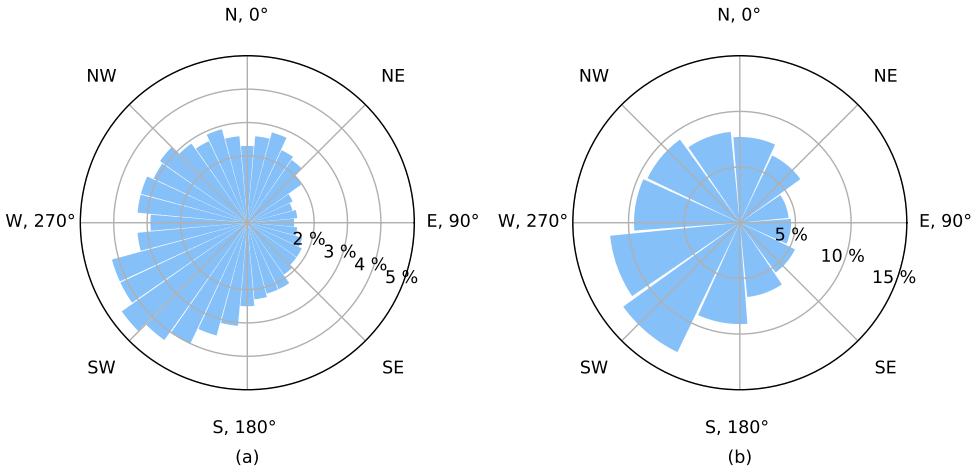

Figure 7. (a) Nantucket frequency wind rose (Western Regional Climate Center, 2012). (b) Nantucket frequency wind rose binned into 12 directions.

We set up and ran the wind farm simulations in Julia (Bezanson et al., 2017) using FLOWFarm.jl[1]. 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 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). 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:

$$\frac{\Delta u_{ij}}{\bar{u}_i} = \left[1 - \sqrt{1 - \frac{C_T \cos\gamma}{8\sigma_y \sigma_z / d^2}}\right] \exp\left(-0.5\left[\frac{y-\delta}{\xi\sigma_y}\right]^2\right) \exp\left(-0.5\left[\frac{z-z_h}{\xi\sigma_z}\right]^2\right), \tag{23}$$

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

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

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

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.jl [3]. 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 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.

To investigate the impact of the number of rotor sample points used during optimization, we ran the high-TI case using the same one hundred starting points as in the directional fidelity study, but varied the number of rotor sample points and held the number of wind directions constant. For a single sample, the point was placed at the hub. For all other numbers of samples we used the sunflower seed packing algorithm to determine the sample locations. We ran this same study using 12 wind 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.

[3] https://github.com/byuflowlab/FLOWMath.jl.git

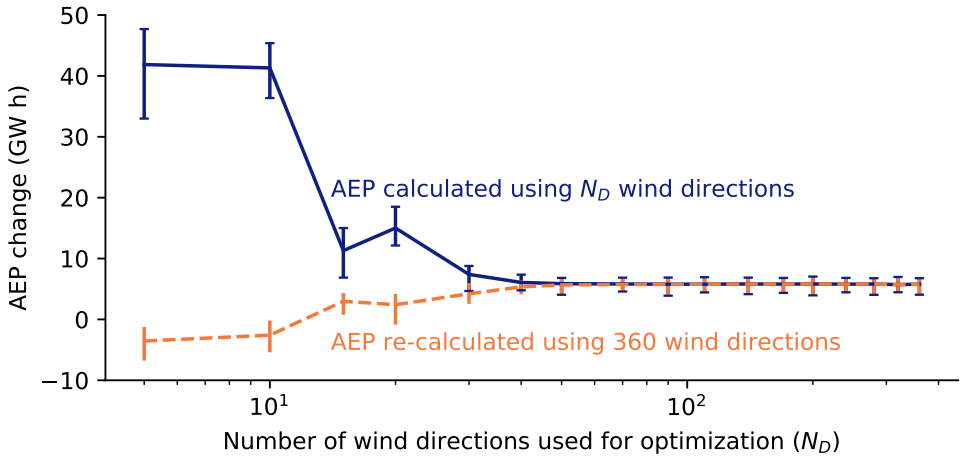

Figure 11. This figure shows that at least 40 or 50 wind directions are needed during optimization based on the average change in AEP from starting layout to optimized layout. Results are shown with AEP calculated using the number of wind directions in the optimizations (N_D) and also with AEP recalculated (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. 13 (a) and Fig. 13 (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 production is shown in Fig. 13 (c) and Fig. 13 (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. 13 (e) and Fig. 13 (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). However, the directional power results can be misleading because we are using a weighted wind rose.

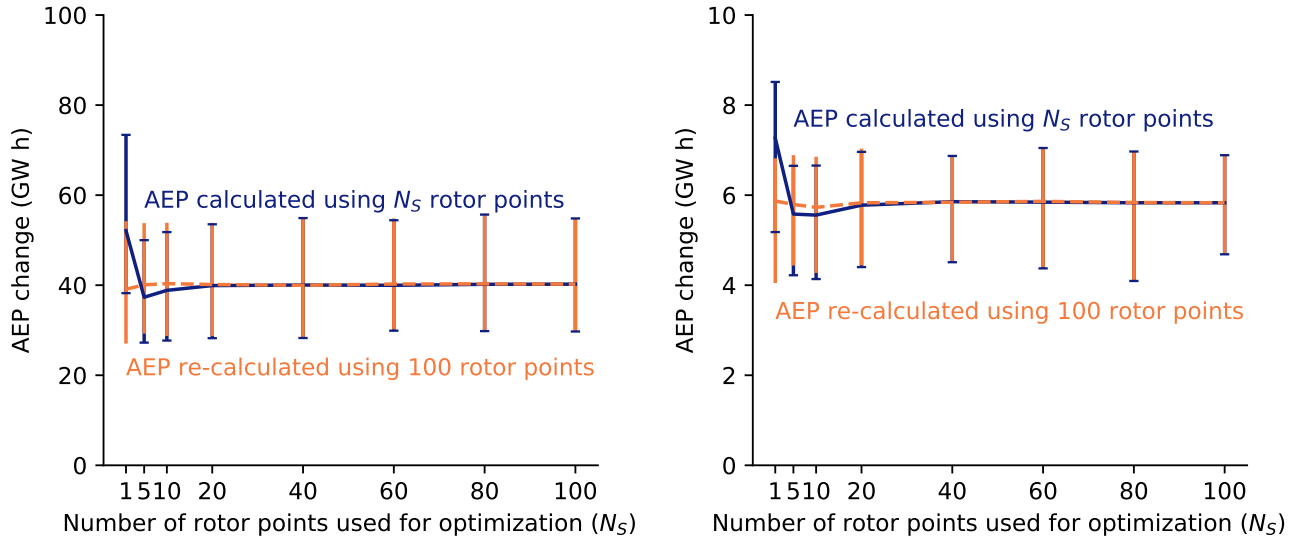

(a) Using 12 wind directions.

(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.

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 according to Eq. (24).

$$L_w = 100\left(1 - \frac{P}{P_I}\right), \tag{24}$$

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 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. 14(a) and Fig. 14(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 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. 14(c) and Fig. 14(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