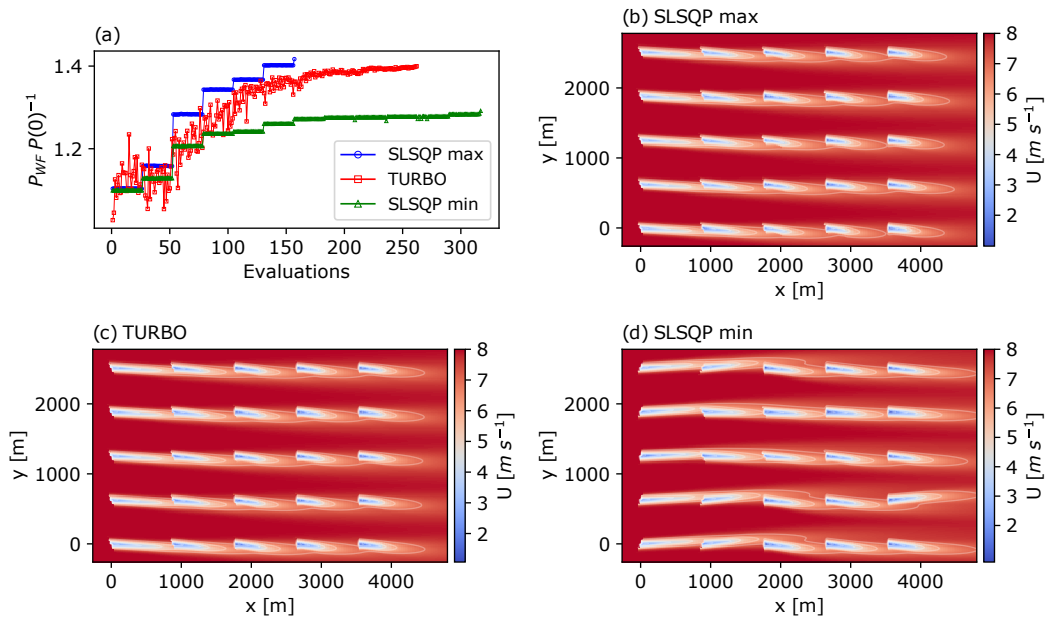


Figure 9. Comparison of optimisation results between SLSQP and TuRBO optimisation algorithms for the GCH wake model. Plot (a): objective function values at each evaluation. Remaining plots: resulting streamwise velocity at hub height with optimised yaw set-points for SLSQP maximum (b), average TuRBO (c), and SLSQP minimum (d) farm improvements.

SLSQP case (blue markers). This again highlights the challenge of multimodality for gradient-based approaches. Finally, it is interesting to note that while the optimal wake velocity fields in Figures 9-b and 9-c are visually similar, they are not identical: the TuRBO solution has 1.5 % lower farm power, and requires 55 % more evaluations to find it. This confirms a possible, albeit counter-intuitive, advantage of the rapid local convergence enjoyed by gradient-based optimisers. Indeed, the quasi-Newton algorithm employed in SLSQP is well-known (Deuffhard, 2011) to possess rapid locally-quadratic convergence rates, which may partially explain this observation.

In the following section, we consider a further increase in problem complexity, and show how using simple constraints can reduce the initial yaw angle sensitivity of the SLSQP algorithm to enable rapid, and robust, convergence to desirable local maxima.

365 4.3 Multiple-turbine case: the Horns Rev Wind Farm

In this section, the results for farm power maximisation via wake steering optimisation are presented for the well-known Horns Rev wind farm. Dou et al. (2020) have previously implemented a wake steering optimisation for this farm layout, finding optimal yaw angles whose signs exhibit high column-inconsistency, reflecting the observations made for the 5×5 layout considered in Section 4.2.3. To the best of the authors' knowledge, the work by Dou et al. (2020) is the only published study in which wake steering optimisation is performed on the Horns Rev farm layout.



The first objective of this high-complexity optimisation case is to confirm whether wake steering sensitivities uncovered in the low and medium complexity cases transfer to a realistic setup. The second objective is to add optimisation constraints to try to mitigate initial yaw angle sensitivity and to achieve coherent and interpretable optimal yaw settings that could eventually be implementable in the field. In this section, we employ only the SLSQP optimisation algorithm, using the parameters in Table 2, motivated by the observation from Section 4.2.3 that SLSQP can exhibit rapid convergence if initialisation sensitivity is appropriately mitigated.

For sensitivity mitigation, we consider two additional sets of optimisation constraints. The constraint “C1” refers to the case in which the yaw angles of every turbine are constrained to be positive. This is achieved by modifying the variable bounds from $[-25^\circ, +25^\circ]$ to $[0^\circ, +25^\circ]$. A second set of constraints, “C2”, corresponds to the case where the yaw angles of each turbine (excluding turbines in the most upstream row) are forced to be equal to or lower than the yaw angles of the upstream turbines in the same column, and is implemented in SLSQP using linear inequality constraints.

The investigated cases are limited to a fully aligned layout, as it is the predominant condition and the one holding the largest potential for the implementation of wake steering strategies. The analysis performed is statistical, where each of the multiple optimisation cases per wake model is initialised with an independent set of random initial conditions sampled from a uniform distribution on $[-25^\circ, +25^\circ]$. For clarity throughout the section, optimisation performed using only the original yaw angle bounds of $[-25^\circ, +25^\circ]$ is referred to as “nominal”. Optimisation cases with additional constraints are indicated as “constrained”. Note that for the additional constrained Horns Rev cases, a 2.5° step-size corresponds to $\text{eps} = 0.1$ due to the modified variable bounds $[0^\circ, +25^\circ]$.

Optimal yaw configurations for a single SLSQP optimisation with a unique random set of initial conditions in the nominal, “C1”, and “C1+C2” optimisation cases are presented in Fig. 10. In line with the results of Dou et al. (2020), the first row of subfigures shows that, in the nominal case, the optimal yaw settings obtained for the four wake models are not practical, with frequent switches between large positive and negative yaw angles within any given row or column of the wind array. The only common trend is that the final downstream row of turbines has optimal yaw angles close to zero, even with the Jensen model. The key message is that, without a sensitivity mitigation strategy, the proposed yaw angle settings cannot be implemented in real life due to obvious operational constraints.

The introduction of additional constraints clearly improves the identification of the optimal variables. For all models, optimal yaw angles represent a more interpretable and practical solution. Overall, similar performance is achieved with the two different constrained approaches, “C1” and “C1+C2”. This successfully proves how the use of simple constraints, motivated by an understanding of the optimisation problem, can lead to considerable improvements in optimisation performance. The “C1” constraint forces the optimiser to avoid the maxima corresponding to within-column yaw angles with alternating signs, of particular importance when such maxima are sub-optimal for models with power asymmetry due to yaw (Multi-zone and GCH). The “C2” constraint, limited to “C1+C2” cases, effectively decreases initialisation sensitivity by imposing column equal or monotonic decreasing optimal yaw angle solutions. This behaviour can be clearly observed in the “C1+C2” Jensen optimisation case (Fig. 10-i). For the GCH model, which captures secondary steering effects, introducing constraints leads to a smoother decrease in optimal decision variables from row to row, and to the identification of an optimal decrease rate.

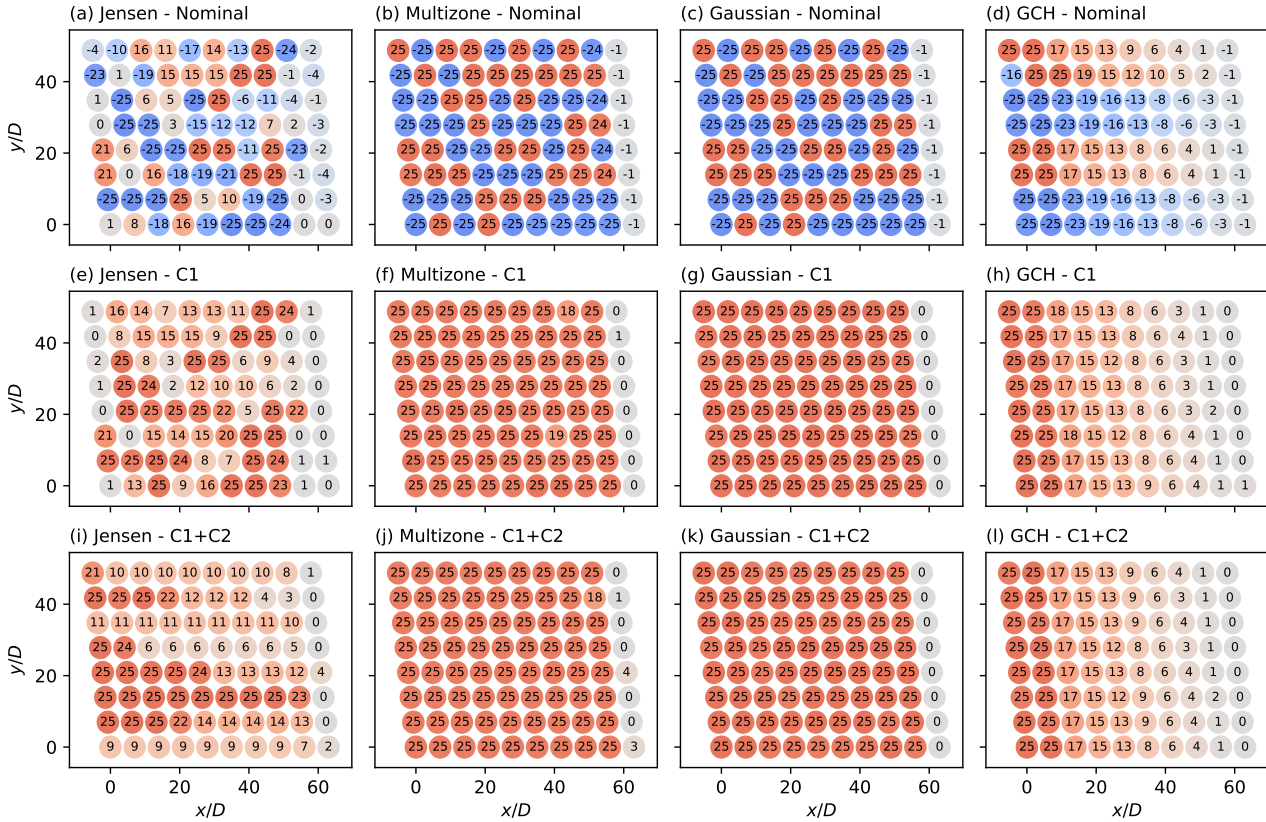


Figure 10. Comparison of the optimal yaw setting for a single optimisation between nominal, “C1”, and “C1+C2” constrained SLSQP optimisation cases for (a,e,i) Jensen, (b,f,j) Multi-zone, (c,g,k) Gaussian, and (d,h,l) GCH wake models. Red and blue indicate positive and negative yaw angles, respectively.

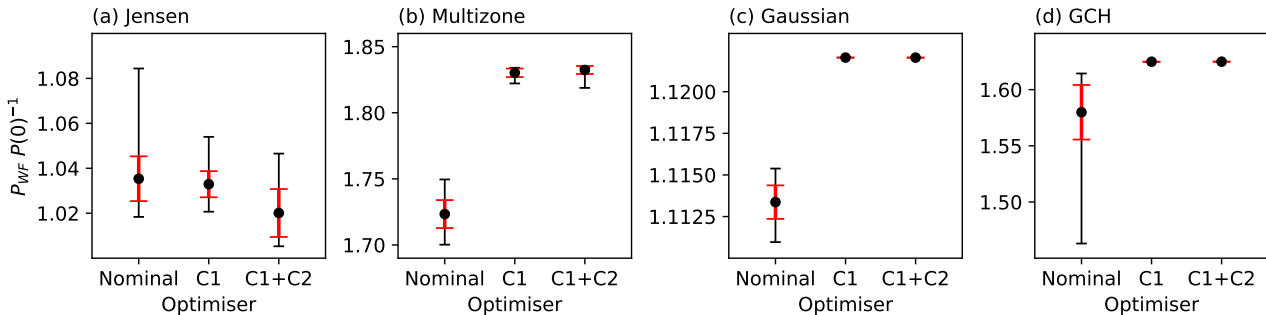


Figure 11. Comparison of the objective function statistics between nominal, “C1”, and “C1+C2” constrained SLSQP optimisation cases for (a) Jensen, (b) Multi-zone, (c) Gaussian, and (d) GCH wake models. Mean values are illustrated with error bars for one standard deviation (red) and minimum and maximum values (black).

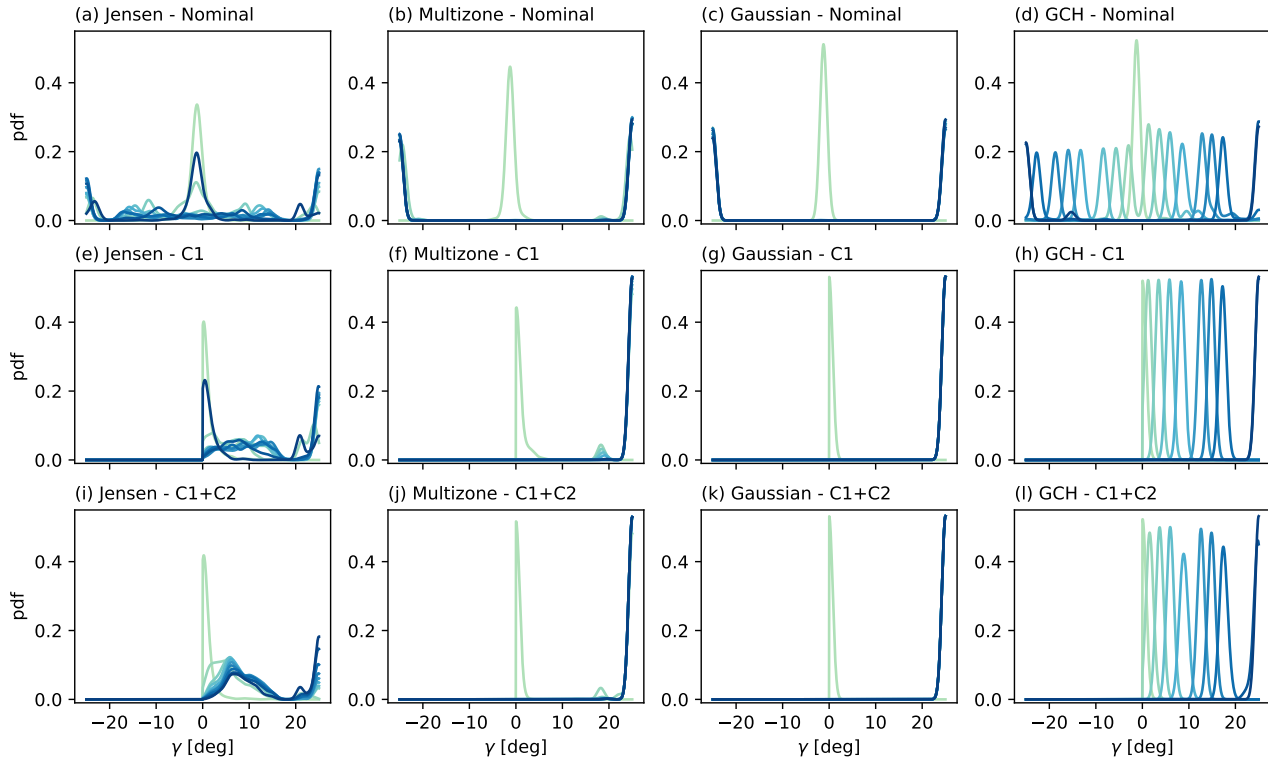


Figure 12. Comparison of the optimal yaw setting statistics between nominal, “C1”, and “C1+C2” constrained SLSQP optimisation cases for (a,e,i) Jensen, (b,f,j) Multi-zone, (c,g,k) Gaussian, and (d,h,l) GCH wake models. Column-averaged probability density functions of the optimal yaw angles are shown, with a transition from dark blue (most upstream row) to light green (most downstream row).

Figure 11 presents the statistical results for farm power improvement in the nominal, “C1”, and “C1+C2” optimisation cases. For the Multi-zone, Gaussian, and GCH models, a significant reduction in initialisation sensitivity can be clearly observed when constraints are imposed, both in terms of standard deviation and minimum and maximum values. For the GCH model, dependency on initial conditions in the nominal case leads to sub-optimal solutions with potential farm power losses up to 15 %. With the introduction of constraints, these losses are below 0.1 %. Further, a higher mean farm power improvement is experienced in the constrained cases for all models except the Jensen one. In the Multi-zone case, avoiding sub-optimal solutions due to constraints leads to a 12 % increase in mean farm power improvement. Regarding the Jensen model, no optimisation setting is found to predict a significant farm power improvement for the investigated cases. Although successfully reducing initialisation sensitivity both in terms of standard deviation and minimum and maximum values, constrained cases lead to a reduction in the maximum farm power improvement, as a result of the flat nature of Jensen’s objective function. Finally, consistent with the conclusions for the medium-complexity 5×5 case, an overall strong dependency of wake steering strategies to wake model choice is observed, exhibiting significant discrepancies in farm power improvement.



Figure 12 shows the statistical results for the optimal yaw settings in the nominal, “C1”, and “C1+C2” optimisation cases. For all wake models, the row-averaged probability density functions indicate an overall reduction in optimal yaw angle dependency to initial conditions for the constrained cases. Higher probability density function magnitudes can be observed, with more than double the values for some row distributions. For the Multi-zone and Gaussian models, optimal yaw settings in the nominal case tend to converge at the optimisation bounds, as shown in Fig. 5 and confirmed by bi-modal probability density functions in Fig. 12-b and 12-c. Caused by initialisation sensitivity and lack of constraints, this undesirable behaviour of bi-modality is easily removed in the “C1” and “C1+C2” optimisation cases. For the GCH model, bi-modal probability density functions can also be identified in the nominal optimisation results. However, the peaks of the modes exhibit a more complex pattern depending on the farm row (Fig. 12-d). As for the previous models, bi-modality is removed by the introduction of constraints. In the Jensen case, nominal results are characterised by low and row-inconsistent probability density function values throughout the whole yaw solution range (Fig. 12-a). Constraints significantly reduce initialisation sensitivity by restricting the solution space to the positive region and, only for the “C1+C2” case, by enabling the optimisation algorithm to identify more consistent optimal decision variables. Finally, consistently with the discussion for the 5×5 farm layout in Section 4.2, significant overall discrepancies between wake models can be identified in terms of optimal yaw angles.

The presented results indicate that the simple strategy proposed for SLSQP sensitivity mitigation is effective for this realistic and complex wake steering optimisation problem. The enforcement of constraints successfully decreases the impact of initial conditions both in terms of the objective function and optimal decision variables, while achieving generally higher farm power improvements and more consistent optimal yaw settings.

5 Conclusions

A detailed investigation of the sensitivity of wake steering optimisation for increased power output of wind farms was carried out in this study. Sensitivity to the choice of analytical wake models, optimisation algorithm, and different operating set-points was assessed in terms of variability of both optimised farm power and optimal yaw angles. The study was performed with four different analytical wake models, with a gradient-based and a global optimisation algorithm. Three wind farm layouts were investigated: two generic layouts with 2×1 and 5×5 turbines, and a layout corresponding to the Horns Rev wind farm with 80 turbines. This hierarchy of test-cases with increasing complexity facilitated the understanding of the different sensitivities present in wake steering optimisation. By conducting analysis on initialisation sensitivity, the nature of the farm power function, and optimiser statistical performance, this investigation provided insight into the impact of modelling, algorithmic, and parametric choices on wake steering strategies to increase the power output of wind farms.

For the 2×1 layout, optimisation was performed by a simple parametric sweep, removing any sensitivity to algorithmic choice or implementation. The results indicated a strong sensitivity of wake steering strategies to the choice of analytical wake model and highlighted the role of flatness and multi-modality of the objective function in wake steering optimisation performance.



450 In the medium-complexity 5×5 case, optimisation was first performed using the gradient-based SLSQP optimisation algo-
rithm. High optimisation sensitivity to wake model choice, both in terms of objective function and optimal decision variables,
was observed. All models exhibited strong variability of the objective function and optimal yaw angles to initial conditions,
with sub-optimal solutions leading to potential power improvement losses. A comparison between the gradient-based SLSQP
and the global optimiser TuRBO showed that, overall, TuRBO required fewer wind farm evaluations and achieved larger im-
455 improvements in wind farm power production, with a reduced sensitivity to initial conditions. Although the multi-modal and
discontinuous nature of the farm power function strongly affected the SLSQP performance, the results also highlighted that
rapid convergence could be achieved if initialisation sensitivity was reduced, motivating the sensitivity mitigation strategy
proposed in the final test-case considered.

In a final, higher complexity, optimisation was performed with the SLSQP algorithm for the farm layout corresponding to
460 the Horns Rev wind farm (80 turbines). Optimal yaw angles were found to exhibit highly column-inconsistent signs, making
them impractical to implement in practice. It was found that the introduction of constraints during the optimisation dramatically
decreased the sensitivity to initial conditions while achieving more consistent and interpretable yaw settings. This work, there-
fore, demonstrates that the sensitivity of wake steering optimisation can be mitigated either by the use of global optimisation
algorithms or by the addition of simple constraints to gradient-based algorithms. The success of a constrained gradient-based
465 approach is especially promising, particularly with a view towards adaptation to new problems, due to the simplicity and
transparency of gradient-based algorithms in contrast to more ad-hoc global strategies.

Future work will look at sensitivity of wind farm wake steering optimisation for a range of wind directions and inflow speeds,
and investigate the impact of atmospheric conditions for more complex optimisation problems, in addition to identifying appro-
priate optimisation constraints to enable sensitivity mitigation in each setting. Future research will also focus on determining
470 wake steering sensitivities in an online control framework, including state estimation and parameter tuning algorithms.

Acknowledgements. The authors would like to acknowledge the Department of Aeronautics, Imperial College London, for supporting this
work funded by a doctoral studentship.

Appendix A

A1 Wake models

475 The most relevant parameter in analytical wake models is the streamwise velocity u (u_w in Eq. (1)) in the downstream region
of a turbine subject to a mean free stream velocity U_∞ . Due to turbine operation and the consequent wake streamwise velocity
deficit δu , the streamwise velocity is defined as

$$u = U_\infty(1 - \delta u) . \tag{A1}$$



Considering the unsteadiness of the wake phenomenon and its evolution in three-dimensional space, in principle, all terms
480 in Eq. (A1) are a function of space and time (x, y, z, t) . However, the main assumption of analytical models is the description
of a steady-state, where time dependency is ignored, and the focus is on equilibrium conditions. The model formulations in
this section are presented based on velocity deficit $\delta u(x, y, z)$.

A1.1 Jensen Model

The Jensen model (Jensen, 1983) is a simple steady wake model widely used in the industry. It is derived by applying mass
485 conservation to a control volume in the wake region of the wind turbine and relating the velocity deficit just behind the rotor to
the thrust coefficient C_T using Betz theory. Hence, turbines are assumed to be actuator disks with rotor diameter D and thrust
coefficient $C_T = 4a(1 - a)$, where a is the turbine induction factor. The Jensen model is based on two main assumptions: the
conservation of the cross-stream integral of the streamwise velocity deficit as the wake linearly expands downstream and the
velocity deficit being simply a function of the downstream distance x . The first main assumption implies a quadratic decay of
490 the velocity deficit with the linear expansion of the wake. The second one implies a uniform velocity deficit in the wake, hence
an axisymmetric wake with a well-defined edge (see Fig. 1-a).

The streamwise velocity deficit induced by a turbine with diameter D , assumed to be operating at an induction factor a , can
be expressed as

$$\delta u(x, r) = \begin{cases} 2a \left(\frac{D}{D+2kx} \right)^2, & \text{if } r \leq \frac{D+2kx}{2} \\ 0, & \text{otherwise} \end{cases}, \quad (\text{A2})$$

495 where a cylindrical coordinate system is defined with the origin at the rotor hub of the first upstream turbine. Radial and
streamwise distances are expressed by r and x , respectively, and k represents the dimensionless expansion coefficient.

The Jensen model is based on a steady description of the wake where turbines are modelled as actuator disks with uniform
loading, and no notion of added turbulence intensity due to upstream turbines' operation is included. Moreover, it does not
conserve momentum, it is limited to far wake predictions, and no calculation involving cross-stream and vertical velocity
500 components is included. Despite these limitations, the Jensen model is simple and inexpensive. It can be used for control and
optimisation studies, and it can provide valuable insights into power production for large wind farm layouts in normal operating
conditions.

A1.2 Multi-zone Model

The Multi-zone model, developed by Gebraad et al. (2014), is a modification of the Jensen model to better describe wake
505 velocity profiles and partial wake overlapping, especially in yawed conditions. Within a turbine wake, three zones q are defined:
near wake zone ($q = 1$), far wake zone ($q = 2$), and mixing wake zone ($q = 3$). They are assumed to expand linearly with
downstream distance x . Different expansion rates are determined based on the tuned model parameters k_e and $m_{e,q}$. The wake
diameter $D_{w,q}$ for each zone is defined as

$$D_{w,q}(x) = \max(D + 2k_e m_{e,q} x, 0). \quad (\text{A3})$$



510 The velocity deficit in the three wake zones is assumed to decay quadratically with the downstream distance rather than being directly related to the wake expansion as in the Jensen model. Due to the presence of different wake zones, a smoother transition from the wake centre to freestream velocity is achieved (see Fig. 1-b), where the velocity deficit spatial variables are the radial position from the wake centre r and turbine downstream distance x . The wake remains axisymmetric. Considering a turbine assumed to be operating at an induction factor a and yaw angle γ , the mean streamwise velocity deficit $\delta u(x, r)$ can be

515 defined as

$$\delta u(x, r) = 2ac(x, r), \quad (\text{A4})$$

where $c(x, r)$ is the wake decay coefficient

$$c(x, r) = \begin{cases} c_1 & \text{if } r \leq D_{w,1}/2 \\ c_2 & \text{if } D_{w,1}/2 < r \leq D_{w,2}/2 \\ c_3 & \text{if } D_{w,2}/2 < r \leq D_{w,3}/2 \\ 0 & \text{if } r > D_{w,3}/2 \end{cases}. \quad (\text{A5})$$

The local wake decay coefficient $c(x)$ for each wake zone q is expressed by

520
$$c_q(x) = \left[\frac{D}{D + 2k_e m_{U,q}(\gamma) x} \right]^2, \quad (\text{A6})$$

where k_e parameter determines wake zones expansion and velocity deficit recovery, while $m_{U,q}$ is an empirically derived coefficient. For each wake zone, q , $m_{U,q}$ is defined as

$$m_{U,q}(\gamma) = \frac{M_{U,q}}{\cos(a_U + b_U \gamma)}, \quad (\text{A7})$$

with the empirical parameters $M_{U,q}$ (ensuring outer zones faster recovery), a_U , and b_U tuned with high-fidelity simulations.

525 The Multi-zone wake model is a computationally inexpensive model, suitable for control and optimisation studies, including yaw applications. However, it only describes equilibrium conditions, and it is limited to far wake predictions. In addition, it does not exhibit any sensitivity to inflow turbulence intensity, and it does not include considerations for added turbulence intensity by upstream turbines. Moreover, it involves many tuned empirical parameters, decreasing the model confidence for a various range of operating conditions. Finally, it does not explicitly conserve momentum.

530 A1.3 Gaussian Model

The Gaussian wake model was originally developed by Bastankhah and Porté-Agel (2014) and recently improved by various studies in the literature (Abkar and Porté-Agel, 2015; Bastankhah and Porté-Agel, 2016; Niayifar and Porté-Agel, 2016; Dilip and Porté-Agel, 2017). This steady wake model consists of a mass and momentum conserving formulation based on a simplification of the Navier-Stokes equations. A Gaussian distribution on the streamwise velocity deficit is imposed by applying the

535 self-similarity theory used in shear flows (see Fig. 1-c). Moreover, a linear wake expansion is assumed. The model accounts for atmospheric stability and added turbulence intensity due to upstream turbines. Defining a three-dimensional coordinate



system at the rotor hub of the turbine, the streamwise velocity deficit in the far wake, dependant on the three spatial dimensions (x, y, z) , is defined as

$$\delta u(x, y, z) = C e^{-(y-\delta)^2/2\sigma_y^2} e^{-(z-z_h)^2/2\sigma_z^2}, \quad (\text{A8})$$

540 where C is the streamwise velocity deficit at the wake centre, δ the wake deflection (refer to Appendix-A2), z_h the hub height of the turbine, and σ_y and σ_z the standard deviations of the Gaussian velocity deficit at each streamwise location, indicating the width of the wake in the cross-stream and vertical direction, respectively. The general concept is to apply a Gaussian distribution to the streamwise velocity deficit at the wake centre, given linearly increasing cross-stream and vertical wake widths and a cross-stream deflection due to turbine yaw misalignment. The wake centre velocity deficit C is defined as

$$545 \quad C(x) = 1 - \sqrt{1 - \frac{C_T \cos \gamma}{8(\sigma_y \sigma_z / D^2)}}, \quad (\text{A9})$$

while the standard deviations σ_y and σ_z as

$$\sigma_y(x) = k_w(x - x_0) + \sigma_{y0}, \quad (\text{A10})$$

$$\sigma_z(x) = k_w(x - x_0) + \sigma_{z0}. \quad (\text{A11})$$

550 All quantities with the subscript “0” represent wake properties at the far wake onset (the end of the near wake) and depend on turbine thrust coefficient C_T and turbulence intensity I (Eq. (A18)). Near wake length is computed with

$$x_0 = \frac{D \cos \gamma (1 + \sqrt{1 - C_T})}{\sqrt{2} (4\alpha I + 2\beta (1 - \sqrt{1 - C_T}))}, \quad (\text{A12})$$

where α and β are the tuning parameters governing the influence of turbulence and thrust coefficient on the near wake end location, respectively. Far wake onset wake widths in the vertical (σ_{z0}) and cross-stream (σ_{y0}) directions are defined as

$$555 \quad \sigma_{z0} = 0.5D \sqrt{\frac{u_R}{U_\infty + u_0}}, \quad (\text{A13})$$

$$\sigma_{y0} = \sigma_{z0} \cos \gamma, \quad (\text{A14})$$

with the velocity at the turbine rotor u_r and the velocity in the near wake u_0 (assumed as constant) computed with

$$u_R = \frac{U_\infty C_T}{2(1 - \sqrt{1 - C_T})}, \quad (\text{A15})$$

$$560 \quad u_0 = U_\infty \sqrt{1 - C_T}. \quad (\text{A16})$$

For additional details on the derivations of far wake onset quantities refer to Bastankhah and Porté-Agel (2016).

In the Gaussian model formulation, k_w represents the main parameter. It defines the rate of linear wake expansion, and it accounts for ambient and added turbulence intensities. It is expressed by

$$565 \quad k_w = k_a I + k_b \quad (\text{A17})$$



where k_a and k_b are tuning parameters defining the weight of turbulence intensity and fundamental wake recovery, and I represents the turbulence intensity, expressed as

$$I = \sqrt{\sum_{j=0}^N (I_j^+)^2 + I_0^2}. \quad (\text{A18})$$

N is the total number of turbines influencing the reference turbine, and I considers both I_0 , the ambient turbulence intensity, and I_j^+ , the added turbulence intensity due to an upstream turbine j . Given an overlap area $A_{overlap}$ between the wake of an upstream turbine j and the turbine of interest, I_j^+ is defined by the empirical expression by Crespo et al. (1999)

$$I_j^+ = A_{overlap} \left(0.5a_j^{0.8} I_0^{0.1} (x/D_j)^{-0.32} \right), \quad (\text{A19})$$

where the coefficients are tuned with LES simulations by King et al. (2021). For additional details on added turbulence calculations, refer to Niayifar and Porté-Agel (2016).

Although the wake model complexity is increased, the Gaussian wake model is still suitable for control and optimisation applications. Additional considerations include ambient and added turbulence intensity, and the wake model explicitly conserves momentum. The main limitations are the formulation based on a free shear approximation of the Navier-Stokes equations and the inaccurate near wake predictions. Moreover, it does not compute cross-wise and streamwise velocity components, critical for modelling wake steering effects, and it relies on multiple empirical coefficients, decreasing the range of conditions for suitable predictions.

A1.4 Gauss-Curl Hybrid Model

The Gauss-Curl Hybrid (GCH) model (King et al., 2021) extends the modelling capabilities of the Gaussian model by implicitly modelling wake rotation and counter-rotating vortices due to yaw misalignment, allowing the description of wake asymmetry, added yaw-based wake recovery, and secondary-steering effects. The Gaussian model formulation is improved by introducing analytical approximations based on the Curl model (Martínez-Tossas et al., 2019), a recently developed linearised RANS formulation with an increased computational complexity of around 1000x compared to the Gaussian model. This section provides an overview of the GCH wake model. For additional details on the calculations, refer to King et al. (2021).

Wake rotation is included by modelling a Lamb-Oseen vortex, where its circulation strength is dependent on the turbine axial induction factor and tip-speed ratio. The counter-rotating vortex system due to turbine yaw misalignment is modelled as two single vortices released at the top and the bottom of the rotor, with a vortex strength dependent on turbine thrust coefficient and yaw angle. For each vortex described, cross-stream V and vertical W velocity components are calculated and linearly combined. Ground effects are considered by adding mirrored vortices below the ground. Vortices combination of the velocity components can be expressed by

$$V_{wake}(x, y, z) = V_{top} + V_{bottom} + V_{wake\ rotation}, \quad (\text{A20})$$

$$W_{wake}(x, y, z) = W_{top} + W_{bottom} + W_{wake\ rotation}.$$



595 Vortices dissipation with downstream distance is computed as

$$V(x, y, z) = V_{\text{wake}}(x, y, z) \left(\frac{\epsilon^2}{4\nu_T \frac{(x-x_0)}{U_\infty} + \epsilon^2} \right), \quad (\text{A21})$$

$$W(x, y, z) = W_{\text{wake}}(x, y, z) \left(\frac{\epsilon^2}{4\nu_T \frac{(x-x_0)}{U_\infty} + \epsilon^2} \right),$$

where ϵ is the vortex core size ($\epsilon = 0.2D$ in this study), and ν_T is the turbulent viscosity, defined with a mixing length model.

As presented in King et al. (2021), added wake recovery due to turbine yaw misalignment is considered by updating turbulence intensity with an additional component I_{mix} due to yaw enhanced mixing. The updated turbulence intensity I_{updated} is expressed as

$$I_{\text{updated}} = I + \phi I_{\text{mix}}, \quad (\text{A22})$$

where the tuning parameter ϕ is introduced due to the analytical approximations at the base of I_{mix} calculations. I_{updated} directly affects the wake expansion as specified in Eq. (A17) (Appendix-A1.3). I_{mix} is calculated with

$$I_{\text{mix}} = \frac{\sqrt{\frac{2}{3} K_{e,\text{tot}}}}{\bar{U}} - I, \quad (\text{A23})$$

605 with $K_{e,\text{tot}}$ representing the total turbulent kinetic energy including yaw effects. Using approximate solutions for velocity fluctuations in three dimensions, $K_{e,\text{tot}}$ can be expressed by

$$K_{e,\text{tot}} = 0.5 \left(u'^2 + (v' + v_{\text{curl}})^2 + (w' + w_{\text{curl}})^2 \right). \quad (\text{A24})$$

Cross-stream and vertical velocity fluctuations are defined as the sum of the average cross-stream v' and vertical w' velocity contributions at the turbine and the upstream turbine contributions v_{curl} and w_{curl} from Eq. (A21). Streamwise velocity fluctuations u' are derived from the turbulent kinetic energy K_e , based on the ambient turbulence intensity, and are defined as

$$u' = \sqrt{2K_e}, \text{ where } K_e = \frac{3}{2} (\bar{U}I)^2. \quad (\text{A25})$$

615 Considering yawed upstream turbines, secondary steering refers to the phenomenon of downstream turbines without yaw misalignment experiencing wake deflection and deformation due to the interaction between upstream-generated vortices and downstream turbine wake rotation (see Fig. 1-d). The GCH wake model captures secondary steering effects by computing an effective yaw angle γ_{eff} , the yaw misalignment that a downstream turbine would require to produce the same cross-stream velocity component V_{eff} as the average one calculated at its rotor due to upstream vortices \bar{V} . γ_{eff} is computed as

$$\gamma_{\text{eff}} = \text{argmin} |\bar{V} - V_{\text{eff}}|, \text{ where } V_{\text{eff}} = V_{\text{wake}}(\gamma). \quad (\text{A26})$$

620 Defining γ_{turb} as the rotor misalignment to the wind, the total yaw angle γ at the turbine rotor, consequently used in the wake deflection calculations (Appendix-A1.3), is expressed as

$$\gamma = \gamma_{\text{turb}} + \gamma_{\text{eff}}. \quad (\text{A27})$$



Although cross-wise and vertical velocity components are not explicitly modelled, the GCH wake model is capable of capturing, through analytical approximations, the increase in wake recovery due to yaw misalignment and secondary steering effects. The latter is particularly important when considering large wind farms and evaluating wake steering control strategies. However, GCH added complexity results in 3.5x higher computational requirements. Moreover, the parameter dependency for the Gaussian model still persists.

A2 Deflection models

When a turbine is yawed, the unbalance of thrust on the rotor leads to a cross-stream momentum gain and a deflection of the wake in the direction of yaw. Deflection models aim at quantifying this shift on the streamwise distance x and applying it to the streamwise velocity deficit field. All analytical models investigated in this study describe the velocity vector in one dimension only, leading to yaw solely affecting the streamwise velocity. Estimations of cross-stream and vertical velocity components are included in the GCH wake model (Appendix-A1.4), where an additional correction to the streamwise velocity field is applied based on these empirical estimations.

A2.1 Jimenez Model

The Jimenez deflection model is based on the empirical formulation proposed by Jiménez et al. (2009). The two main assumptions are: first, the streamwise velocity deficit values are almost negligible compared to inflow velocity $U_\infty - \delta u \simeq U_\infty$, second, the wake skew angle ζ is small enough that $\cos(\zeta) \simeq 1$ and $\sin(\zeta) \simeq \zeta$. Considering a turbine offset to the freestream by a yaw angle γ , the wake angle at the centreline is defined as

$$\zeta(x, \gamma, C_T) \approx \frac{(\zeta_{\text{init}}(C_T, \gamma))^2}{1 + 2k_d \frac{x}{D}}, \quad (\text{A28})$$

$$\zeta_{\text{init}}(\gamma, C_T) = 0.5 \cos^2 \gamma \sin \gamma C_T,$$

where k_d is a tuneable deflection parameter. The amount of wake deflection $\delta(x)$ can be calculated by skew angle integration. Ambient turbulence intensity can be indirectly taken into account through the k_d parameter. However, no considerations for added turbulence intensity and wake rotation are provided. Moreover, wake meandering is not accounted for due to the model's steady-state description.

A2.2 Bastankhah Model

The Bastankhah deflection model is based on the budget analysis of the continuity and Reynolds-averaged Navier-Stokes equations conducted by Bastankhah and Porté-Agel (2016). Using vortex theory, the wake deflection angle at the rotor is defined by

$$\theta \approx \frac{0.3\gamma}{\cos \gamma} \left(1 - \sqrt{1 - C_T \cos \gamma}\right). \quad (\text{A29})$$

By assuming the validity of Eq. (A29) is not limited to the calculation of flow angle at the rotor but it extends to the near wake region, and by approximating the near wake deflection to a constant value, the far wake onset wake deflection δ_0 can be



expressed as

$$\delta_0 = x_0 \tan \theta, \tag{A30}$$

where x_0 is defined in Appendix-A1.3. The total far wake deflection due to wake steering is defined as

$$r \frac{\delta(x)}{D} = \frac{\delta_0}{D} + \frac{\theta}{14.7} \sqrt{\frac{\cos \gamma}{k_w^2 C_T}} \left(2.9 + 1.3 \sqrt{1 - C_T} - C_T \right) \ln \left[\frac{(1.6 + \sqrt{C_T}) \left(1.6 \sqrt{\frac{8\sigma_y \sigma_z}{D^2 \cos \gamma} - \sqrt{C_T}} \right)}{(1.6 - \sqrt{C_T}) \left(1.6 \sqrt{\frac{8\sigma_y \sigma_z}{D^2 \cos \gamma} + \sqrt{C_T}} \right)} \right], \tag{A31}$$

655 with the main affecting factors being turbine operating conditions and ambient and added turbulence intensity.

Expressions for the symbols in the above equation are provided in Appendix-A1.3. Refer to Bastankhah and Porté-Agel (2016) for details on the model derivation. The Bastankhah deflection model does not include considerations about asymmetry due to wake rotation, calculations of vertical and cross-stream velocity components, or the effect of wake meandering due to the steady nature of the model predictions.

660 Appendix B

In this appendix, the validation of the presented optimisation framework is provided. The optimisation problem in Rak and Santos Pereira (2022) is replicated, and results for Jensen, Gaussian and GCH wake models are compared. A single column of aligned NREL 5 MW turbines is optimised for farm power maximisation. The yaw angle of each turbine in the farm layout is set as a variable bounded by $[-50^\circ, +50^\circ]$. SLSQP optimisation algorithm by Kraft (1988) is used. Moreover, optimiser
 665 parameters and initial yaw settings are matched to the values provided in Rak and Santos Pereira (2022). The farm layouts and atmospheric conditions optimised are shown in Table B1. Results are provided in Fig. B1. Overall, a satisfactory agreement is reached. The trends for the optimal variable solutions are captured for all wake models. Moreover, a close match for optimal yaw angles can be observed, except for a few isolated conditions. These marginal discrepancies are attributed to some of the wake model parameters not being specified in Rak and Santos Pereira (2022) and to the use of different FLORIS versions (2.4
 670 in the current work, 2.2 in Rak and Santos Pereira (2022)).

Table B1. Investigated wind farm layouts and atmospheric conditions from Rak and Santos Pereira (2022).

Case name	ws [$m s^{-1}$]	wd [deg]	Ti ₀ [%]	spacing [D]
Reference case (RC)	8.0	270.0	7.5	7.0
High wind speed (HWS)	13.0	270.0	7.5	7.0
Wind direction 275 ° (WD 275)	8.0	275.0	7.5	7.0
Low turbulence intensity (LTI)	8.0	270.0	5.0	7.0
Small spacing (SS)	8.0	270.0	7.5	5.0




Figure B1. Validation of the presented optimisation framework. Results are shown for the optimal turbine yaw angle comparison with Rak and Santos Pereira (2022) for (a) Jensen, (b) Gaussian, and (c) GCH wake models.

Author contributions. Filippo Gori: optimisation computations, Methodology, Validation, Data Analysis, Writing. Andrew Wynn: Methodology, Data Analysis, Funding acquisition, Supervision, Writing. Sylvain Laizet: Methodology, Data Analysis, Funding acquisition, Supervision, Writing.

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



675 **References**

- Abkar, M. and Porté-Agel, F.: Influence of atmospheric stability on wind-turbine wakes: A large-eddy simulation study, *Physics of Fluids*, 27, 035 104, 2015.
- 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, 680 2009.
- Barthelmie, R. J., Pryor, S. C., Frandsen, S. T., Hansen, K. S., Schepers, J. G., Rados, K., Schlez, W., Neubert, A., Jensen, L. E., and Neckelmann, S.: Quantifying the impact of wind turbine wakes on power output at offshore wind farms, *Journal of Atmospheric and Oceanic Technology*, 27, 1302–1317, 2010.
- Bastankhah, M. and Porté-Agel, F.: A new analytical model for wind-turbine wakes, *Renewable Energy*, 70, 116–123, 2014.
- 685 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., Porté-Agel, F., Port E-Agel, F., and Affiliations, .: Wind farm power optimization via yaw angle control: A wind tunnel study, *Journal of Renewable and Sustainable Energy*, 11, 023 301, 2019.
- Bastankhah, M., Shapiro, C. R., Shamsoddin, S., Gayme, D. F., and Meneveau, C.: A vortex sheet based analytical model of the curled wake 690 behind yawed wind turbines, *Journal of Fluid Mechanics*, 933, 36, 2022.
- Campagnolo, F., Petrović, V., Bottasso, C. L., and Croce, A.: Wind tunnel testing of wake control strategies, in: *Proceedings of the American Control Conference*, pp. 513–518, 2016.
- Crespo, A., Hernández, J., and Frandsen, S.: Survey of modelling methods for wind turbine wakes and wind farms, *Wind Energy*, 2, 1–24, 1999.
- 695 Deuffhard, P.: *Newton Methods for Nonlinear Problems*, Springer Berlin, Heidelberg, 2011.
- Dilip, D. and Porté-Agel, F.: Wind Turbine Wake Mitigation through Blade Pitch Offset, *Energies* 2017, Vol. 10, Page 757, 10, 757, 2017.
- Doekemeijer, B., Van Wingerden, J. W., and Fleming, P. A.: A tutorial on the synthesis and validation of a closed-loop wind farm controller using a steady-state surrogate model, in: *2019 American Control Conference (ACC)*, pp. 2825–2836, 2019.
- Doekemeijer, B., van der Hoek, D., and van Wingerden, J. W.: Closed-loop model-based wind farm control using FLORIS under time-varying 700 inflow conditions, *Renewable Energy*, 156, 719–730, 2020.
- Dou, B., Qu, T., Lei, L., and Zeng, P.: Optimization of wind turbine yaw angles in a wind farm using a three-dimensional yawed wake model, *Energy*, 209, 2020.
- Eriksson, D., Pearce, M., Gardner, J., Turner, R. D., and Poloczek, M.: Scalable Global Optimization via Local Bayesian Optimization, in: *Advances in Neural Information Processing Systems*, pp. 1–12, 2019.
- 705 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.
- Fleming, P., Annoni, J., Churchfield, M., Martinez-Tossas, L. A., Gruchalla, K., Lawson, M., and Moriarty, P.: A simulation study demon- 710 strating the importance of large-scale trailing vortices in wake steering, *Wind Energ. Sci*, 3, 243–255, 2018.



- Gebraad, P. M., Teeuwisse, F. W., Van Wingerden, J. W., Fleming, P., Ruben, S. D., Marden, J. R., and Pao, L. Y.: A data-driven model for wind plant power optimization by yaw control, in: Proceedings of the American Control Conference, pp. 3128–3134, 2014.
- Göçmen, T., Campagnolo, F., Duc, T., Eguinoa, I., Andersen, S., Petrović, V., Imširović, L., Braunbehrens, R., Liew, J., Baungaard, M., Laan, M., Qian, G., Aparicio-Sanchez, M., González-Lope, R., Dighe, V., Becker, M., Broek, M., Wingerden, J. W., Stock, A., and Meyers, J.:
715 FarmConnors wind farm flow control benchmark – Part 1: Blind test results, *Wind Energy Science*, 7, 1791–1825, 2022.
- Houck, D. R.: Review of wake management techniques for wind turbines, *Wind Energy*, 25, 195–220, 2021.
- Howland, M. F., Lele, S. K., and Dabiri, J. O.: Wind farm power optimization through wake steering, *Proceedings of the National Academy of Sciences of the United States of America*, 116, 14 495–14 500, 2019.
- Howland, M. F., Ghate, A. S., Lele, S. K., and Dabiri, J. O.: Optimal closed-loop wake steering – Part 1: Conventionally neutral atmospheric
720 boundary layer conditions, *Wind Energy Science*, 5, 1315–1338, 2020.
- Howland, M. F., Ghate, A. S., Quesada, J. B., Pena Martínez, J. J., Zhong, W., Larrañaga, F. P., Lele, S. K., and Dabiri, J. O.: Optimal closed-loop wake steering – Part 2: Diurnal cycle atmospheric boundary layer conditions, *Wind Energy Science*, 7, 345–365, 2022.
- Jensen, L., Morch, C., Sorensen, P., and Svendsen, K. H.: Wake measurements from the Horns Rev wind farm, in: EWEC 2004, London, United Kingdom, 2004.
- 725 Jensen, N.: A note on wind generator interaction, Tech. Rep. RISO-M-2411, Risoe National Laboratory, Roskilde (Denmark), 1983.
- Jiménez, Á., Crespo, A., and Migoya, E.: Application of a LES technique to characterize the wake deflection of a wind turbine in yaw, *Wind Energy*, 13, 559–572, 2009.
- Jonkman, J.: FAST, NREL, 2021.
- Jonkman, J., Butterfield, S., Musial, W., and Scott, G.: Definition of a 5-MW Reference Wind Turbine for Offshore System Development,
730 National Renewable Energy Laboratory (NREL), 2009.
- Katic, I., Højstrup, J., and Jensen, N.: A Simple Model for Cluster Efficiency, EWEC’86. Proceedings, 1, 407–410, 1987.
- Kheirabadi, A. C. and Nagamune, R.: A quantitative review of wind farm control with the objective of wind farm power maximization, *Journal of Wind Engineering and Industrial Aerodynamics*, 192, 45–73, 2019.
- King, J., Fleming, P., King, R. N., Martínez-Tossas, L. A., Bay, C., Mudafort, R., and Simley, E.: Control-oriented model for secondary
735 effects of wake steering, *Wind Energy Science*, 6, 701–714, 2021.
- 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., 1988.
- Martínez-Tossas, L. A., Annoni, J., Fleming, P., and Churchfield, M. J.: The aerodynamics of the curled wake: a simplified model in view of flow control, *Wind Energy Science*, 4, 127–138, 2019.
- 740 Niayifar, A. and Porté-Agel, F.: Analytical Modeling of Wind Farms: A New Approach for Power Prediction, *Energies*, 9, 741, 2016.
- NREL: FLORIS. Version 2.4, GitHub, 2021.
- Nygaard, N. G.: Wakes in very large wind farms and the effect of neighbouring wind farms, *Journal of Physics: Conference Series*, 524, 0–10, 2014.
- Pryor, S. C., Barthelmie, R. J., and Shepherd, T. J.: Wind power production from very large offshore wind farms, *Joule*, 5, 2663–2686, 2021.
- 745 Rak, B. P. and Santos Pereira, R. B.: Impact of the wake deficit model on wind farm yield: A study of yaw-based control optimization, *Journal of Wind Engineering and Industrial Aerodynamics*, 220, 104 827, 2022.