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, 2022) 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). In our case in particular, using the Sparse Nonlinear OPTimizer (SNOPT), linear constraints may be used to define where the functions may be safely evaluated but the constraints on turbine positions are non-linear which means they may be enforced by minimizing the constraint violations (Gill et al., 2005). SNOPT may viloate the constraints at certain steps in the optimization or even in the solution if the constraints turn out to be infeasible. Another consideration in defining the near wake is that the near-wake region of the BP model varies in length and may sometimes be within the defined optimization constraints.

where 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 defined by Bastankhah and Porté-Agel (2016) and shown below.

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

$$\sigma_z = k_z[x - x_0] + \frac{d}{\sqrt{8}}, \tag{3}$$

where 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 by Bastankhah and Porté-Agel (2016) and reproduced below.

$$\frac{x_0}{d} = \frac{\cos\gamma\left[1+\sqrt{1-C_T}\right]}{\sqrt{2}\left[\alpha^* I + \beta^*\left[1-\sqrt{1-C_T}\right]\right]}, \tag{4}$$

where 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

$$u = u_r\left[\frac{z - z_o}{z_r - z_o}\right]^{\psi}, \tag{5}$$

where 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 surface roughness height, 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 Eq. (1) is undefined when the term inside the square root becomes negative, which means that Eq. (1) is undefined when

$$\frac{C_T\cos\gamma}{8\sigma_y\sigma_z/d^2} > 1. \tag{6}$$

In order to be useful, the discontinuity must be defined in terms of the downstream position in the wake. By substituting Eqs. (2) and (3) into Eq. (6), and separating the terms based on the powers of $[x - x_0]$, we obtain

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

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

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

The model, then, is undefined for regions where $x < x_d$. To remove the discontinuity, we set σ_y and σ_z to be constant, their values at $x = x_0$, everywhere between the rotor hub and x_0. With this change, Eqs. 2 and 3 become

$$\sigma_y = \begin{cases} k_y[x - x_0] + \frac{d\cos\gamma}{\sqrt{8}} & x > x_0 \\ \frac{d\cos\gamma}{\sqrt{8}} & 0 \le x \le x_0 \end{cases} \tag{9}$$

and

$$\sigma_z = \begin{cases} k_z[x - x_0] + \frac{d}{\sqrt{8}} & x > x_0 \\ \frac{d}{\sqrt{8}} & 0 \le x \le x_0. \end{cases} \tag{10}$$

Because $x_0 > x_d$, this change removes the discontinuity. While Eqs. 9 and 10 do not create an accurate near-wake model, they do provide good enough information for the optimization algorithm to continue the optimization and move turbines out of nearby wakes.

The results of the model changes discussed above are shown in Fig. 1. These changes are only intended to make the model compatible with gradient-based optimization and are not intended to provide an accurate near-wake model for simulation purposes. The 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 an accurate near-wake model is needed it may be possible to replace our rough near-wake model with something better, such as the near-wake model proposed by Keane et al. (2016), for final calculations.

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.

where $A_{w_{ij}}$ is the area of the wake of turbine i at the downstream location of turbine j, A_{r_j} 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 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 as

$$I_{+_{ij}} = 0.73 a_i^{0.8325} I_i^{0.0325} \left[\frac{x}{d_i} \right]^{-0.32}, \tag{18}$$

where 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 as

$$a = \begin{cases} 0.143 + \sqrt{0.0203 - 0.6427[0.889 - C_T]} & C_T > 0.96 \\ 0.5[1 - \sqrt{1 - C_T}] & \text{otherwise.} \end{cases} \tag{19}$$

We finally calculated total TI as (Niayifar and Porté-Agel, 2016)

$$I_j = \sqrt{I_0^2 + I_{+_j}^2}. \tag{20}$$

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 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. To generate the pattern, we used a set of equations from the work of Vogel (1979) as presented by de Jong (2013), with a slight modification from an answer on Stack Overflow (user3717023, 2016). The sample points on the rotor-swept area were calculated as

$$\theta = c \frac{2\pi}{\phi^2}, \tag{21}$$

$$r = \frac{\sqrt{c - 1/2}}{\sqrt{n - 1/2}}, \tag{22}$$

where θ is the angle to the point location from the horizontal axis in the plane of the rotor-swept area, c is the current point number, ϕ is the golden ratio, r is the radial distance from the center of the rotor-swept area, and n is the total number of points. This approach allowed our point function to return a normalized set of well-distributed points within a radius of 1 regardless of how many points we requested. We then scaled the pattern to cover the full rotor-swept area. The resulting sample point pattern is 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 shown in Fig. 2(b). We performed all analyses and provided all results using 100 sample points across the rotor-swept areas.

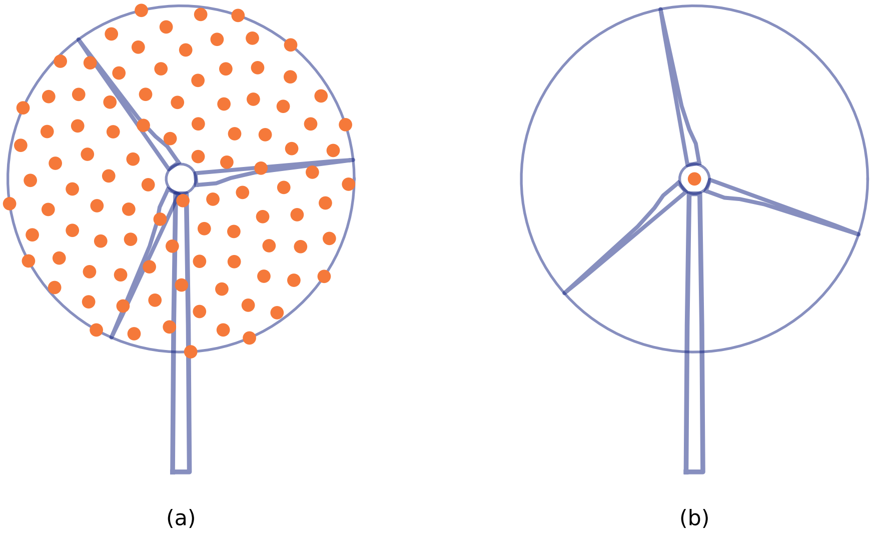

Figure 2. Location of wind speed sampling points, shown as orange dots, for approximating the inflow wind speed at each turbine. (a) Pattern used for LES comparisons and all provided results. (b) Single point used for approximating effective inflow wind speed during optimization.

2.2.4 AEP and wind turbine power

We calculated AEP in W h as

$$AEP = \left[\frac{hours}{day}\right]\left[\frac{days}{year}\right]\sum_{k=1}^{N_D}\left[f_k\sum_{j=1}^{N_T}P_{j_k}\right], \tag{23}$$

where 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_{j_k}, is calculated based on the definition of the power coefficient as done by Gebraad et al. (2014) and shown here as

$$P_{j_k} = \begin{cases} 0 & \bar{u}_j < u_{cutin} \\ 0.5\rho A_j C_{P_j}\bar{u}_j^3 & \bar{u}_j \geq u_{cutin} \text{ and } P_j < P_{rated} \\ P_{rated} & P_j \geq P_{rated}, \end{cases} \tag{24}$$

where ρ is the air density, A_j represents the rotor-swept area of turbine j, and C_{P_j} 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.

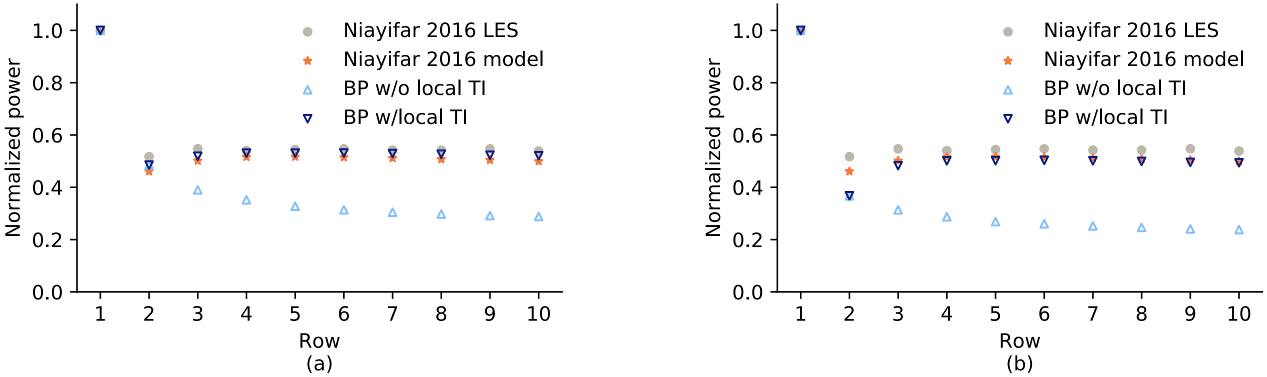

Figure 3. Averaged power of turbines in columns two, three, and four of the Horns Rev wind farm by turbine row. (a) Comparing our code using 100 points to approximate the effective wind speed at the rotor. (b) Comparing our code using a single sample at the rotor hub to approximate the effective wind 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 buoyancy and Earth rotation (Coriolis) effects in the atmosphere. The bottom boundary is modeled using Schumman's model for surface stress. Aerodynamic surface roughness values of 0.0002 m and 0.15 m for the low- and high-TI cases respectively were chosen to generate flows with approximately 5% and 10% TI. The resulting planar-

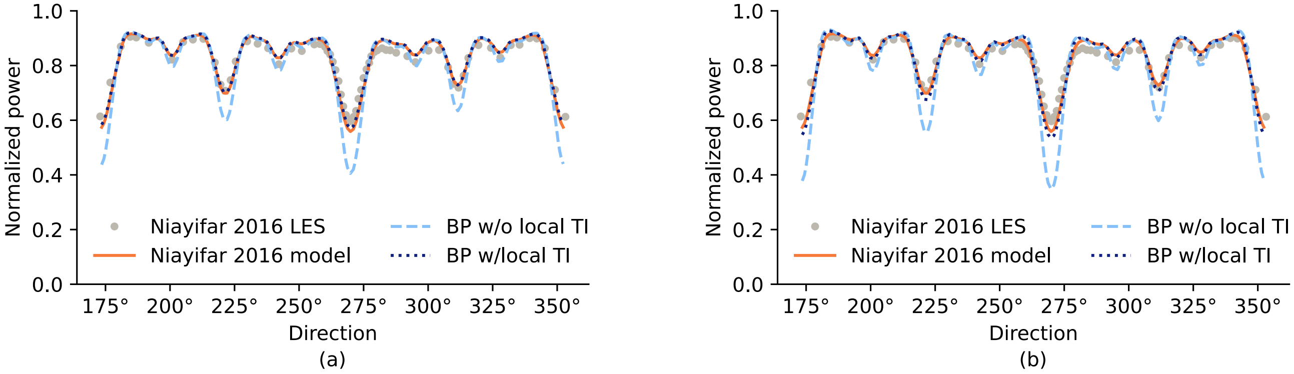

(a)

(b)

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.

averaged velocity profiles are shown in Fig. 5. A two-step approach is used with SOWFA simulations containing turbines. First, turbulence is spun-up using a computational domain with periodic boundary conditions on the lateral boundaries. This step, known in the literature as precursor, is executed until the turbulence is fully developed, identified by convergence of planar-average metrics such as stresses and turbulence intensity. In this work, for the atmospheric conditions and turbulence desired, this step took 20,000 s. Next, the precursor is further executed for an additional 7,200 s with its boundary information saved at each time step. Finally, a second simulation is setup, containing the turbine models, and executed for the 7,200 s of available boundary information. At this step, the inflow boundary conditions are given by information saved from the precursor stage, and a zero gradient boundary condition for the velocity is used on the outflow boundaries. Domains are $5 \times 5 \times 1$ km in extent, with 10 m uniform grid resolution in the x, y, and z directions.

SOWFA uses the time-accurate 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 a supercomputer using a few hundred to a few thousand processors, depending on the size of the wind plant. Here, 25-million-cell simulations were executed, totaling 36 hours of walltime, using 14 36-core nodes. 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 the need of different precursor for each of the 12 directions, the layout was rotated to a reference wind direction of 270°. The precursor cases executed for these analyses have a canonical neutrally stratified atmospheric boundary layer. The boundary layer height is enforced by providing the potential temperature gradient at the desired height. Here, we set a capping inversion of a 5 K difference over a 100 m thick inversion layer, effectively capping the boundary layer at 750 m. We set the mean wind speed to be 8 m s-1 at a height of 90 m.

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_o to 0 m. We fit the model of Eq. (5) to the LES data by varying u_r and ψ 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).

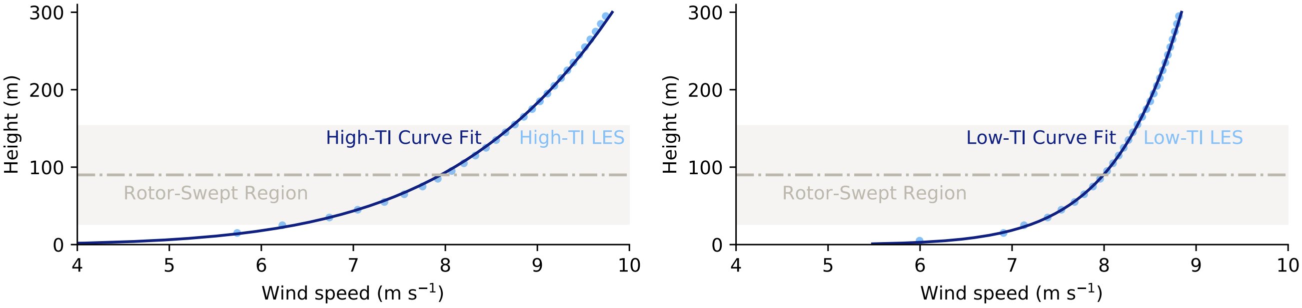

(a)

(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 estimated as

$$I = \frac{\sqrt{\sum u_x'^2}}{\bar{u}_x}, \tag{25}$$

where u_x' is the zero-mean turbulent fluctuation of the streamwise component at the hub height for all instants, and $\bar{u}_x$ is the average wind speed at hub height in the wind direction. Only the streamwise fluctuations were considered for the computation of the TI. The simulations were setup such that, while Coriolis is active, its reference height is the hub height, effectively inducing no veer at the analysis height of 90 m. The lack of temperature stratification in the neutral boundary layer simulated means no buoyancy effects are present and the vertical fluctuations are only due to the low-roughness planetary surface. Thus, the observed fluctuations in the spanwise and vertical directions were small when compared to those in the streamwise direction. Average TI was estimated to be 5.1% for the low-TI precursor and 10.1% for the high-TI precursor. We set the inflow wind speed in the BP model 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. 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. 16 and 17.

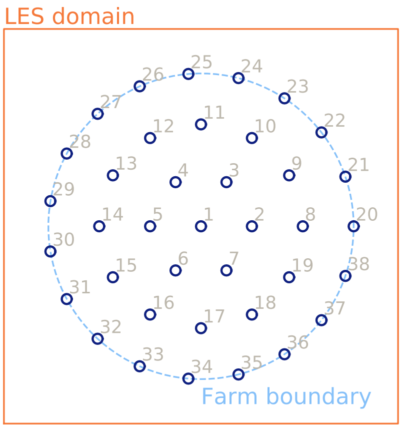

Figure 6. Base case wind farm layout. The circles marking turbine locations are to scale, with diameters equal to the rotor diameter. Turbine numbers are provided for comparison to results figures (Figs. 16 and 17).

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.

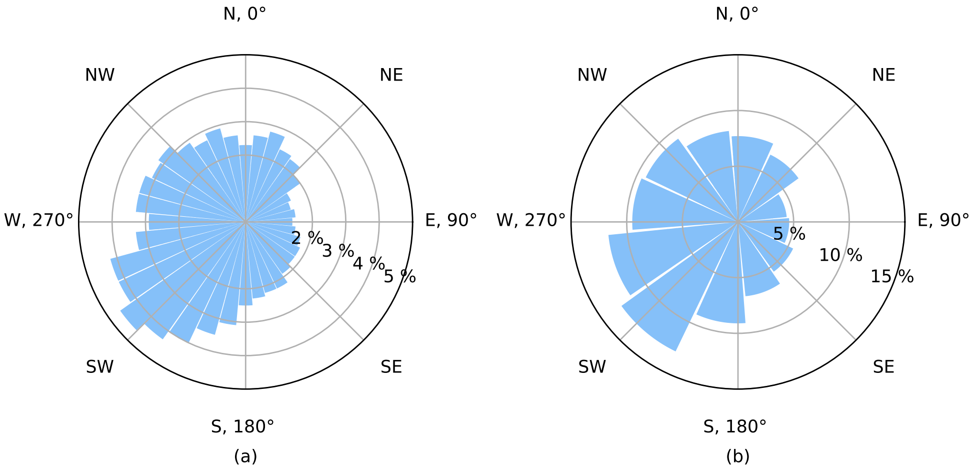

Figure 7. (a) Nantucket frequency wind rose (Western Regional Climate Center, 2012). (b) Nantucket frequency wind rose binned into 12 directions.

2.7 Optimization

We optimized the wind farm layout using AEP as the objective. Our optimization problem was formulated as

$$
\begin{aligned}
\text{maximize} \quad & AEP(x_i, y_i) \;\; i = 1 \ldots 38 \\
\text{subject to} \quad & S_{ij} \geq 2d \;\; i,j = 1 \ldots 38 \;\; i \neq j \\
& [x_c - x_i]^2 + [y_c - y_i]^2 \leq r_b^2 \;\; i = 1 \ldots 38
\end{aligned}
\tag{26}
$$

where (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.

We set up and ran the wind farm simulations in Julia (Bezanson et al., 2017) using FLOWFarm.jl[1]. Exact derivatives of the full simulation model and the constraints were obtained using forward-mode algorithmic differentiation (AD) provided by

---

[1] https://github.com/byuflowlab/FLOWFarm.jl

the ForwardDiff.jl package (Revels et al., 2016). Derivatives obtained through AD are more accurate (Griewank and Walther, 2008) and faster to calculate (Gray et al., 2014) than derivatives estimated using finite difference methods. The derivatives of the objective function and the constraints were scaled to be between ± 1. We then optimized the final problem using 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 SNOW.jl[2] to connect FLOWFarm.jl, ForwardDiff.jl, and SNOPT.

We used the wake expansion continuation (WEC) method for reducing the multimodality of the wind farm layout optimization problem (Thomas et al., 2022). The WEC method works by widening the turbine wakes, while maintaining the original velocity deficits at the wake centers. The widened wakes with larger deficits then dominate the wakes with smaller deficits, removing at least some local optima from the design space in a manner similar to Gaussian continuation optimization as presented by Mobahi and Fisher (2015). When using WEC, the optimization problem is run once for each WEC factor from largest to smallest and the wakes are brought back to their standard widths for the final optimization. The result of each optimization in the series is used as the starting layout for the next optimization. In this way we avoid many local optima while not losing any accuracy in the final result. We used the same WEC factors as Thomas et al. (2022), [3.0 2.6 2.2 1.8 1.4 1.0 1.0]. The WEC factors, when applied to the BP model as in this work, are wake diameter multipliers. We applied the WEC factors to the BP model (see Eq. 1) as follows:

$$\frac{\Delta u_{ij}}{\bar{u}_i} = \left[1 - \sqrt{1 - \frac{C_T \cos\gamma}{8\sigma_y\sigma_z/d^2}}\,\right] \exp\left(-0.5\left[\frac{y-\delta}{\xi\sigma_y}\right]^2\right) \exp\left(-0.5\left[\frac{z-z_h}{\xi\sigma_z}\right]^2\right), \tag{27}$$

where ξ is the WEC factor being applied.

We 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.

We ran the full optimization process with 400 different starting locations for each TI level. The starting locations were determined by randomly and individually rotating each of the concentric circles of turbines in the starting layout around the boundary center (see Fig. 6). We used this method to obtain multiple starting layouts while maintaining good coverage of the available wind farm area and significant separation between turbines.

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 recovery in the high-TI case makes it easier to avoid wake interactions than in the low-TI case. The final optimized layouts

---

[2] https://github.com/byuflowlab/SNOW.jl

directions and then again using 50 wind directions. The AEP results for each number of sample points were then recalculated using 100 sample points for comparison. Results are shown in Fig. 12. We found that the number of sample points had only minimal impact on the optimization results. The largest difference in the average AEP change as recalculated using 100 rotor sample points (see Fig. 12) was 2.79 % when using 12 wind directions and 1.75 % when using 50 wind directions. While the number of rotor samples does not appear to significantly impact the average optimization results, the results distribution did vary depending on the number of rotor sample points. We also found that about 20 rotor sample points are needed to obtain the same AEP change predictions as calculated using 100 rotor sample points. Based on these results, we consider a single rotor sample point to be sufficient for wind farm layout optimization, but recommend that at least 20 rotor sample points be used for AEP estimation when using the sunflower seed packing algorithm for rotor sample point placement. More or fewer points may be needed for AEP estimation if a different rotor sample point placement algorithm were used.

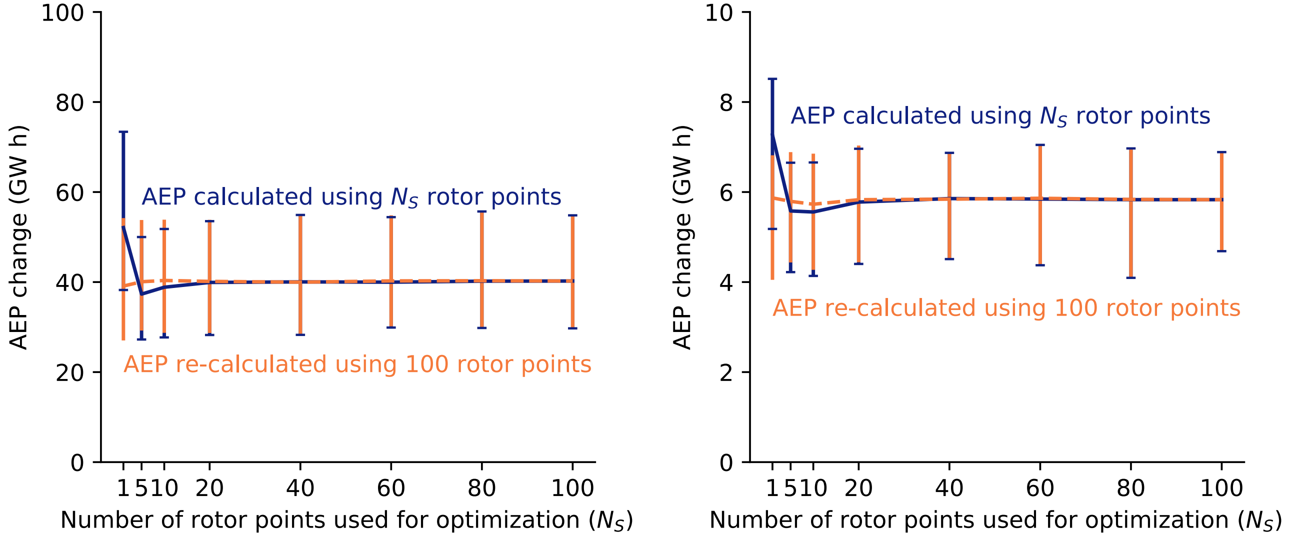

(a) Using 12 wind directions.

(b) Using 50 wind directions.

Figure 12. This figure shows that the number of rotor sample points has only minimal impact on the optimization results, but that about 20 rotor sample points are needed to predict the same AEP as when using 100 rotor sample points. Results are shown with AEP calculated using the number of rotor sample points in the optimizations (N_S) and also with AEP recalculated (not re-optimized) using 100 rotor sample points 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 sample points.

Due to limited computational resources, we were constrained to an LES simulation domain that may be too small and could have caused artificial speedup effects. The 5 km-by-5 km domain had a total CPU time per simulation of approximately 756 days, or 36 hours on fourteen 36-core nodes. This study required 48 simulations (one for each direction, for each TI level, for both the baseline and the optimized layouts), for a total CPU time of about 72 days on fourteen 36-core nodes. We estimate that to reduce the artificial speedup effects sufficiently would require a domain size of approximately 15 km-by-15 km. However, even just increasing the domain to 10 km-by-10 km would require approximately 6 days on fourteen 36-core nodes per simulation, or around 288 days on fourteen 36-core nodes to complete all the simulations for all directions, which is more than the computational resources we had available for this project.

As we will show later, the turbine powers from the LES results are higher than those from the BPA model. The higher power may be due to higher wind speeds stemming from the domain size and artificial speedup effects. The speedup effects are related to the blockage. As a proxy for blockage, we used the percent of the outlet plane blocked by a perpendicular projection of the rotor swept areas. The approximate blockage by direction for each case is shown in Fig. 13. While the approximate blockage values of around 5 to 8 percent are high, the average change across all directions between the base and optimized cases is only 1.19 and 0.90 percentage points for the low- and high-TI cases respectively. While blockage effects are almost certainly present, the optimization algorithm cannot exploit them because the optimization algorithm has no knowledge of the LES. Because the change in estimated blockage between the base and optimized cases is small, and the optimizer was unable to exploit the blockage effects, we believe that these simulations are sufficient for our purpose in this work to demonstrate that the improvements from the baseline layout to the optimized layout are reasonable. However, we have adjusted the SOWFA results for the optimized layouts to account for the additional blockage.

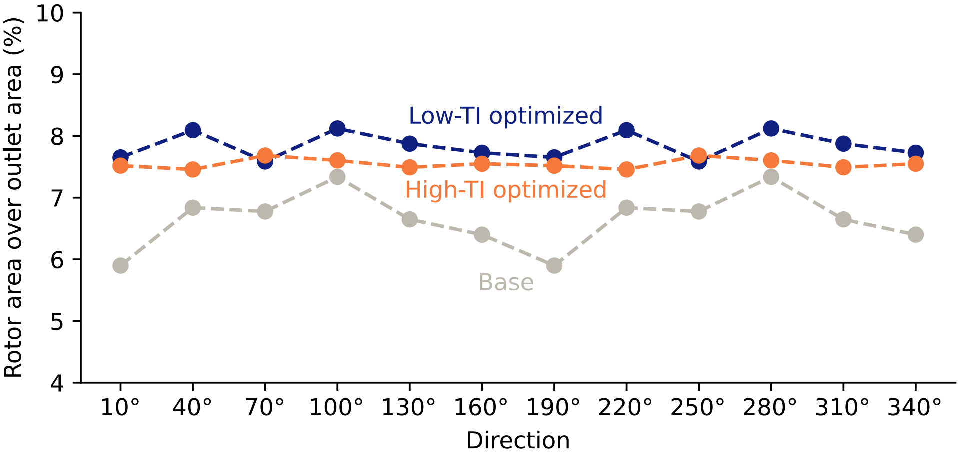

Figure 13. Cumulative rotor area projection divided by the outlet area as a percentage, used as a proxy for blockage of the LES domain. Plot includes values for each wind direction for each case. Towers were ignored for this calculation.

To account for added blockage in the SOWFA results for the optimized layouts, we estimated the increase in wind speed due to blockage based on the conservation of mass. We chose A_1 and $\bar{u}_1$ to be the original flow area and velocity, and A_2 and $\bar{u}_2$ to be the reduced flow area and velocity. As SOWFA solves the incompressible equations, assuming constant density, we can write the conservation of mass in this case to be

$$A_1\bar{u}_1 = A_2\bar{u}_2. \tag{28}$$

If we assume that the reduced flow area is related to the original flow area through the blockage by

$$A_2 = (1-b)A_1, \tag{29}$$

where b is the blockage, then we get that $\bar{u}_2$ should be

$$\bar{u}_2 = \frac{\bar{u}_1}{1-b}. \tag{30}$$

We can further estimate the expected percent change in power due to blockage (P_b') based on Eq. 24 as

$$P_b' = \frac{0.5\rho A C_P \bar{u}_2^3 - 0.5\rho A C_P \bar{u}_1^3}{0.5\rho A C_P \bar{u}_1^3}. \tag{31}$$

If we assume equal C_P and constant density in the numerator and denominator, then Eq. 31 reduces to

$$P_b' = \frac{\bar{u}_2^3 - \bar{u}_1^3}{\bar{u}_1^3}. \tag{32}$$

If we then substitute in from Eq.30 we get that

$$P_b' = \frac{(\bar{u}_1/(1-b))^3}{\bar{u}_1^3} - 1, \tag{33}$$

which reduces to

$$P_b' = \frac{1}{(1-b)^3} - 1. \tag{34}$$

By adjusting the percent change in power for the optimized layouts' SOWFA results using Eq. (34) with the change in estimated blockage between the base and optimized layouts from Fig. 13, we were able to approximate what some of the results would have been had we been able to use a sufficiently large LES domain. Blockage adjustments were applied to the results on a per-direction basis, not to individual turbines. In the following discussion, results do not include adjustments for blockage unless otherwise specified.

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. 14 (a) and Fig. 14 (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. It is probable that a similar offset would be seen in Figs. 4 and 3 if the results were not normalized. 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.

The error in the directional power is shown in Fig. 14 (c) and Fig. 14 (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 is shown in Fig. 14 (e) and Fig. 14 (f). The directional power improvement 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).

Improvements in the directional power, adjusted to account for additional blockage in the SOWFA results for the optimized layout, are shown in Fig. 14 (g) and Fig. 14 (h). Here we see that accounting for blockage reduces the directional power improvements according to SOWFA to below the BP model. This makes sense because the BP model does not capture deep array and unsteady effects that would be captured in the SOWFA results. Of particular importance is that the trends in the adjusted SOWFA results are much closer to the trends in the BP results than the unadjusted SOWFA results. In particular, the adjusted SOWFA results for the high-TI case show the same direction for greatest improvement as the BP results, that is 220°.

The directional power results can be misleading because we are using a weighted wind rose. To make more general insights into what is going on across the wind directions, we first switch from AEP to wake loss, which is non-dimensional. We calculated wake loss as

$$L_w = 100\left(1 - \frac{P}{P_I}\right), \tag{35}$$

where the ideal power, P_I, was determined as the power of the wind farm in a given direction with the same inflow conditions but as if there were no wake interactions. We determined the free stream power in each direction as the average power of 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 including the blockage adjustment for the optimized layouts is shown in Fig. 15(a) and Fig. 15(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 curves, and we can see that the BP model is certainly capturing important effects.

Using the adjusted wake loss and the directional ideal power, we calculated the directional annual energy loss, shown in Fig. 15(c) and Fig. 15(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 less wake loss in the base case. While the spike in directional energy loss at 220° was mostly flattened through the optimization, the low-TI case exhibits a slight 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 (including adjustments for blockage in the SOWFA results for the optimized layouts) to indicate how much the total energy produced increased in each direction due to the optimization. The results are shown in Fig. 15(e) and Fig. 15(f). The close match of the BP and SOWFA

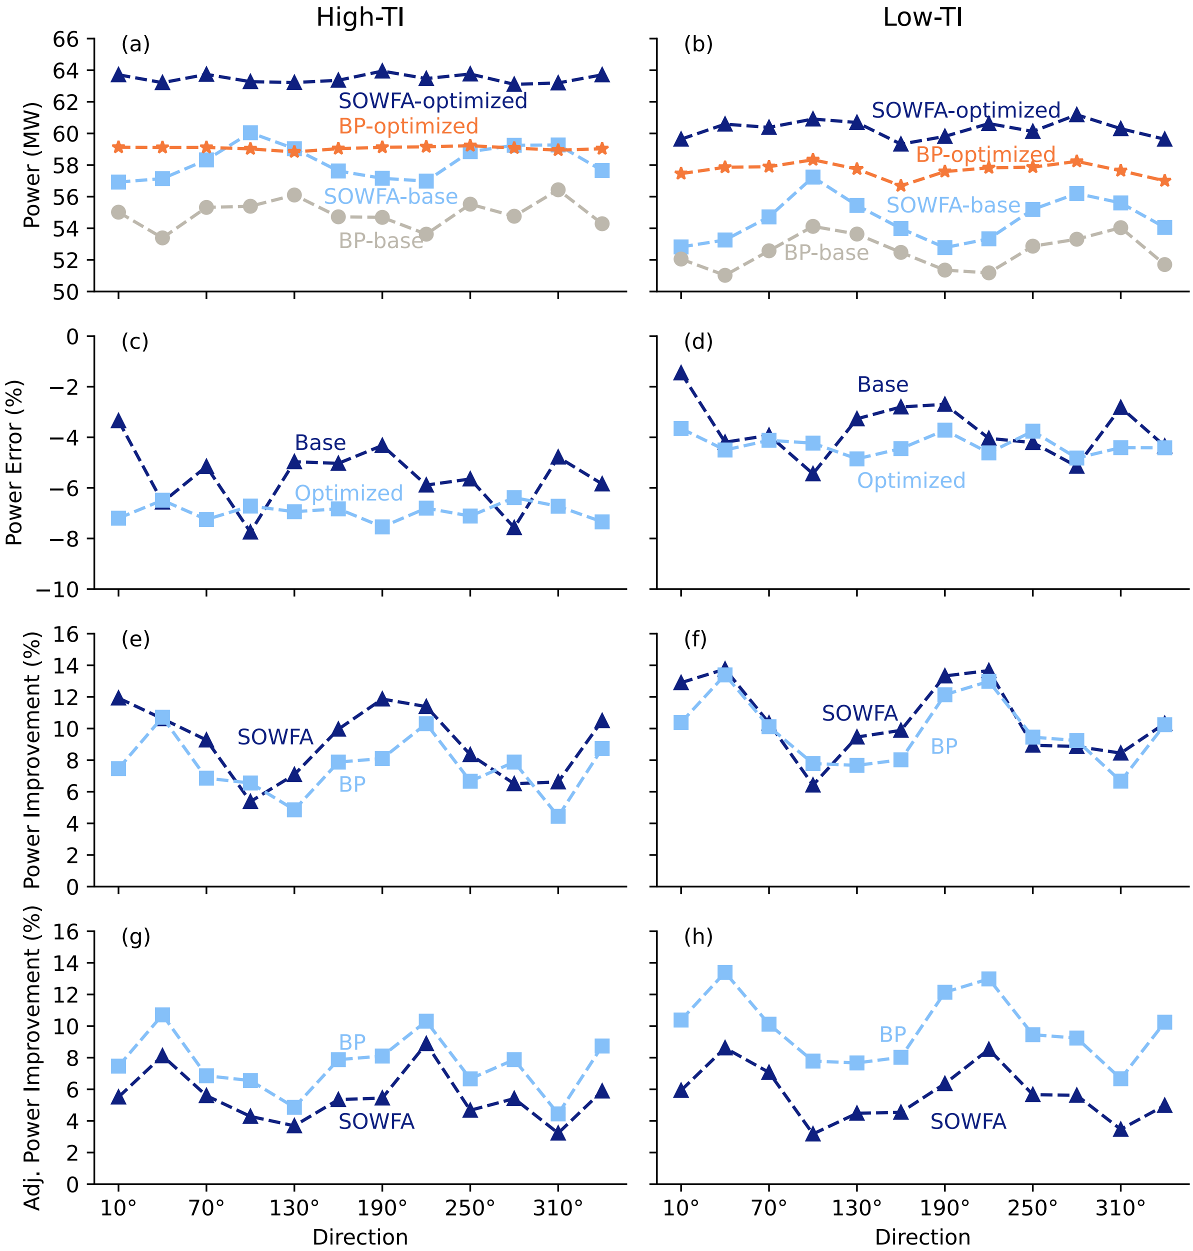

Figure 14. Power ((a) and (b)), power error ((c) and (d)), power improvement ((e) and (f)), and power improvement adjusted to account for speedup effects due to added blockage in the SOWFA results for the optimized case ((g) and (h)) in each direction for both the high-TI ((a), (c), (e), and (g)) and low-TI ((b), (d), (f), and (h)) cases using 100 samples across the rotor for both the base layout (Fig. 6) and the optimized layouts (Fig. 9).

results in directional annual energy loss is striking. Even though the actual power 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. The SOWFA results showed somewhat less improvement in each direction than the BP model, possibly due to the unsteady and deep array effects captured by SOWFA but not represented by the BP model.

Despite the directional power prediction differences, both the BP model and SOWFA predicted increased power 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 unadjusted high-TI case, the BP model predicted a 7.7% AEP increase, while SOWFA predicted a 9.3% AEP increase. For the unadjusted low-TI case, the BP model predicted a 10.0% AEP increase, while SOWFA predicted a 10.7% AEP increase. Adjusting the SOWFA results for the optimized layouts to account for the additional blockage shows a 5.2% and 5.6% AEP increase for the high- and low-TI cases respectively. The unadjusted 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 unadjusted 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 for each turbine in each wind direction are shown in Fig. 16 and Fig. 17. Power 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 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 substantial according to LES simulations.

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

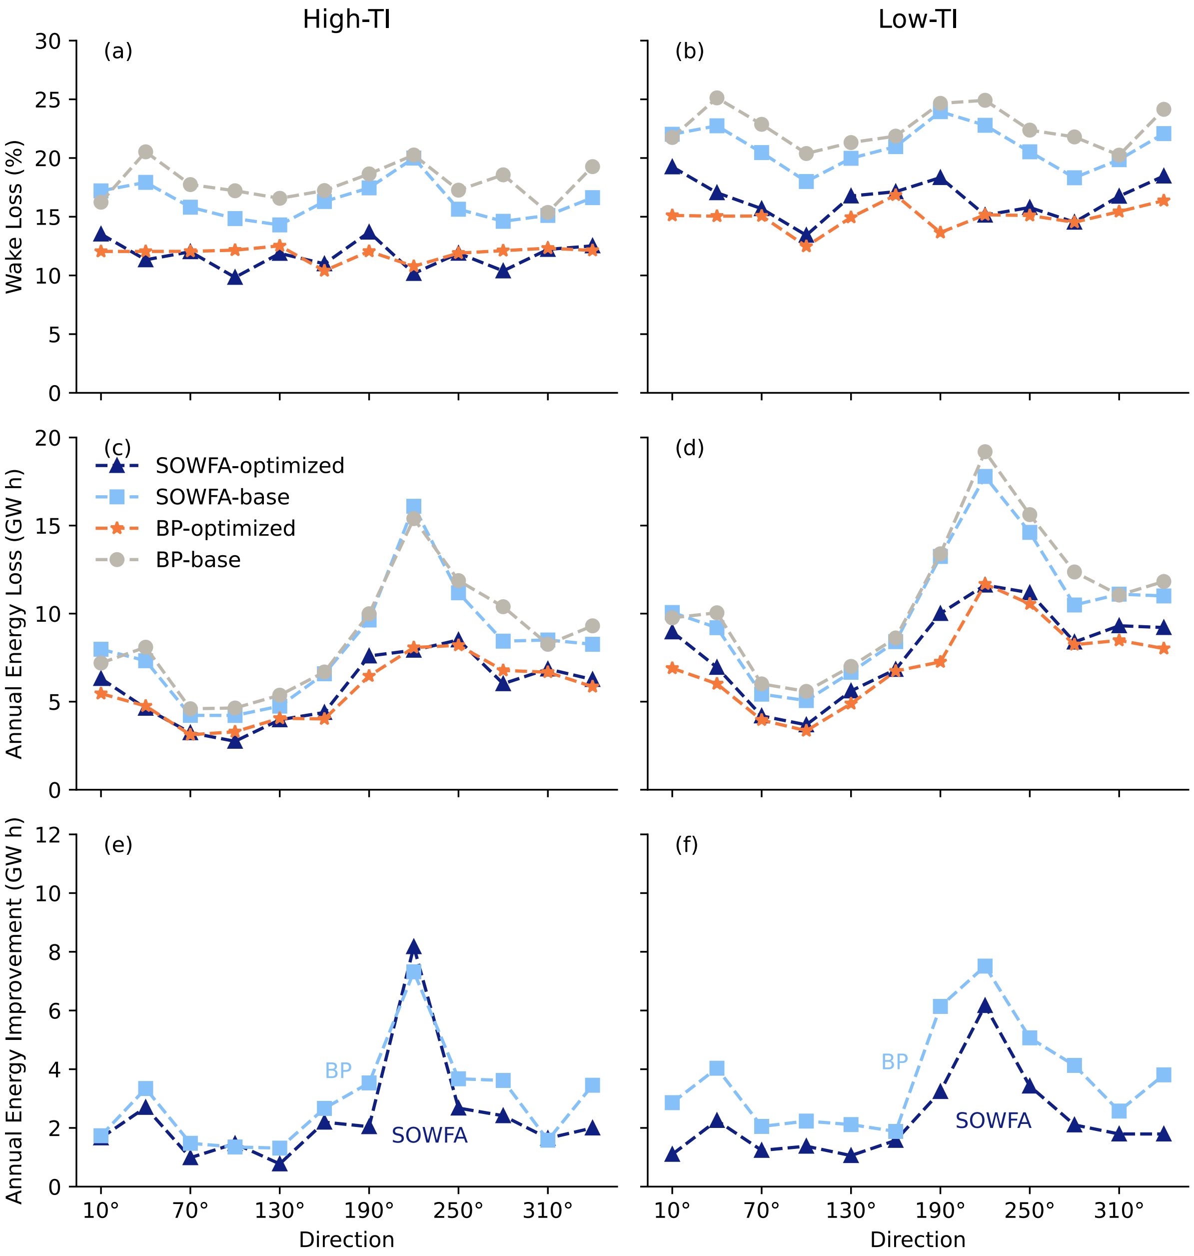

Figure 15. Wake loss ((a) and (b)), annual energy loss ((c) and (d)), and annual energy improvement ((e) and (f)) in each direction for both the high-TI ((a),(c), and (e)) and low-TI ((b), (d), and (f)) cases for both the base layout (Fig. 6) and the optimized layouts (Fig. 9). The legend in (c) applies to ((a) through (d)). The SOWFA results for the optimized layouts presented in this figure have been adjusted to account for speedup effects due to added blockage in the optimized case.