



Gradient-Based Wind Farm Layout Optimization Results Compared with Large-Eddy Simulations

Jared J. Thomas¹, Christopher J. Bay², Andrew P. J. Stanley², and Andrew Ning¹

¹Brigham Young University, Provo, UT 84602

²National Renewable Energy Laboratory, Golden, CO, 80401

Correspondence: Jared J. Thomas (jaredthomas@byu.net)

Abstract. The physics models commonly used during wind farm layout optimization include simplifying assumptions that can alter the design space compared to reality and higher-fidelity simulations. Some characteristics of these simple models may negatively influence the resulting layouts. In this paper, we perform wind farm layout optimization using a simple engineering wake model and then simulate the base and optimized layouts using large-eddy simulation (LES) to confirm that the layout was actually improved and not just an artifact of the simplifying assumptions in the low-fidelity wind farm simulation. We begin by describing the physics models used, including changes specific for use with gradient-based optimization. We then compare the simple model's output to previously published model and LES results. Using the simple models described, we performed gradient-based wind farm layout optimization using exact gradients. We optimized the wind farm twice, with high- and low-turbulence intensity (TI), respectively. We then recalculated annual energy production (AEP) using LES for the original and optimized layouts in each TI scenario and compared the results. For the high-TI case, the simple model predicted an AEP improvement of 7.7%, while the LES reported 9.3%. For the low-TI case, the simple model predicted a 10.0% AEP improvement, while the LES reported 10.7%. We concluded that the improvements found by optimizing with the simple model are not just an artifact of the model, but are real improvements assuming appropriate wind rose fidelity. We also found that the optimization did take advantage of the number of wind directions used, often aligning wind turbines in directions that were not included in the simulation. We found that, for the case studied, at least 50 wind directions are needed to avoid having the number of wind directions in the optimization significantly impact the optimized results. Future work should investigate further LES comparisons and wind rose fidelity in wind speed.

1 Introduction

Wind farm layouts are usually optimized using simple engineering models because the simple models require fewer computational resources and less time than the higher-fidelity models that are used for final analysis of a wind farm design. The models commonly used for wind farm layout optimization problems are based on simplifying physical assumptions that could potentially be exploited to find layouts that are good for the simple wake model but not good in reality. Some of the common assumptions include axisymmetric wakes and constant wake spread rates. In addition to the simplifying physical assumptions, the geometry of the design space is often simplified by dividing it into a grid space of possible locations (Mosetti et al., 1994;



- 25 Grady et al., 2005; Elkinton et al., 2008; González et al., 2010; Moorthy et al., 2014; Turner et al., 2014) or imposing a regular structure by only allowing array-like layouts and optimizing with array parameters, such as turbine row spacing and array rotation angle (Perez-Moreno et al., 2018). However, regular and grid layouts impose artificial limitations, potentially forcing the optimization algorithms away from the most optimal layouts (Chowdhury et al., 2012). Due to the limiting effect of grid and regular layout approaches, there is interest in unstructured, or irregular, wind farm layout optimization (Chowdhury et al., 2012; Réthoré et al., 2014; Fleming et al., 2015; Guirguis et al., 2016; Gebraad et al., 2017; Perez-Moreno et al., 2018). However, with the restrictions removed, unrestricted optimization could also be more susceptible to unrealistic characteristics of the models used during optimization. Because of the limitations inherent in these simplified models, it is uncertain whether the improvements reported by unstructured wind farm layout optimizations, using these simple models, are actual improvements or if the improvements are merely products of the models' weaknesses.
- 30 Because it is not practical to build and rebuild a wind farm to confirm the legitimacy of the improvements due to optimization with simple models, we are left with wind tunnel experiments and higher-order models, such as large-eddy simulation (LES). While many studies have compared simple wake models to LES and wind tunnel results (Crespo et al., 1999; Barthelmie et al., 2014, 2009; Gaumond et al., 2012; Beaucage et al., 2012; Moriarty et al., 2014; Annoni et al., 2014; Andersen et al., 2014; Göçmen et al., 2016; Machefaux et al., 2015), few have investigated the validity of the optimization results obtained with the models. One study (Bokharaie et al., 2016) involved LES in hybrid, or multifidelity, optimization with the Jensen "top-hat" wake model. LES was used for validation and for tuning during optimization using 30 turbines and up to 8 wind directions. This study showed similar improvements, according to both the LES and the Jensen model, but avoided some of the inherent inaccuracy of the simple model by adjusting it iteratively based on LES results. While this approach worked well, it was still much more costly computationally than optimization using only simple models. The only other validation activity we are aware of for irregular or unstructured wind farm layout optimization problems was performed using wind tunnel experiments and involved a small wind farm (9 turbines) and unidirectional flow (Chowdhury et al., 2012). The complexity of the wind farm layout optimization problem comes from large numbers of wind turbines and probabilistic wind in many directions, so it is important to investigate the results of unrestricted wind farm layout optimization problems involving larger wind farms and many wind directions.
- 45 Large numbers of wind turbines cause a corresponding increase in the number of design variables and constraints used by the optimization algorithm, which can reduce the effectiveness and efficiency of the algorithm. While the wind farm layout optimization problem is typically solved using gradient-free optimization methods (Herbert-Acero et al., 2014) (as done by Grady et al. (2005); Elkinton et al. (2008); González et al. (2010); Kusiak and Song (2010); Moorthy et al. (2014); Turner et al. (2014); Graf et al. (2016); Perez-Moreno et al. (2018)), gradient-free optimization methods exhibit reduced performance when applied to high-dimensional problems (Rios and Sahinidis, 2013). In the work of Rios and Sahinidis (2013), most gradient-free solvers were unable to solve more than 15% of the test problems studied with 10–30 design variables, and even the best methods tested could solve only 28% of the test problems with 31–300 design variables. New wind farms can include well over 100 turbines, meaning that unrestricted wind farm layout optimization problems can have more than 200 variables and



thousands of constraints. In short, the variables and constraints of the unrestricted wind farm layout optimization problems of interest exceed the abilities of most gradient-free optimization methods.

In contrast to gradient-free optimization methods, some gradient-based methods can solve problems with thousands of variables and constraints (Gill et al., 2005). However, gradient-based methods are prone to premature convergence when used with multimodal problems, such as the wind farm layout optimization problem. Because of the tendency to converge to local optima in multimodal problems, gradient-based optimization methods have not historically been seen as viable for solving the wind farm layout optimization problem (Herbert-Acero et al., 2014). However, recent studies have demonstrated that significant improvements to wind farm layouts, if not global optimality, can be obtained by using gradient-based methods either alone (Fleming et al., 2015; Guirguis et al., 2016; Gebraad et al., 2017; Thomas et al., 2017, 2021) or with gradient-free methods in a hybrid approach (Réthoré et al., 2014).

Solving a wind farm layout optimization problem requires several different types of models. Two of the most important of these are the wake model and the farm model. Generally speaking, wind turbine wake models define the wind speed in a turbine wake given a set of inflow conditions, while farm models combine the wakes of multiple turbines to determine the cumulative effect of their wakes. Many different wind turbine wake and wind farm models have been presented, all with varying levels of accuracy and computational cost. Many of the wake models, such as the “top-hat” Jensen/Park model (Jensen, 1983), the Frandsen model (Frandsen et al., 2006), and the original FLOW Redirection and Induction in Steady State (FLORIS) model (or Multi-Zone model) (Gebraad et al., 2014), have regions of nonphysical zero-valued gradients where their gradients go to zero while the gradients of the real design space do not. If a model with nonphysical zero-valued gradients is used with gradient-based optimization, the optimization may converge to an especially poor local optimum due to the lack of information provided in the gradients (Thomas et al., 2017).

The Bastankhah and Porté-Agel (BP) wake model presented in (Bastankhah and Porté-Agel, 2014, 2016) is well suited to gradient-based optimization with the exception of the near-wake region. Most of the BP model is smooth, continuous, and free from flat regions, but the near-wake definition of the BP model can be either flat or undefined. The near-wake region is rarely needed for evaluating a wind farm because wind turbines are usually placed far enough from each other that they are in the far wake. While the final wind farm design may not place turbines in the near-wake, it is important to have the near-wake region defined and not flat during gradient-based optimization because optimization algorithms may attempt infeasible solutions during the optimization process (Belegundu and Chandrupatla, 2011), and the near-wake region of the BP model may sometimes be within the defined optimization constraints. The BP wake model uses a Gaussian distribution to define the far wake and is easily applied to different turbines because it only has one independent parameter. Niayifar and Porté-Agel (2015, 2016) proposed a wind farm model intended to be used in concert with the BP wake model as published in 2014. We will refer to this wind farm model as the NP farm model. The NP farm model makes the one independent variable of the BP wake model dependent on turbulence intensity (TI). In this study, we use the NP farm model with the 2016 version of the BP wake model. While the NP farm model seems to provide good agreement with LES data, there are some minor problems with the NP farm model for use with gradient-based optimization that need to be addressed.



In this work, we address the lack of near-wake definition in the BP wake model for the purposes of using gradient-based optimization methods to solve the wind farm layout optimization problem, apply the NP wind farm model with some adjustments, and obtain exact gradients across the entire system. We then optimize a test case with a gradient-based approach using high (10.1%) and low (5.1%) TI values. Finally, we perform an LES study of the base case and optimized wind farm layouts with each TI level and compare the LES results with the model results.

This paper is an extension and revision of a previously published conference paper (Thomas et al., 2019). This version includes improved methods, completely new results, an additional TI scenario, revised and additional analyses, and updated conclusions. The low-fidelity wind farm simulation code was also completely rewritten and rigorously tested.

2 Methods

First, we provide details of the BP wake model and relevant adjustments made for compatibility with gradient-based optimization. Next, we detail our implementation of, and adjustments to, the NP farm model, including wake combination, TI, sampling over the rotor-swept area, and power calculations. Following the model descriptions, we verify the models using previously published LES results. We then present our model tuning approach, our optimization test case wind farm, followed by our optimization approach. We conclude this section with key LES details.

2.1 Wake model

While the BP wake model exhibits many characteristics necessary for use with gradient-based optimization, there are some small problems that must be overcome. First, the main BP model is undefined from the turbine location up to several diameters downstream of the wind turbine. Bastankhah and Porté-Agel (2016) suggested using regions of constant velocity in the near-wake, but using regions of constant velocity in the wind turbine wake has been shown to cause premature convergence to poor local optima when optimizing with gradient-based methods (Thomas et al., 2017). The standard BP model cannot be used in the near-wake region because there is a discontinuity due to a negative value in the square-root term of the velocity deficit calculation of the BP wake model, as shown in Eq. (1) (Bastankhah and Porté-Agel, 2016).

$$\frac{\Delta u_{ij}}{\bar{u}_i} = \left[1 - \sqrt{1 - \frac{C_T \cos \gamma}{8\sigma_y \sigma_z / d^2}} \right] \exp \left(-0.5 \left[\frac{y - \delta}{\sigma_y} \right]^2 \right) \exp \left(-0.5 \left[\frac{z - z_h}{\sigma_z} \right]^2 \right) \quad (1)$$

In Eq. (1), Δu_{ij} is the velocity deficit in the wake due to turbine i at the point j ; \bar{u}_i is the average inflow velocity of turbine i ; C_T is the thrust coefficient; γ is the yaw relative to the inflow wind direction; d is the rotor diameter; y is the horizontal distance from the line through the rotor hub to the point of interest perpendicular to the wind direction; δ is the horizontal wake offset; z and z_h are the heights of the point of interest and the rotor hub, respectively; and σ_y and σ_z are the horizontal and vertical wake spread as defined in Eqs. (2) and (3) (Bastankhah and Porté-Agel, 2016).

$$\sigma_y = k_y [x - x_0] + \frac{d \cos \gamma}{\sqrt{8}} \quad (2)$$



$$\sigma_z = k_z[x - x_0] + \frac{d}{\sqrt{8}} \quad (3)$$

In Eqs. (2) and (3), k_y and k_z are the horizontal and vertical wake growth rates (assumed to be equal), x is the downstream distance from the turbine, and x_0 is the length of the potential core as defined in Eq. (4) (Bastankhah and Porté-Agel, 2016).

$$\frac{x_0}{d} = \frac{\cos \gamma [1 + \sqrt{1 - C_T}]}{\sqrt{2} [\alpha^* I + \beta^* [1 - \sqrt{1 - C_T}]]} \quad (4)$$

In Eq. (4), I is the TI, and α^* and β^* are constants used to estimate the length of the near-wake potential core. We set $\alpha^* = 2.32$ and $\beta^* = 0.154$, the tuned values used by Bastankhah and Porté-Agel (2016).

Wind shear was added to the wake model using a power law defined as shown in Eq. (5).

$$u = u_r \left[\frac{z - z_o}{z_r - z_o} \right]^\psi \quad (5)$$

In Eq. (5), u_r is the reference wind speed, z is the height of interest, z_r is the height at which u_r was measured, z_o is the height of local ground, and ψ is the shear exponent.

The discontinuity mentioned previously can occur for various combinations of k_y , k_z , and $(x - x_0)$. When local TI is calculated in the NP wind farm model, as done by Niayifar and Porté-Agel (2016) (discussed further in Sect. 2.2), the values of k_y and k_z vary, making the discontinuity even more mobile than in the original model. We can determine where the model is undefined as follows: First, we know the model is undefined when the condition in Eq. (6) is met.

$$\frac{C_T \cos \gamma}{8 \sigma_y \sigma_z / d^2} > 1 \quad (6)$$

By substituting Eqs. (2) and (3) into Eq. (6), and separating the terms based on the powers of $[x - x_0]$, we obtain Eq. (7).

$$k_y k_z [x - x_0]^2 + \frac{d[k_y + k_z \cos \gamma]}{\sqrt{8}} [x - x_0] - \frac{d^2 \cos \gamma [C_T - 1]}{8} < 0 \quad (7)$$

To find the exact downstream location where the model begins to be defined, x_d , we apply the quadratic formula to Eq. (7), select the relevant root, then solve for x to obtain Eq. (8).

$$x_d = x_0 + \frac{d \left[k_y + k_z \cos \gamma - \sqrt{[k_y + k_z \cos \gamma]^2 - 4 k_y k_z [C_T - 1] \cos \gamma} \right]}{2 k_y k_z \sqrt{8}} \quad (8)$$

The model, then, is undefined for regions where $x < x_d$. To remove the discontinuity, we used a simple linear interpolation on both the velocity value and the wake spread, leaving velocity and wake spread constant at their corresponding values at x_d from x_d to the hub location.

The results of the model changes discussed above are shown in Fig. 1. These changes make the model continuous throughout the domain and differentiable everywhere except at the exact turbine location. With a separation constraint in place, the



probability of turbines landing exactly on top of each other is extremely small, as noted by Thomas et al. (2017). With the linear interpolation method, optimizations succeed even when turbines are placed less than 1 diameter from each other. While the linear interpolation makes little attempt to be physically accurate in the near-wake, the accuracy of the model for AEP calculation purposes should be unaffected because turbines will almost never be placed close enough together to be in the linear interpolation region in a final design. If greater accuracy is desired in the near-wake, linear interpolation could be used during optimization, and a more accurate near-wake model, such as the one proposed by Keane et al. (2016), could be used for final calculations.




Figure 1. Vertical contour plot using the Bastankhah and Porté-Agel wake model (a) without and (b) with the added linear interpolation for the near-wake. The turbine and flow conditions correspond to the high-turbulence intensity case study presented in this paper.

2.2 Wind farm model

Niayifar and Porté-Agel (2016) proposed a wind farm model that used the BP model as published by Bastankhah and Porté-Agel (2014) for calculating the wind speed in the wakes. We applied the same wind farm model to the 2016 version of the BP model as published by Bastankhah and Porté-Agel (2016), with some variation as discussed in the following sections.



2.2.1 Wake combination

160 Niayifar and Porté-Agel (2016) found the most accurate method of wake combination to be a linear velocity deficit superposition that uses the inflow velocity of the upstream turbine as shown in Eq. (9).

$$u_j = \bar{u}_\infty - \sum_{i=1}^{N_T} [\bar{u}_i - u_{ij}] = \bar{u}_\infty - \sum_{i=1}^{N_T} \Delta u_{ij} \quad (9)$$

In Eq. (9), u_j is the flow velocity at the point of interest, \bar{u}_i is the average inflow velocity of upstream turbine i , u_{ij} is the wind speed that would exist at the point j if only turbine i were accounted for, and \bar{u}_∞ is the average free-stream velocity.

165 For the upstream turbine velocity to be used, the inflow velocity of the turbines must be calculated, in order, from upstream to downstream. We used the heapsort algorithm to obtain a sorted index list of the turbines prior to solving the system for each wind direction. Using the sorted index list allowed turbine sorting without overcomplicating the gradients of the system. Our tests of the accuracy of various wake combination methods also found this method to be the most accurate in comparison to the LES results used by Niayifar and Porté-Agel (2016).

170 2.2.2 Turbulence intensity

Niayifar and Porté-Agel calculated local TI using the model presented by Crespo and Hernández (1996). The local TI is then applied to the BP model through an empirical model for calculating the wake expansion coefficient as shown in Eq. (10) (Niayifar and Porté-Agel, 2016).

$$k^* = 0.3837I + 0.003678 \quad (10)$$

175 In Eq. (10), I is the TI. According to Niayifar and Porté-Agel (2016), this model is only valid for TI values between 0.065 and 0.15. We chose to use this model for both our high-TI and low-TI cases even though our low-TI case falls below the given range. It is possible that using the TI model outside the specified range has introduced some error into our model. While the initial BP wake model does not include TI (Bastankhah and Porté-Agel, 2014), TI was included by Bastankhah and Porté-Agel in their 2016 version of the model (Bastankhah and Porté-Agel, 2016). We used the same TI values for the BP 2016 wake
 180 model and the NP wind farm model.

To avoid increasing the number of local optima, we used two different approaches to calculating TI at various points in the optimization process: (1) ambient TI only and (2) local TI with a smooth maximum function. Further motivation for the two TI calculation methods is discussed in Sect. 2.7. The smooth maximum function we used is shown in Eqs. (11), (12), and (13), and is derived from the Log-Sum-Exponential function (Cook, 2010).

$$185 \quad t_{max} = \min(t_1, t_2) \quad (11)$$

$$t_{min} = \max(t_1, t_2) \quad (12)$$



$$\max_{smooth}(t_{max}, t_{min}, s) = \frac{\ln(1 + \exp(s[t_{min} - t_{max}])) + st_{max}}{s} \quad (13)$$

190 In Eqs. (11), (12), and (13), s controls the smoothness of the transition between terms and influences the error of the calculation; t_{max} and t_{min} represent the respective maximum and minimum of the two values being compared. The maximum error of Eq. (13) occurs when $t_{min} = t_{max}$ and can be calculated as in Eq. (14).

$$\text{Error}_{max} = \frac{\ln(2)}{s} \quad (14)$$

To keep the error from the smooth maximum below 0.001, or about 1% of the total TI, we used $s = 700$.

195 In calculating local TI, Eq. (13) is applied, as in Eq. (15), to compare the product of the added local TI and the wake overlap ratio. Niayifar and Porté-Agel (2016) used a hard maximum in Eq. 15.

$$I_{+j} = \max_{smooth}\left(\frac{I_{+ij}A_{wij}}{A_{rj}}\right) \quad (15)$$

In Eq. (15), A_{wij} is the area of the wake of turbine i at the downstream location of turbine j , A_{rj} is the area of the rotor of turbine j , and I_{+ij} is the added TI from turbine i at the downstream location of turbine j . In our implementation, only two
 200 values are compared at a time: the current value and the previous result of the smooth maximum function.

We calculated added TI as proposed by Crespo and Hernández (1996) and reproduced in Eq. (16).

$$I_{+ij} = 0.73a_i^{0.8325}I_i^{0.0325}\left[\frac{x}{d_i}\right]^{-0.32} \quad (16)$$

In Eq. (16), a_i is the axial induction of turbine i and I_i is the TI at turbine i . To account for varying axial induction, we based the axial induction on the C_T curve of the wind turbine as done by Gebraad et al. (2017) and reproduced in Eq. (17).

$$205 \quad a = \begin{cases} 0.143 + \sqrt{0.0203 - 0.6427[0.889 - C_T]} & C_T > 0.96 \\ 0.5[1 - \sqrt{1 - C_T}] & \text{otherwise} \end{cases} \quad (17)$$

We finally calculated total TI using Eq. (18) (Niayifar and Porté-Agel, 2016).

$$I_j = \sqrt{I_0^2 + I_{+j}^2} \quad (18)$$

2.2.3 Sampling over the rotor-swept area

Various methods can be used to approximate the effective inflow wind speed at the rotor. The ideal method is a full averaged
 210 integration across the rotor-swept area. We approximated the integral by sampling the velocity at various points and taking the average. When comparing our code to LES and other data, we used 100 sampling points across the rotor-swept area arranged in the sunflower seed pattern as shown in Fig. 2(a). We found 100 sample points on the rotor-swept area more than sufficient to achieve the level of accuracy demonstrated by Niayifar and Porté-Agel (2016). For efficiency during optimization, we approximated the effective inflow wind speed of each wind turbine using a single sample point located at the rotor hub as
 215 shown in Fig. 2(b). We performed all analyses and provided all results using the 100 sample points across the rotor.




Figure 2. Location of wind speed sampling points, shown as orange dots, for approximating the inflow wind speed at each turbine. (a) Pattern used for LES comparisons and all provided results. (b) Single point used for approximating effective inflow wind speed during optimization.

2.2.4 AEP and wind turbine power

We calculated AEP as in Eq. (19).

$$AEP = \left[\frac{\text{hours}}{\text{day}} \right] \left[\frac{\text{days}}{\text{year}} \right] \sum_{k=1}^{N_D} \left[f_k \sum_{j=1}^{N_T} P_{jk} \right] \quad (19)$$

In Eq. (19), N_D is the number of wind directions, N_T is the number of wind turbines, and f_k is the probability of wind in direction k . The power of turbine j in direction k , P_{jk} , is calculated based on the definition of the power coefficient as done by Gebraad et al. (2014) and shown in Eq. (20).

$$P_{jk} = \begin{cases} 0 & \bar{u}_j < u_{cutin} \\ 0.5\rho A_j C_{Pj} \bar{u}_j^3 & \bar{u}_j \geq u_{cutin} \text{ and } P_j < P_{rated} \\ P_{rated} & P_j \geq P_{rated} \end{cases} \quad (20)$$

In Eq. (20), ρ is the air density, A_j represents the rotor-swept area of turbine j , and C_{Pj} is the power coefficient of turbine j for the given effective wind speed at the rotor (\bar{u}_j).

2.3 Model verification

To verify our implementation of the models described, we compared the output of our code with data for the Horns Rev wind farm from Niayifar and Porté-Agel (2016). We show the normalized power comparison along rows of turbines in Fig. 3. We



show the normalized power comparison of the whole wind farm by wind direction in Fig. 4. The Horns Rev comparisons using 100 samples on the rotor-swept area are shown in Figs. 3(a) and 4(a). The Horns Rev comparisons using one sample on the rotor-swept area are shown in Figs. 3(b) and 4(b). While using only one sample on the rotor-swept area resulted in large errors, especially for the second row of turbines and wind directions where wake effects were highest, using 100 sample points resulted in excellent agreement with Niayifar and Porté-Agel (2016). Based on our model error when we used both 100 samples and local TI being within the range reported by Niayifar and Porté-Agel (2016), we determined that our implementation of the BP wake and NP farm models was sufficient and could be used to perform wind farm layout optimization case studies. For discussion of when and why each of the TI types were used, see Sect. 2.7.




Figure 3. Averaged power of turbines in columns two, three, and four of the Horns Rev wind farm by turbine row. (a) Comparing our code using 100 points to approximate the effective wind speed at the rotor. (b) Comparing our code using a single sample at the rotor hub to approximate the effective wind speed at the rotor. Niayifar 2016 refers to LES and model results from Niayifar and Porté-Agel (2016). See Sec. 2.7 for discussion on when and why we used each TI type.

2.4 Large-eddy simulation

The Simulator fOr Wind Farm Applications (SOWFA) is a high-fidelity, LES tool that was developed at the National Renewable Energy Laboratory (NREL) for wind plant studies (Churchfield and Lee, 2015; Churchfield et al., 2012a; Fleming et al., 2013). It is a computational fluid dynamics solver based on OpenFOAM (Jasak et al., 2007) and can model turbines as actuator disks or actuator lines. This study uses turbines modeled as actuator disks to reduce computational cost. Separate studies have been conducted that demonstrate that the steady-state power is similar in actuator disk and actuator line cases (Martinez et al., 2012).

SOWFA solves the three-dimensional, incompressible, Navier-Stokes equations and transport of potential temperature equations, which take into account the thermal buoyancy and Earth rotation (Coriolis) effects in the atmosphere. The inflow conditions for these simulations are generated using a periodic atmospheric boundary layer precursor with no turbines.

SOWFA calculates the unsteady flow field to compute the time-varying power, velocity deficits, and aerodynamic loads at each turbine in a wind plant. This level of computation, with high-fidelity accuracy, takes on the order of hours to days to run on

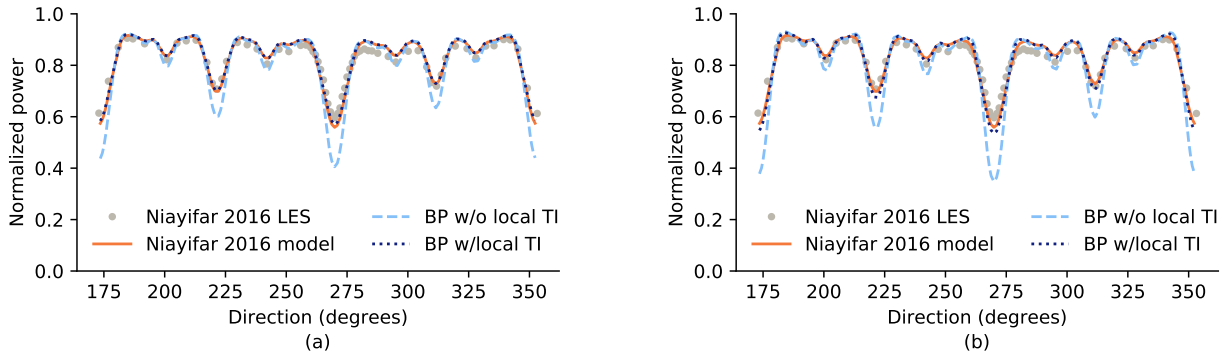


Figure 4. Power by direction for the Horns Rev wind farm. (a) Comparing our code using 100 points to approximate the effective wind speed at the rotor. (b) Comparing our code using a single sample at the rotor hub to approximate the effective wind speed at the rotor. Niayifar 2016 refers to LES and model results from Niayifar and Porté-Agel (2016). See Sect. 2.7 for discussion on when and why we used each TI type.

a supercomputer using a few hundred to a few thousand processors, depending on the size of the wind plant. The simulations run for this study were performed on Eagle, NREL’s high-performance computer.

It should be noted that studies have been performed to validate SOWFA. For example, it has been compared with 48-turbine Lillgrund wind plant field data and shows good agreement through the first five turbines in a row aligned with the wind direction (Churchfield et al., 2012b). In addition, SOWFA has been tested to verify that it captures the inertial range in the turbulent energy spectra and log layer in the mean flow, both of which characterize a real atmospheric boundary layer (Churchfield et al., 2012a). Further validation studies are ongoing.

Actuator disk simulations of the 38-turbine layout described in Sect. 2.6, as well as the optimized layout, were performed using SOWFA. We limited the LES simulation to a 5 km-square area to keep necessary computational resources to a reasonable level. In order to avoid edge effects in the simulation, we allowed for a 0.5 km spacing between the edge of the wind farm and the edge of the simulation space. To avoid needing a different precursor for each of the 12 directions, the layout was rotated to a reference wind direction of 270° . The scenarios were simulated under neutral atmospheric conditions with an 8 m s^{-1} mean wind speed at 80 m and approximately 6.5% TI at hub height.

The simulations each ran for 7200 s of simulated time. A structured mesh was used in this study with a grid spacing of 10 m in the x , y , and z directions, resulting in a grid that is $500 \times 500 \times 100$ (i.e., 25,000,000 grid cells). Using 500 cores, simulations took on the order of 1.5 days to complete.

2.5 Tuning to LES

The reference wind speed (u_r), shear exponent (ψ), and the ambient TI (I) were determined based on analysis of the LES precursor. We found that tuning model inputs based on matching the model outputs to SOWFA resulted in a close match to SOWFA for the base case, but a poor match to SOWFA for other layouts (e.g., optimized) using the same tuned inputs and



wind conditions. For this reason, we chose to base our model inputs on values determined from the SOWFA precursor rather than run the risk of overfitting our model.

We determined the shear exponent (ψ) by fitting Eq. (5) to a series of wind speeds in the LES precursor. We set z_r to 90 m, z_0 to 0 m. We fit the model of Eq. (5) to the LES data by varying u_r and ψ to vary while minimizing the least squared errors. The curve fit results are shown in Fig. 5. We found the best fit values of the high-TI case for u_r and ψ to be 7.95 m s^{-1} and 0.175, respectively, as shown in Fig. 5(a). We found the best fit values of the low-TI case for u_r and ψ to be 7.99 m s^{-1} and 0.084, respectively, as shown in Fig. 5(b).




Figure 5. The calculated wind speed using the power law from Eq. (5) compared with the values obtained from the LES precursors. The reference height is 90 m, with ground height at 0 m. (a) High-TI case with a shear exponent (ψ) of 0.175 and a hub height wind speed of 7.95 m s^{-1} . (b) Low-TI case with a shear exponent (ψ) of 0.084 and a hub height wind speed of 7.99 m s^{-1} .

The ambient TI was determined as shown in Eq. (21).

$$I = \frac{\sqrt{\sum u_x^2}}{\bar{u}_x}. \quad (21)$$

In Eq. (21), u_x is the set of all the wind speeds at hub height in the primary wind direction, and \bar{u}_x is the average wind speed at hub height in the wind direction. Average TI was found to be 5.1% for the low-TI precursor and 10.1% for the high-TI precursor. We set the inflow wind speed to 8.0 m s^{-1} , based on the specified inflow velocity of the LES precursor.

2.6 Test case

For optimization, we used the NREL 5 MW reference turbine (Jonkman et al., 2009). We used the same thrust and power coefficient curves for the NREL 5 MW turbine as used in (Gebraad et al., 2017), with linear interpolation between points.

Our test wind farm was constrained by the domain of the LES simulation precursor, which was a 5 km-by-5 km square. Allowing for sufficient space between the turbines and the simulation boundaries, the available space for the wind farm was a circle with a 2 km radius in the middle of the simulation space. Within the resulting circular region, we placed as many turbines as possible in a concentric circular pattern with a minimum allowable distance between turbines of five times the rotor



diameter. The resulting baseline layout is shown in Fig. 6, where turbine numbers are provided for reference to results in Figs. 14 and 15.

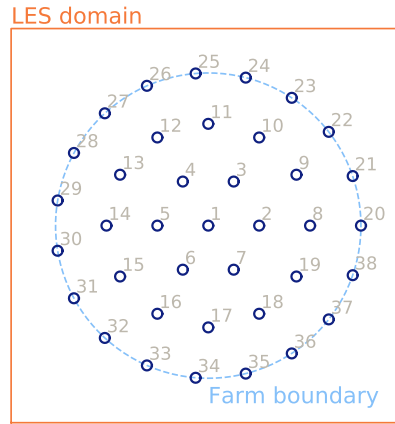


Figure 6. Base case wind farm layout. The circles marking turbine locations are to scale, with diameters equal to the rotor diameter. Turbine numbers are provided for comparison to results figures (Figs. 14 and 15).

Because we used a multistart approach, we needed to generate multiple starting layouts. The other starting layouts were generated by applying a random rotation angle between 0 and 2π radians to each of the concentric circles of the base layout, individually. We generated 399 layouts besides the base layout shown in Fig. 6 for a total of 400 starting layouts.

For the wind frequencies, we chose to use the Nantucket wind rose (Western Regional Climate Center, 2012). To reduce the computational cost of the LES simulations, we generated a wind rose with 12 directions by binning every 3 directions from the Nantucket wind rose starting with wind from the north. We used the 12-direction wind rose for both the optimizations and the LES. The original and final wind roses are shown in Fig. 7.

2.7 Optimization

We optimized the wind farm layout using AEP as the objective. Our optimization problem formula is shown in Eq. (22).

$$\begin{aligned}
 &\underset{x_i, y_i}{\text{maximize}} && AEP(x_i, y_i) \quad i = 1 \dots 38 \\
 &\text{subject to} && S_{ij} \geq 2d \quad i, j = 1 \dots 38 \quad i \neq j \\
 &&& [x_c - x_i]^2 + [y_c - y_i]^2 \leq r_b^2 \quad i = 1 \dots 38
 \end{aligned} \tag{22}$$

In Eq. (22), (x_i, y_i) is the position of each turbine i , S_{ij} represents the separation distance between each pair of turbines i and j , (x_c, y_c) is the location of the center of the wind farm, and r_b is the radius of the wind farm boundary.




Figure 7. (a) Nantucket frequency wind rose (Western Regional Climate Center, 2012). (b) Nantucket frequency wind rose binned into 12 directions.

We set up and ran the wind farm simulations in Julia (Bezanson et al., 2017) using FLOWFarm.jl¹. Exact gradients of the full simulation model were obtained using the ForwardDiff.jl package (Revels et al., 2016). The gradients of both the objective function and constraints were scaled to be between ± 1 . We then optimized the final problem using the Sparse Nonlinear OPTimizer (SNOPT), a gradient-based optimization algorithm that uses a sequential quadratic programming approach. We used SNOPT in this case because it is well suited to nonlinear problems with high dimensionality (Gill et al., 2005). We used SNOPT to connect FLOWFarm.jl, ForwardDiff.jl, and SNOPT.

Along with SNOPT, we used the wake expansion continuation (WEC) method for reducing the multimodality of the wind farm layout optimization problem (Thomas et al., 2021). We used the same WEC relaxation factors as Thomas et al. (2021), [3.0 2.6 2.2 1.8 1.4 1.0 1.0]. We also used two methods for calculating TI: (1) constant ambient TI only, (2) local TI calculated with a smooth maximum function. Calculating local TI, even with the smooth maximum function, adds additional local optima that are exaggerated when using WEC. The additional local optima drastically reduce the effectiveness of the WEC method, so we neglected local TI during the WEC optimization series (TI method (1)). After completing the WEC series, we ran a final optimization, including local TI, from the optimized layout found using the WEC series.

3 Results and discussion

The distributions of AEP for all 400 optimizations for both the high and low-TI cases are provided in Fig. 8. The distributions for the high-TI case are much less spread than those for the low-TI case. This makes sense because the more rapid wake

¹<https://github.com/byuflowlab/FLOWFarm.jl>

²<https://github.com/byuflowlab/SNOW.jl>



recovery in the high-TI case makes it easier to avoid wake interactions than in the low-TI case. The final optimized layouts were chosen as the layouts from the multistart set with the highest optimized AEP. These final optimized layouts are shown in Fig. 9. While the high- and low-TI optimized layouts are different, they share some important characteristics.




Figure 8. AEP distributions (determined using the BP model with 100 rotor sample points) for both the high-TI (a) and low-TI (b) cases for all 400 starting and optimized layouts. Start refers to the actual starting layouts used for the optimizations. Bin width is 1 GW h. Legend applies to both (a) and (b).




Figure 9. Optimized wind farm layouts for both the high-TI (a) and low-TI (b) cases. The circles marking turbine locations are to scale, with diameters equal to the rotor diameter. Numbers are provided for reference to other figures.



The first shared characteristic is that in each optimized layout the wind turbines form a pattern of nested hexagons (see Fig. 10) instead of the nested circle pattern that they started with. The hexagonal pattern is clearer in the high-TI case and more approximate in the low-TI case. The wind directions are at [10 40 70 100 130 160 190 220 250 280 310], while in the optimized results the faces of the two outer hexagons combined are at [25 55 85 115 145 175 205 235 265 295 325 355]. The sides of the two outer hexagons face almost exactly between the wind directions, with an offset of about 15° between the hexagon faces and the wind directions. Interestingly, the sides of the outer hexagon of the high-TI case are nearly aligned with the sides of the inner hexagon of the low-TI case and vice versa. While not all sides of the low-TI case layout match the hexagons, the turbines that break the pattern are the turbines on sides where direct wind is less likely and turbines on the inside of the farm where direct wind is unlikely to ever occur.




Figure 10. High-TI (a) and low-TI (b) optimized layouts with overlaid rotated hexagons. Angles are rounded to the nearest degree.

The second shared characteristic is that in both optimized layouts the turbine spacing near the wind farm boundary is reduced along the faces of the hexagons. The close spacing of these turbines is along a direction not included in the wind rose. So, it appears that the optimization did take advantage of large gaps in the wind rose. Third, many of the turbines on the edges of the farm were not placed directly on the boundary. Rather, the optimization favored straight lines along directions not included in the wind rose over increased turbine separation. While TI did have an impact on the optimized layouts, the general characteristics of the layouts optimized using high and low TI are very similar. The wind rose fidelity appears to have a significant impact on the wind farm layout, with the optimization taking advantage of unused directions by aligning turbine rows with, and reducing turbine spacing along, directions not included in the simulations. A more refined wind rose would reduce the AEP gains of these optimized layouts.

To investigate the impacts of the number of wind directions on the optimization, we ran a series of high-TI optimizations using the same parameters already discussed, but with a varying number of wind directions. To obtain the wind rose for each number of wind directions we interpolated the probabilities of the 36-direction wind rose shown in Fig. 7(a) using an Akima spline from the Julia package FLOWMath.jl³. The directions used were [5 10 15 20 30 40 50 70 90 110 140 170 200 240 280 320 360]. We ran 100 optimizations for each number of wind directions, using the first 100 layouts of the set of starting

³<https://github.com/byuflowlab/FLOWMath.jl.git>



layouts already discussed. The average change in AEP from starting layouts to optimized layouts, as calculated using 100 rotor sample points, are shown in Fig. 11. The same figure also shows the average AEP change, for the same starting and optimized layouts, re-evaluated using 360 wind directions. We found that for about 50 or more wind directions, the number of wind directions in the optimization no longer impacted the final AEP as calculated with 360 wind directions. We also found that for the layouts optimized with less than 30 wind directions, some or all of the optimized AEPs calculated with 360 wind directions were actually lower than the starting AEPs calculated with 360 wind directions. Therefore, optimizing with less than 30 directions could provide some layouts that are actually worse than the start. Optimizing with less than 50 directions may provide exaggerated improvement results and fail to provide as much AEP improvement as is available according to the wind rose with 360 directions. Optimizations should be performed with at least 50 wind directions to achieve as much improvement as possible. It is likely that these results would change significantly if a wind speed distribution was also considered. Despite the insufficient directional fidelity of the wind rose in our primary study case, comparing optimization improvement between the simple BP model and SOWFA in this case may still provide some insight into the validity of optimization results in terms of the BP model impacts.




Figure 11. This figure shows that at least 40 or 50 wind directions are needed during optimization based on the average change in AEP from starting layout to optimized layout. Results are shown with AEP calculated using the number of wind directions in the optimizations (N_D) and also with AEP re-calculated (not re-optimized) using 360 wind directions for the same starting and optimized layouts. The error bars indicate the maximum and minimum AEP change. We ran 100 optimizations from different starting layouts for each number of wind directions.

The directional power for our primary case study with both high and low-TI for both SOWFA and the BP model are shown in Fig. 12 (a) and Fig. 12 (b). While the directional power in both the high and low-TI cases show similar trends between the BP model and SOWFA, there is a significant offset. The BP model appears to consistently underpredict the power as compared



to SOWFA. The general shapes of the curves are also similar for both models between high and low-TI, but the low-TI seems to accentuate the power reductions in wind directions with more wake interactions, resulting in sharper peaks and valleys.

360 The error in the directional power production is shown in Fig. 12 (c) and Fig. 12 (d). In these figures, we can see that the error in the base case varied across directions much more than the error in the optimized case. The low-TI case shows less error across wind directions for both the base and optimized layouts, but otherwise similar directional error trends are apparent as in the high-TI case.

Improvement in the directional power production is shown in Fig. 12 (e) and Fig. 12 (f). The directional power improvement
365 shows that the power was increased the most for the directions that originally had the lowest power. SOWFA and the BP model each show similar trends. The two directions with the most improvement, according to SOWFA, are opposing directions (10° and 190° for the high-TI case and 40° and 220° for the low-TI case), so it makes sense that the pairs show similar improvement. It is interesting to note that the greatest improvement in directional power in the high-TI case, according to SOWFA, was not at 220° (the primary wind direction). However, the directional power results can be misleading because we are using a weighted
370 wind rose.

To gain greater insight into what is going on across the wind directions, we first switch from AEP to wake loss. We calculated wake loss according to Eq. (23).

$$L_w = 100 \left(1 - \frac{P}{P_I} \right), \quad (23)$$

where the ideal power, P_I , was determined as the power of the wind farm in a given direction if there were no wake interactions
375 given the same inflow conditions. We determined the free stream power in each direction as the average power being produced by all the turbines in free stream wind for that direction. We calculated P_I for the BP model and SOWFA separately.

The directional wake loss is shown in Fig. 13(a) and Fig. 13(b). As in the power results, the trends are very similar for the BP model and SOWFA. However, the BP model and SOWFA wake loss curves are more similar than the corresponding power
380 curves, and we can see that even though the power may be underpredicted by the BP model, the model is certainly capturing important effects.

Using the wake loss and the directional ideal power, we calculated the directional annual energy loss, shown in Fig. 13(c) and Fig. 13(d). Even though the greatest power increase was at 190° for the high-TI case, we can now see that 220° (the primary wind direction) is by far the most important direction for the overall optimization in both cases. Most of the other directions had much less wake loss in the base case. While the spike in directional energy loss at 220° was mostly flattened through the
385 optimization, the low-TI case exhibits the spike in the optimized curves as well. This is likely because of the slower wake recovery in low-TI wind conditions.

From the directional annual energy loss, we determined the directional annual energy improvement to indicate how much the total energy produced increased in each direction due to the optimization. The results are shown in Fig. 13(e) and Fig. 13(f). The close match of the BP and SOWFA results in directional annual energy loss is striking. Even though the actual power
390 predictions were significantly underpredicted by the BP model, the improvements in each direction were very well predicted with few exceptions. We can see that the power improvement at 220° was the highest because it is the primary wind direction.

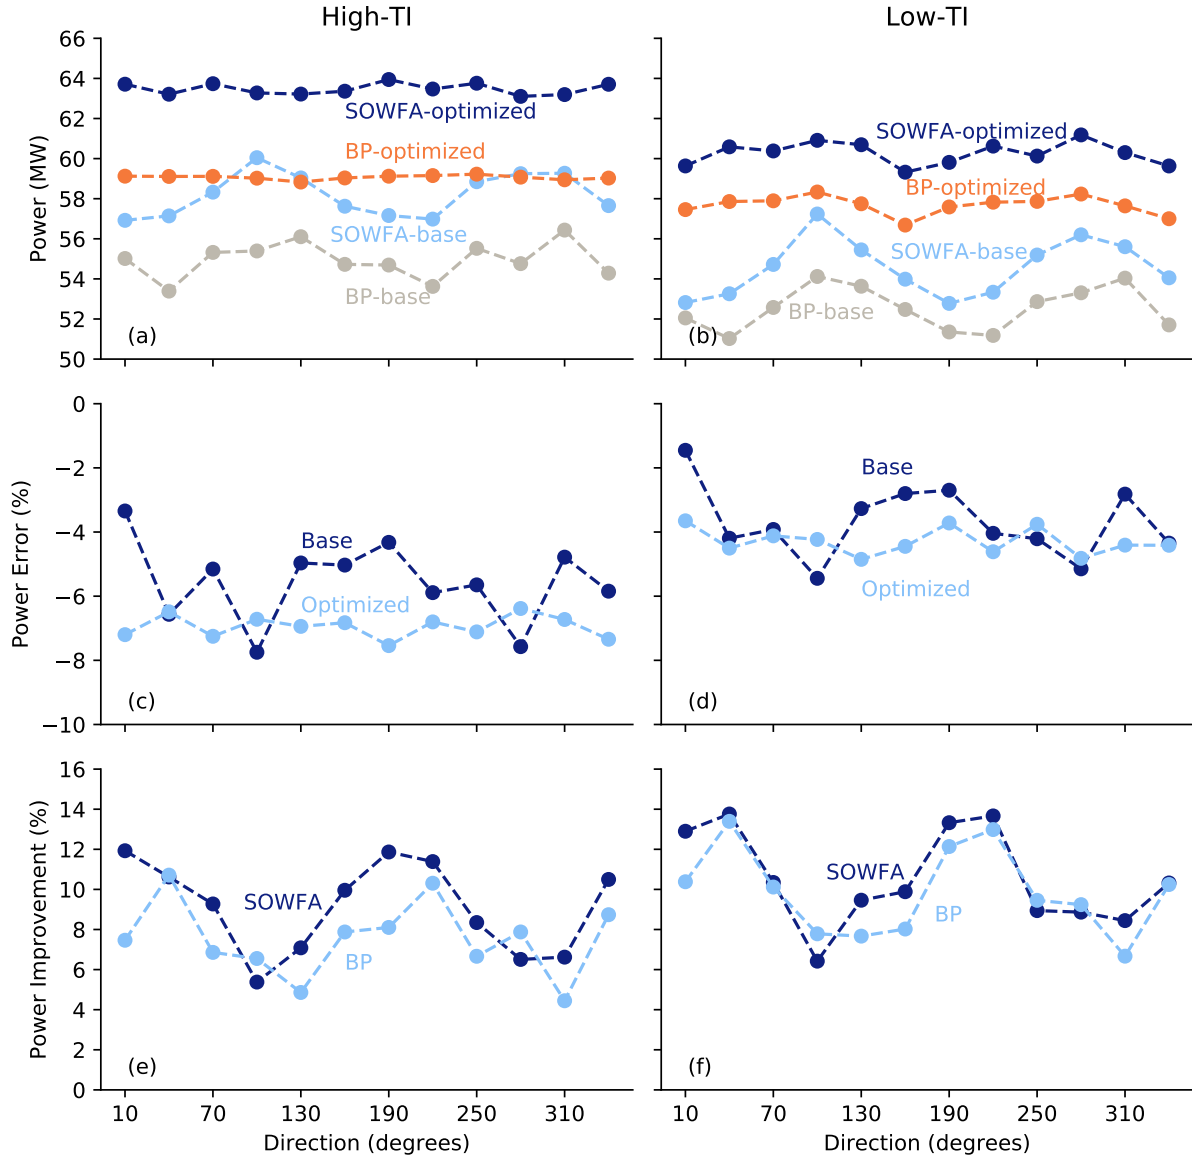


Figure 12. Power ((a) and (b)), power error ((c) and (d)), and power improvement ((e) and (f)) in each direction for both the high-TI ((a), (c), and (e)) and low-TI ((b), (d), and (f)) cases using 100 samples across the rotor for both the base layout (Fig. 6) and the optimized layouts (Fig. 9).

Despite the directional power prediction differences, both the BP model and SOWFA predicted increased power production for every direction for the optimized layout as compared to the base case layout. As logically follows, both the BP model and SOWFA predict significant AEP improvement, with SOWFA predicting greater improvements than the BP model. For the

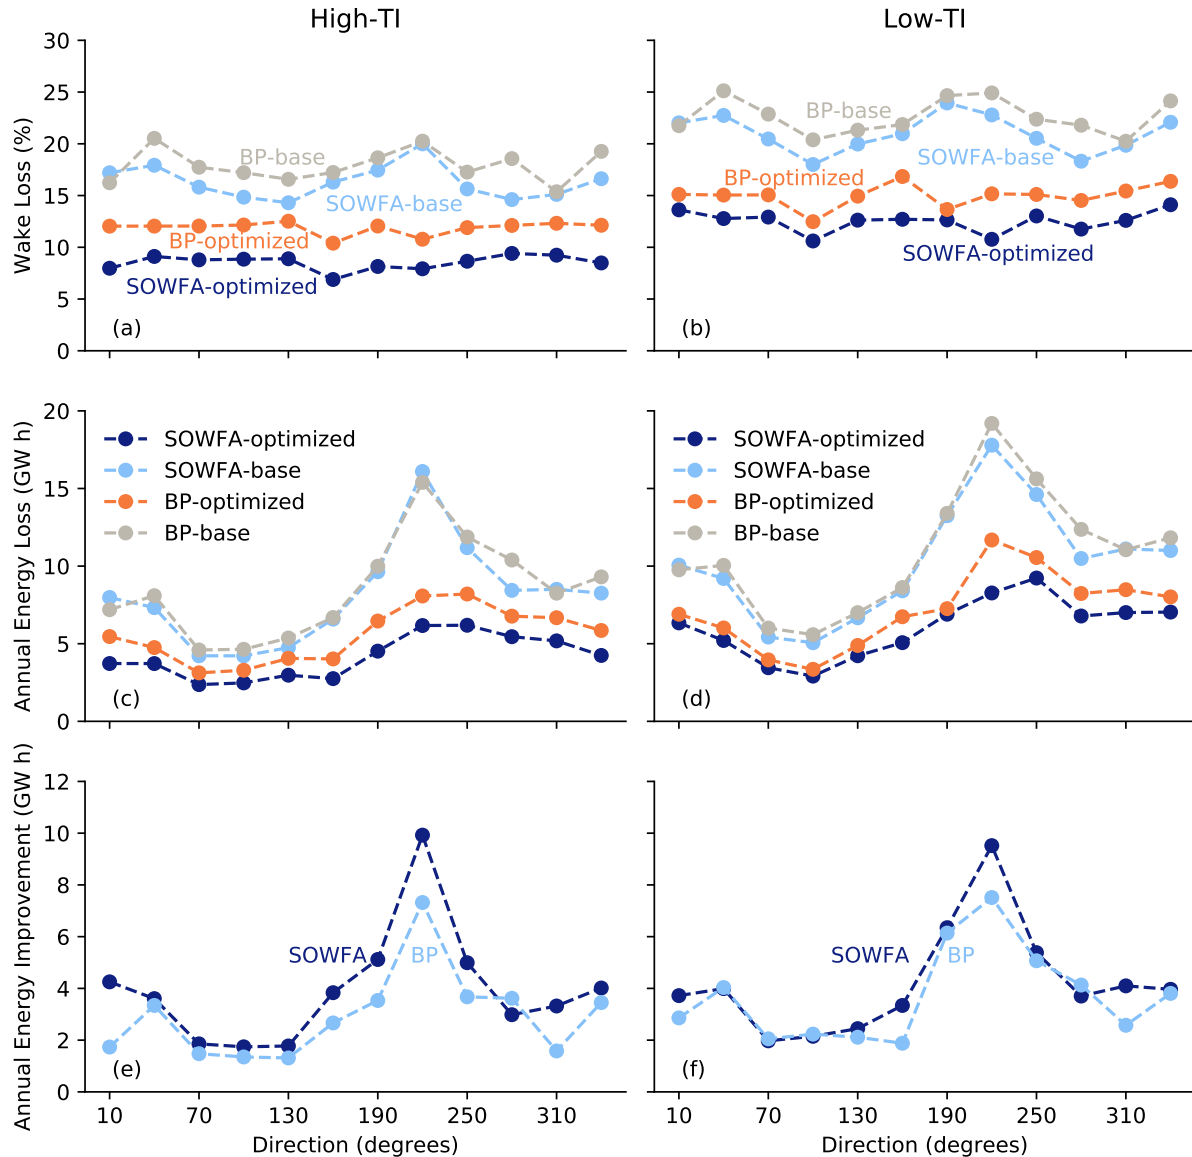


Figure 13. Wake loss ((a) and (b)), annual energy loss ((c) and (d)), and annual energy improvement ((e) and (f)) in each direction for both the high-TI ((a),(c), and (e)) and low-TI ((b), (d), and (f)) cases for both the base layout (Fig. 6) and the optimized layouts (Fig. 9).

high-TI case, the BP model predicted a 7.7% AEP increase, while SOWFA predicted a 9.3% AEP increase. For the low-TI case, the BP model predicted a 10.0% AEP increase, while SOWFA predicted a 10.7% AEP increase. The AEP percent errors between the BP model and SOWFA are similar to the directional percent errors, falling in the 3%–8% range for the high-TI case and the 1%–6% range for the low-TI case. While the directional percent errors are similar to the AEP errors, the range



of individual wind turbine power errors is greater than the directional errors. The errors for power production for each turbine in each wind direction are shown in Fig. 14 and Fig. 15. Power production errors for individual turbines, normalized by the maximum turbine power from SOWFA, ranged from -11% – 28% for the high-TI case base layout, -10% – 20% for the high-TI optimized layout, -18% – 18% for the low-TI case base layout, and -10% – 17% for the low-TI optimized layout. It is notable that the low-TI turbine errors are more evenly distributed (higher and lower than SOWFA) than the high-TI turbine errors. Despite the turbine errors, the BP model still allows the optimization to reduce wake effects, which is the primary driver for increasing AEP in a real wind farm.

We have found that each time we average turbine powers to get directional power, or average directional power to calculate AEP, the error decreases. The maximum errors at each level are: 28% for wind turbine power, 7.7% for directional power, and 6.9% for AEP. This is as expected because we know from the law of large numbers that the more samples that are combined to get an averaged result, the more precise that result tends to be. The directional power and AEP values can be seen as scaled means of the power production of the individual wind turbines. We can then place some confidence in AEP and directional power predictions from the BP model. We have shown that the overall improvements in directional power for this case are not just a factor of the simplified wake model because the AEP and directional improvements were shown to be even greater, according to SOWFA simulations, than predicted by the BP model.

4 Conclusions

In this work, we made some adjustments to the 2016 Bastankhah and Porté-Agel wake model (Bastankhah and Porté-Agel, 2016) and the 2016 Niayifar and Porté-Agel wind farm model (Niayifar and Porté-Agel, 2016) for improved compatibility with gradient-based optimization methods. We then demonstrated that the implementation of the altered models met or exceeded the accuracy of the models presented by Bastankhah, Niayifar, and Porté-Agel as compared to the LES data presented by Niayifar and Porté-Agel (2016). We described our test case, which we ran with both high- and low-turbulence intensity (TI), detailed our gradient-based optimization methods, and described the SOWFA setup used to produce the LES results that we used to check the model and optimization results. We found that the individual turbine predictions based on the models had errors up to 28% as compared to our SOWFA simulations. The directional and AEP errors were much less, up to only 7.7% and 6.9% , respectively. The model predicted an improvement in AEP of 7.7% for the high-TI case, while SOWFA predicted 9.3% . For the low-TI case, the model predicted a 10.0% AEP improvement, while SOWFA predicted 10.7% . Based on these results, we conclude that the improvements in directional power and AEP, at least for this case, can be trusted and are not simply an artifact of the simplified models being used during optimization.

While the simulation models did not appear to be exploited for unrealistic gains in AEP, the wind rose did have a significant impact. The final layouts exhibited a nested hexagonal pattern with the 12 combined sides of two concentric hexagons formed by the wind turbines being offset about halfway between the 12 wind directions. Many turbines were aligned in rows along unused directions in the wind rose. Upon further investigation, we found that at least 50 wind directions would be needed to remove the impact of wind direction fidelity on the final optimized AEP.




Figure 14. High-TI power production error of individual turbines in each direction as compared to SOWFA for the base layout (a) (see Fig. 9(a)) and optimized layout (b) (see Fig. 9(c)). Cells with the maximum and minimum errors as compared to SOWFA are shown with black borders. The color bar applies to both (a) and (b).

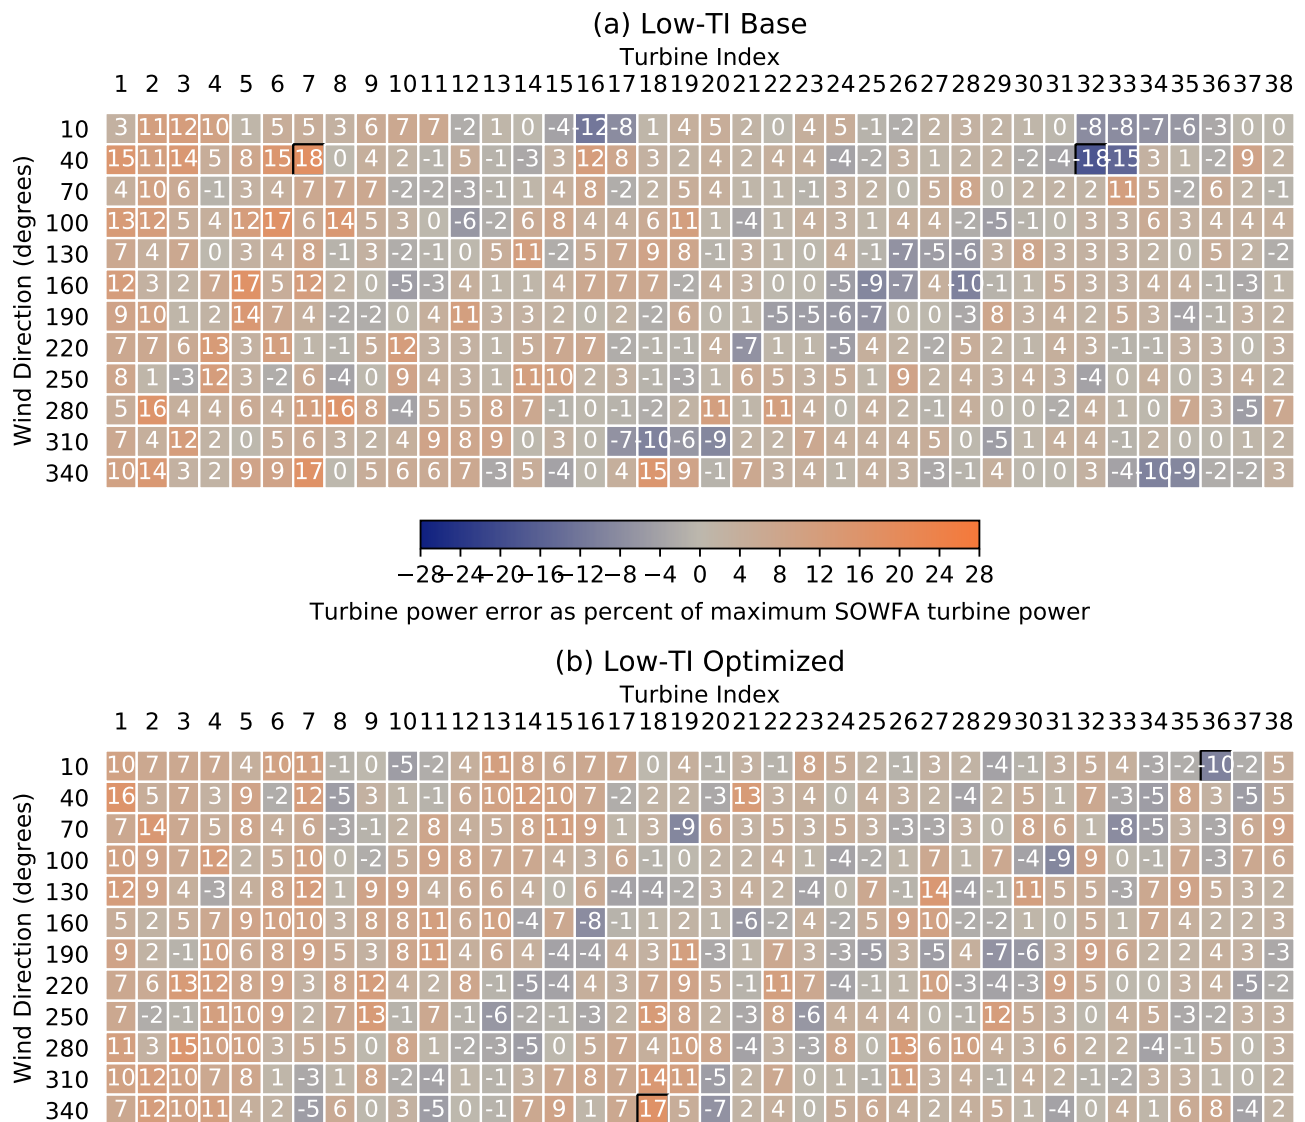


Figure 15. Low-TI power production error of individual turbines in each direction as compared to SOWFA for the base layout (a) (see Fig. 9(d)) and optimized layout (b) (see Fig. 9(f)). Cells with the maximum and minimum errors as compared to SOWFA are shown with black borders. The color bar applies to both (a) and (b).



Future work should include LES comparisons of optimization results using higher-fidelity wind roses that provide less opportunity for the optimization to take advantage of unused wind directions. The impact of wind rose fidelity in both wind direction and wind speed on optimization results should be investigated. It would be useful to investigate optimizations that account for TI as a varied quantity, using a similar approach to wind speed, wind direction, and wind frequency. It would also be beneficial to extend the data-based TI model given in Eq. (10) to account for a wider range of TI values.



Code and data availability. The numerical data presented in this paper, and the associated code, are available for reference in an online repository located at <https://zenodo.org/record/5786815>. The archive is hosted by Zenodo (European Organization For Nuclear Research and OpenAIRE, 2013).

440 *Author contributions.* Jared J. Thomas set up and ran the BP model, performed the results analysis, generated most of the figures, and wrote the bulk of the paper. Christopher J. Bay set up and performed the SOWFA runs and did related analysis and writing. Andrew P. J. Stanley contributed to code development, figure generation, and writing. Andrew Ning provided direction, ideas, and feedback throughout the development of this work. All authors contributed conceptually to this work.

Competing interests. The authors declare no competing interests.

445 *Acknowledgements.* This work was authored [in part] by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by the U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Wind Energy Technologies Office. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or
450 reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

A portion of the research was performed using computational resources sponsored by the DOE Office of Energy Efficiency and Renewable Energy and located at the National Renewable Energy Laboratory.

A portion of the research was performed using computational resources managed by the Office of Research Computing at Brigham Young University.

455 The authors would like to thank Jennifer King, Paul Fleming, Garrett Barter, and John Jasa for their feedback and discussion of the research presented here.



References

- Andersen, S. J., Sørensen, J. N., Ivanell, S., and Mikkelsen, R. F.: Comparison of Engineering Wake Models with CFD Simulations, *J PHYS CONF SER* (Online), 524, 2014.
- 460 Annoni, J., Seiler, P., Johnson, K., Fleming, P., and Gebraad, P.: Evaluating Wake Models for Wind Farm Control, in: American Control Conference, American Automatic Control Council, New York, 2014.
- Barthelmie, R. J., Hansen, K., Frandsen, S. T., Rathmann, O., Schepers, J. G., Schlez, W., Phillips, J., Rados, K., Zervos, A., Politis, E. S., and Chaviaropoulos, P. K.: Modelling and Measuring Flow and Wind Turbine Wakes in Large Wind Farms Offshore, *Wind Energy*, 12, 431–444, 2009.
- 465 Barthelmie, R. J., Folkerts, L., Larsen, G. C., Rados, K., Pryor, S. C., Frandsen, S. T., Lange, B., and Schepers, G.: Comparison of Wake Model Simulations with Offshore Wind Turbine Wake Profiles Measured by Sodar, *J ATMOS OCEAN TECH*, 23, 2014.
- Bastankhah, M. and Porté-Agel, F.: A New Analytical Model for Wind-Turbine Wakes, *RENEW ENERG*, 70, 116–123, 2014.
- Bastankhah, M. and Porté-Agel, F.: Experimental and Theoretical Study of Wind Turbine Wakes in Yawed Conditions, *J FLUID MECH*, 806, 506–541, 2016.
- 470 Beaucage, P., Brower, M., Robinson, N., and Alonge, C.: Overview of Six Commercial and Research Wake Models for Large Offshore Wind Farms, in: *EWEA Proceedings*, 2012.
- Belegundu, A. D. and Chandrupatla, T. R.: *Optimization Concepts and Applications in Engineering*, Cambridge University Press, 2nd edn., 2011.
- Bezanson, J., Edelman, A., Karpinski, S., and Shah, V. B.: Julia: A fresh approach to numerical computing, *SIAM REV*, 59, 65–98, <https://doi.org/10.1137/141000671>, 2017.
- 475 Bokharaie, V. S., Bauweraerts, P., and Meyers, J.: Wind-Farm Layout Optimisation Using a Hybrid Jensen–LES Approach, *Wind Energy Science*, 1, 311–325, <https://doi.org/10.5194/wes-1-311-2016>, 2016.
- Chowdhury, S., Zhang, J., Messac, A., and Castillo, L.: Unrestricted Wind Farm Layout Optimization (UWFLO): Investigating Key Factors Influencing the Maximum Power Generation, *RENEW ENERG*, 38, 16–30, 2012.
- 480 Churchfield, M. and Lee, S.: NWTC Information Portal (SOWFA), <https://nwtc.nrel.gov/SOWFA>, 2015.
- Churchfield, M. J., Lee, S., Michalakes, J., and Moriarty, P. J.: A Numerical Study of the Effects of Atmospheric and Wake Turbulence on Wind Turbine Dynamics, *J TURBUL*, 13, N14, <https://doi.org/10.1080/14685248.2012.668191>, 2012a.
- Churchfield, M. J., Lee, S., Moriarty, P. J., Martinez, L. A., Leonardi, S., Vijayakumar, G., and Brasseur, J. G.: A Large-Eddy Simulation of Wind-Plant Aerodynamics, *AIAA paper*, 537, 2012, 2012b.
- 485 Cook, J. D.: How to Compute the Soft Maximum, <https://www.johndcook.com/blog/2010/01/20/how-to-compute-the-soft-maximum/>, 2010.
- Crespo, A. and Hernández, J.: Turbulence Characteristics in Wind-Turbine Wakes, *J WIND ENG IND AEROD*, 61, 71 – 85, [https://doi.org/http://dx.doi.org/10.1016/0167-6105\(95\)00033-X](https://doi.org/http://dx.doi.org/10.1016/0167-6105(95)00033-X), 1996.
- Crespo, A., Hernández, J., and Frandsen, S.: Survey of Modelling Methods for Wind Turbine Wakes and Wind Farms, *Wind Energy*, 2, 1–24, 1999.
- 490 Elkinton, C. N., Manwell, J. F., and McGowan, J. G.: Algorithms for Offshore Wind Farm Layout Optimization, *Wind Engineering*, 32, 67–83, 2008.
- European Organization For Nuclear Research and OpenAIRE: Zenodo, <https://doi.org/10.25495/7GXK-RD71>, 2013.



- Fleming, P., Gebraad, P., van Wingerden, J.-W., Lee, S., Churchfield, M., Scholbrock, A., Michalakes, J., Johnson, K., and Moriarty, P.: The SOWFA Super-Controller: A High-Fidelity Tool for Evaluating Wind Plant Control Approaches, in: Proceedings of the EWEA Annual Meeting, Vienna, Austria, 2013.
- Fleming, P., Ning, A., Gebraad, P., and Dykes, K.: Wind Plant System Engineering Through Optimization of Layout and Yaw Control, Wind Energy, 2015.
- Frandsen, S., Barthelmie, R., Pryor, S., Rathmann, O., Larsen, S., and Højstrup, J.: Analytical Modelling of Wind Speed Deficit in Large Offshore Wind Farms, Wind Energy, 9, 39–53, 2006.
- Gaumond, M., Réthoré, P.-E., Bechmann, A., Ott, S., Larsen, G. C., Peña Diaz, A., and Hansen, K. S.: Benchmarking of Wind Turbine Wake Models in Large Offshore Wind Farms, in: Making Torque from the Wind Conference, Oldenburg, Germany, 2012.
- Gebraad, P., Thomas, J. J., Ning, A., Fleming, P., and Dykes, K.: Maximization of the Annual Energy Production of Wind Power Plants by Optimization of Layout and Yaw-Based Wake Control, Wind Energy, 20, 97–107, <https://doi.org/10.1002/we.1993>, 2017.
- Gebraad, P. M. O., Teeuwisse, F. W., van Wingerden, J. W., Fleming, P. A., Ruben, S. D., Marden, J. R., and Pao, L. Y.: Wind Plant Power Optimization Through Yaw Control Using a Parametric Model for Wake Effects—a CFD Simulation Study, Wind Energy, 2014.
- Gill, P., Murray, W., and Saunders, M.: SNOPT: an SQP Algorithm for Large-Scale Constrained Optimization, SIAM REV, 47, 99–131, 2005.
- Göçmen, T., van der Laan, P., Réthoré, P.-E., Peña Diaz, A., Larsen, G. C., and Ott, S.: Wind Turbine Wake Models Developed at the Technical University of Denmark: A Review, RENEW SUST ENERG REV, 60, 752–769, 2016.
- González, J. S., Rodríguez, A. G. G., Mora, J. C., Santos, J. R., and Payan, M. B.: Optimization of Wind Farm Turbines Layout Using an Evolutive Algorithm, RENEW ENERG, 35, 1671 – 1681, <https://doi.org/10.1016/j.renene.2010.01.010>, 2010.
- Grady, S., Hussaini, M., and Abdullah, M.: Placement of Wind turbines Using Genetic Algorithms, RENEW ENERG, 30, 259 – 270, <https://doi.org/10.1016/j.renene.2004.05.007>, 2005.
- Graf, P., Dykes, K., Scott, G., Fields, J., Lunacek, M., Quick, J., and Rethore, P.-E.: Wind Farm Turbine Type and Placement Optimization, J PHYS CONF SER, 753, 062004, <http://stacks.iop.org/1742-6596/753/i=6/a=062004>, 2016.
- Guirguis, D., Romero, D. A., and Amon, C. H.: Toward Efficient Optimization of Wind Farm Layouts: Utilizing Exact Gradient Information, APPL ENERG, 179, 110 – 123, <https://doi.org/10.1016/j.apenergy.2016.06.101>, 2016.
- Herbert-Acero, J. F., Probst, O., Réthoré, P.-E., Larsen, G. C., and Castillo-Villar, K. K.: A Review of Methodological Approaches for the Design and Optimization of Wind Farms, Energies, 7, 2014.
- Jasak, H., Jemcov, A., Tukovic, Z., et al.: OpenFOAM: A C++ Library for Complex Physics Simulations, in: International workshop on coupled methods in numerical dynamics, vol. 1000, pp. 1–20, IUC Dubrovnik, Croatia, 2007.
- Jensen, N. O.: A Note on Wind Generator Interaction, Tech. rep., Risø National Laboratory, DK-4000 Roskilde, Denmark, 1983.
- Jonkman, J., Butterfield, S., Musial, W., and Scott, G.: Definition of a 5-MW Reference Wind Turbine for Offshore System Development, National Renewable Energy Laboratory, Golden, CO, Technical Report No. NREL/TP-500-38060, 2009.
- Keane, A., Aguirre, P. E. O., Ferchland, H., Clive, P., and Gallacher, D.: An Analytical Model for a Full Wind Turbine Wake, J PHYS CONF SER, 753, 032039, <http://stacks.iop.org/1742-6596/753/i=3/a=032039>, 2016.
- Kusiak, A. and Song, Z.: Design of Wind Farm Layout for Maximum Wind Energy Capture, RENEW ENERG, 35, 685–694, 2010.
- Machefaux, E., Larsen, G. C., and Leon, J. P. M.: Engineering Models for Merging Wakes in Wind Farm Optimization Applications, J PHYS CONF SER, 625, 012037, <http://stacks.iop.org/1742-6596/625/i=1/a=012037>, 2015.



- 530 Martinez, L., Leonardi, S., Churchfield, M., and Moriarty, P.: A Comparison of Actuator Disk and Actuator Line Wind Turbine Models and Best Practices for Their Use, in: 50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition, p. 900, 2012.
- Moorthy, C. B., Deshmukhb, M., and Mukherejeeec, D.: New Approach for Placing Wind Turbines in a Wind Farm Using Genetic Algorithm, *Wind Engineering*, 38, 633–642, 2014.
- 535 Moriarty, P., Rodrigo, J. S., Gancarski, P., Chuchfield, M., Naughton, J. W., Hansen, K. S., Machefaux, E., Maguire, E., Castellani, F., Terzi, L., Breton, S.-P., and Ueda, Y.: IEA-Task 31 WAKEBENCH: Towards a Protocol for Wind Farm Flow Model Evaluation. Part 2: Wind Farm Wake Models, *J PHYS CONF SER*, 524, 012 185, <http://stacks.iop.org/1742-6596/524/i=1/a=012185>, 2014.
- Mosetti, G., Poloni, C., and Diviacco, B.: Optimization of Wind Turbine Positioning in Large Windfarms by Means of a Genetic Algorithm, *J WIND ENG IND AEROD*, 51, 105 – 116, [https://doi.org/https://doi.org/10.1016/0167-6105\(94\)90080-9](https://doi.org/https://doi.org/10.1016/0167-6105(94)90080-9), 1994.
- 540 Niayifar, A. and Porté-Agel, F.: A New Analytical Model for Wind Farm Power Prediction, in: *J PHYS CONF SER*, 2015.
- Niayifar, A. and Porté-Agel, F.: Analytical Modeling of Wind Farms: A New Approach for Power Prediction, *Energies*, 9, 1–13, 2016.
- Perez-Moreno, S. S., Dykes, K., Merz, K. O., and Zaaijer, M. B.: Multidisciplinary Design Analysis and Optimisation of a Reference Offshore Wind Plant, *J PHYS CONF SER*, 1037, 042 004, <http://stacks.iop.org/1742-6596/1037/i=4/a=042004>, 2018.
- Réthoré, P.-E., Fuglsang, P., Larsen, G. C., Buhl, T., Larsen, T. J., and Aagaard Madsen, H.: TOPFARM: Multi-Fidelity Optimization of
- 545 Wind Farms, *Wind Energy*, 17, 1797–1816, 2014.
- Revels, J., Lubin, M., and Papamarkou, T.: Forward-Mode Automatic Differentiation in Julia, arXiv:1607.07892 [cs.MS], <https://arxiv.org/abs/1607.07892>, 2016.
- Rios, L. M. and Sahinidis, N. V.: Derivative-Free Optimization: a Review of Algorithms and Comparison of Software Implementations, *J GLOBAL OPTIM*, 56, 1247–1293, <https://doi.org/10.1007/s10898-012-9951-y>, 2013.
- 550 Thomas, J., Gebraad, P., and Ning, A.: Improving the FLORIS Wind Plant Model for Compatibility with Gradient-Based Optimization, *Wind Engineering*, 41, 313–329, <https://doi.org/10.1177/0309524X17722000>, 2017.
- Thomas, J. J., Annoni, J., Fleming, P. A., and Ning, A.: Comparison of wind farm layout optimization results using a simple wake model and gradient-based optimization to large eddy simulations, in: *AIAA Scitech 2019 Forum*, p. 0538, 2019.
- Thomas, J. J., McOmber, S., and Ning, A.: Wake Expansion Continuation: Multi-Modality Reduction in the Wind Farm Layout Optimization
- 555 Problem, *Wind Energy*, (in press), 2021.
- Turner, S., Romero, D., Zhang, P., Amon, C., and Chan, T.: A New Mathematical Programming Approach to Optimize Wind Farm Layouts, *RENEW ENERG*, 63, 674–680, 2014.
- Western Regional Climate Center: Station Wind Rose: Nantucket, https://wrcc.dri.edu/cgi-bin/wea_windrose.pl?laKACK, 2012.