



Massive Simplification of the Wind Farm Layout Optimization Problem

Andrew P. J. Stanley and Andrew Ning

Brigham Young University Department of Mechanical Engineering, 701 E University Pkwy, 350 EB, Provo, UT, 84602

Correspondence: Andrew P. J. Stanley (stanley_andrewpj@byu.net)

Abstract. The wind farm layout optimization problem is notoriously difficult to solve because of the large number of design variables and extreme multimodality of the design space. Because of the multimodality of the space and often discontinuous models used in wind farm modeling, the wind industry is heavily dependent on gradient-free techniques for wind farm layout optimization. Unfortunately, the computational expense required with these methods scales poorly with increasing numbers of variables. Thus, many companies and researchers have been limited in the size of wind farms they can optimize. To solve these issues, we present the boundary-grid parameterization. This parameterization uses only five variables to define the layout of a wind farm with any number of turbines. For a 100 turbine wind farm, we show that optimizing the five variables of the boundary-grid method produces wind farms that perform within 0.5% of farms where the location of each turbine is optimized individually, which requires 200 design variables. Our presented method unlocks the ability to optimize and study large wind farms, something that has been mostly infeasible in the past.

1 Introduction

In 2018, wind energy produced 6.6% percent of the electricity use in the United States¹. With current market trends and technology, the U.S. Energy Information Administration projects that this number will rise by 1% in both 2019 and 2020 (U.S. Energy Information Administration, 2019a), and the installed capacity will increase by 4% every year through 2050 (U.S. Energy Information Administration, 2019b). In order for the U.S. and the rest of the world to meet and exceed these projections, it is necessary to be able to create efficient turbine layouts for large wind farms. The wind farm layout optimization problem is notoriously difficult to solve because of the large number of design variables, computationally expensive models for high fidelity simulations, and extreme multimodality of the design space (see Fig. 1).

Because of the multimodality of the space and often discontinuous models used in wind farm modeling, the wind industry is heavily dependent on gradient-free techniques for wind farm layout optimization (Herbert-Acero et al., 2014). Although these methods can be highly effective for small numbers of design variables, the computational expense required to converge scales poorly, approximately quadratically, with increasing numbers of variables. Because of this poor computational scaling, many companies and researchers have been limited in the size of wind farms they can optimize, as the number of variables typically increases with the number of turbines. Figure 2 demonstrates this principle. This figure shows the number of function

¹<https://www.eia.gov/tools/faqs/faq.php?id=427&t=3>




Figure 1. The complexity and multimodality of wind farm layout design space. Shown is the normalized annual energy production of a 100 turbine wind farm as a function of the location of one turbine. 99 turbines remain fixed, while one is moved throughout the wind farm. (a) A 2-dimensional view of the design space. (b) A 3-dimensional surface which highlights the extreme variation of the peaks and valleys. This figure shows only the multimodality from two dimensions, where the true design space has 200 design variables

25 evaluations required to optimize the multi-dimensional Rosenbrock function verses the number of variables. To give a sense of what these numbers mean, if this problem with 64 variables and exact-analytic gradients takes one hour to optimize, using finite difference gradients would take almost four days, while a gradient-free method would take over 20 years! The trends, not the exact numbers, shown in this figure are general for other optimization problems, such as wind farm layout. As the size of the problem increases, the computational expense with certain optimization methods can become unmanageable.

30 Despite its difficulty, layout optimization is an essential step in wind farm development in order to maximize power production. Power losses of 10–20% are typical from turbine interactions within a wind farm (Barthelmie et al., 2007, 2009; Briggs, 2013), and can be as high as 30–40% for farms with closely spaced wind turbines (Stanley et al., 2019). However, because the difficulties in finding optimal turbine placement increase with the number of turbines, layout optimization can quickly become infeasible for large wind farms (Ning and Petch, 2016). Even so, accelerated research and understanding of the principles
35 governing wind energy, as well as public demand for renewable energy sources are encouraging developers and communities to install farms with more wind turbines than have been typical in the past. Current turbine layout definitions and optimization methods are woefully inadequate for these increasingly large farms.

The most common current wind farm layout definitions include defining the location of each turbine directly (Feng and Shen, 2015; Guirguis et al., 2016; Gebraad et al., 2017), preassigning some locations in a wind farm as suitable turbine locations to
40 limit size of the design space (Emami and Noghreh, 2010; Parada et al., 2017; Ju and Liu, 2019), and parameterizing the

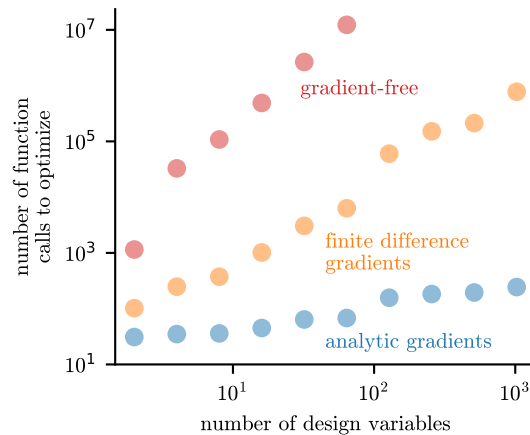


Figure 2. The number of function calls required to optimize the multi-dimensional Rosenbrock function versus the number of variables. The computational expense of gradient-free and finite difference gradients scale poorly with the number of variables.

turbines as a grid (González et al., 2017; Perez-Moreno et al., 2018). Defining the location of every wind turbine directly allows the most freedom, but also requires two variables for each turbine. In addition, the design space is the most multimodal. If one limits the design space by predetermining acceptable turbine locations or parameterizing the turbine locations with a simple grid, they are able to optimize larger wind farms. However, these methods produce simplistic wind farm designs, which
45 underperform for most realistic scenarios.

In this paper we present the boundary-grid (BG) layout parameterization, a new wind farm layout parameterization. This new method solves the challenges that have previously made wind farm layout optimization so difficult. BG parameterization uses only five variables, and can produce layouts that perform within 0.5% of the layouts achieved by directly optimizing the location of each wind turbine. With some of the most advanced wind farm optimization methods that have previously been
50 available, we can directly optimize the location of every turbine in a 100 turbine wind farm in 4–5 hours. More common methods take on the order of days or longer. With BG parameterization, we can optimize a 100 turbine wind farm in 3 minutes. Additionally, this new parameterization dramatically reduces the multimodality of the design space compared to direct layout optimization (compare Figs. 1 and 11b). Finally, BG parameterization has additional benefits, including a regular, aesthetically pleasing layout, naturally defined roads or shipping lanes, and easily defined cabling between turbines. This technique can
55 immediately be applied to wind farm design to obtain excellent wind farm layouts with limited computational resources.

2 Boundary-Grid Parameterization

When the locations of wind turbines in a farm are optimized directly, the final layout often follows two general rules. First, a large fraction of turbines are grouped on or near the wind farm boundary. Second, the turbines that are not positioned on the




Figure 3. Example 100 turbine wind farm layouts, and parameterized wind turbine layout definition. Each dot is to scale, representing the wind turbine diameter. (a) Wind farm layout when the position of each turbine has been optimized directly. This optimization required 200 design variables, the x and y location of each turbine. (b) Wind farm layout optimized with boundary-grid parameterization. This optimization required five design variables, shown in (c) and (d). (c) The start location design variable, s . (d) The four variables defining the inner grid: the grid spacing, dx and dy , the grid offset b , and the rotation, θ .

boundary are loosely arranged in rows throughout the farm (Fig. 3a). By observing these patterns in optimal wind farm layouts,
60 we defined our new layout parameterization such that it would create wind farms that filled these requirements.

2.1 New Layout Variables

In BG parameterization, the turbines are divided into two groups: the boundary and the inner grid (Fig. 3b). The boundary
turbines are spaced around the circumference of the wind farm and are defined with one design variable. The rest of the
turbines in the farm make up the inner grid, which is defined with four design variables for a total of five variables to describe
65 the location of every turbine in the farm. The boundary turbines are placed on the wind farm boundary, spaced equally traversing
the perimeter. These are defined by one variable, s , which is the distance along the perimeter where the first turbine, or start
turbine, is placed. This in turn defines the position of every turbine around the boundary (Fig. 3c). The inner grid turbines are
defined by four design variables: dx , dy , b , and θ . The grid spacing, dx and dy , are the distance between columns and rows in
70 the grid, b is the offset distance, which defines how far consecutive rows are offset, θ is the grid rotation angle, which rotates
the entire grid (Fig. 3d). The grid offset could also be defined as an angle, however we have used a distance as the gradients
are more conducive to optimization.



2.2 Selection of Discrete Values

There are some discrete values which are important in our formulation, namely the number of turbines which are placed along the boundary and how many are in the grid, how many rows and columns are in the grid, and how the rows and columns are organized. We present some rules that we have found effective in determining these discrete values for all wind roses, wind farm boundaries, and wake models that we tested. Each individual case may benefit slightly from a more specialized selection of these values but our method works well across all cases tested.

The number of turbines placed on the boundary is determined by the wind farm perimeter and turbine rotor diameter. If the perimeter is large enough, 45% of the wind turbines are placed on the boundary. In some cases, the wind farm perimeter is small, and would result in turbines that are too closely spaced if 45% were placed around the boundary. In this case, the number of boundary turbines is reduced until the minimum desired turbine spacing in the wind farm is preserved. No matter how many turbines are placed around the boundary, they are always spaced equally traversing the perimeter, and all of the remaining turbines are placed in the inner grid.

The number of rows, columns, and their organization in the grid is determined with the following procedure. First, dy is set to be four times dx , b is set such that turbines are offset twenty degrees from those in adjacent rows, and θ is initialized randomly. Then, dx is varied with θ remaining constant, and dy and b changing to fulfill the requirements prescribed in the initialization definition, until the correct number of turbines are within the wind farm boundary. During optimization, each of the grid variables can change individually, however the discrete values remain fixed.

3 Wind Farm Modeling

3.1 Wind Turbine Parameters

In the testing of BG wind farm layout parameterization method, we modeled the turbine parameters after the IEA 3.35-MW reference turbine². The relevant parameters are a rotor diameter of 130 meters, hub height of 110 meters, and a rated electrical power of 3.35 MW. The thrust coefficient was assumed idealized and constant, $C_T = 8/9$. The power curve was defined as a piecewise equation in Eq. 1.

$$P_i(V) = \begin{cases} 0 & V < V_{\text{cut-in}} \\ P_{\text{rated}} \left(\frac{V}{V_{\text{rated}}} \right)^3 & V_{\text{cut-in}} \leq V < V_{\text{rated}} \\ P_{\text{rated}} & V_{\text{rated}} \leq V < V_{\text{cut-out}} \\ 0 & V \geq V_{\text{cut-out}} \end{cases} \quad (1)$$

In this power curve definition, P_i is the power produced by an individual wind turbine, V is the hub velocity at that turbine, P_{rated} is 3.35 MW, V_{rated} is 10 m/s, $V_{\text{cut-in}}$ is 3 m/s, and $V_{\text{cut-out}}$ is 25 m/s.

²<https://www.nrel.gov/wind/assets/pdfs/se17-9-iea-wind-task-37-systems-engineering.pdf>



3.2 Wind Farm Details

The major benefit of wind turbine layout parameterization comes for large wind farms. For farms with just a few turbines, the layout can be optimized directly with a small amount of design variables. In such cases with few design variables, there is little to no benefit gained from intelligently parameterizing the design space. In this study, each wind farm layout that we optimized had 100 wind turbines, to demonstrate the benefits of BG parameterization for large wind farms.

We tested the performance of our parameterization method on wind farms with different average turbine spacing: four, six, and eight rotor diameters shown in Fig. 8. In addition to testing wind farms with different turbine spacing, we modeled and optimized several different wind farm boundaries in this study: the boundary of the Princess Amalia wind farm, a real farm in the North Sea (Van Dam et al., 2012; Gebraad and Van Wingerden, 2015; Kanev et al., 2018), a circle, and a square to demonstrate the sharp angles that can occur in wind farm boundaries. These boundaries are shown in Fig. 10.

3.3 Wake Model

We calculated the wind speed deficit in the wake behind a wind turbine with a simple Gaussian wake model, modified from that originally developed by Bastankhah and Porté-Agel (Bastankhah and Porté-Agel, 2016; Baker et al., 2019). The wake model is defined with Eqs. 2 and 3.

$$\sigma = kx + \frac{D}{\sqrt{8}} \quad (2)$$

$$l_i = \left(1 - \sqrt{1 - \frac{C_T}{8\sigma^2/D^2}}\right) \exp\left[\frac{-1}{2}\left(\frac{y}{\sigma}\right)^2\right] \quad (3)$$

In these equations, x is the distance between turbines for which the wake loss is being calculated in line with the wind direction, y is the cross stream distance between turbines, D is the rotor diameter of the wind turbines, k is a parameter calculated from the turbulence intensity and is equal to 0.0325, C_T is the thrust coefficient, and l_i is the wake loss contribution from a single turbine. This model is a simplification of the model proposed by Bastankhah and Porté-Agel by calculating all wake losses in the horizontal plane, and by assuming a potential core length of zero. The total loss at any point in the wind farm was given by the L2-norm of the individual loss contributions from each turbine shown in Eq. 4.

$$L = \sqrt{\sum_{i=1}^{n\text{Turbines}} l_i^2} \quad (4)$$

Finally, the wind speed at any point is expressed in Eq. 5.

$$V = U(1 - L) \quad (5)$$

In this equation, V is the wind speed affected by wake losses, U is the free-stream wind speed, and L is the total wake loss.

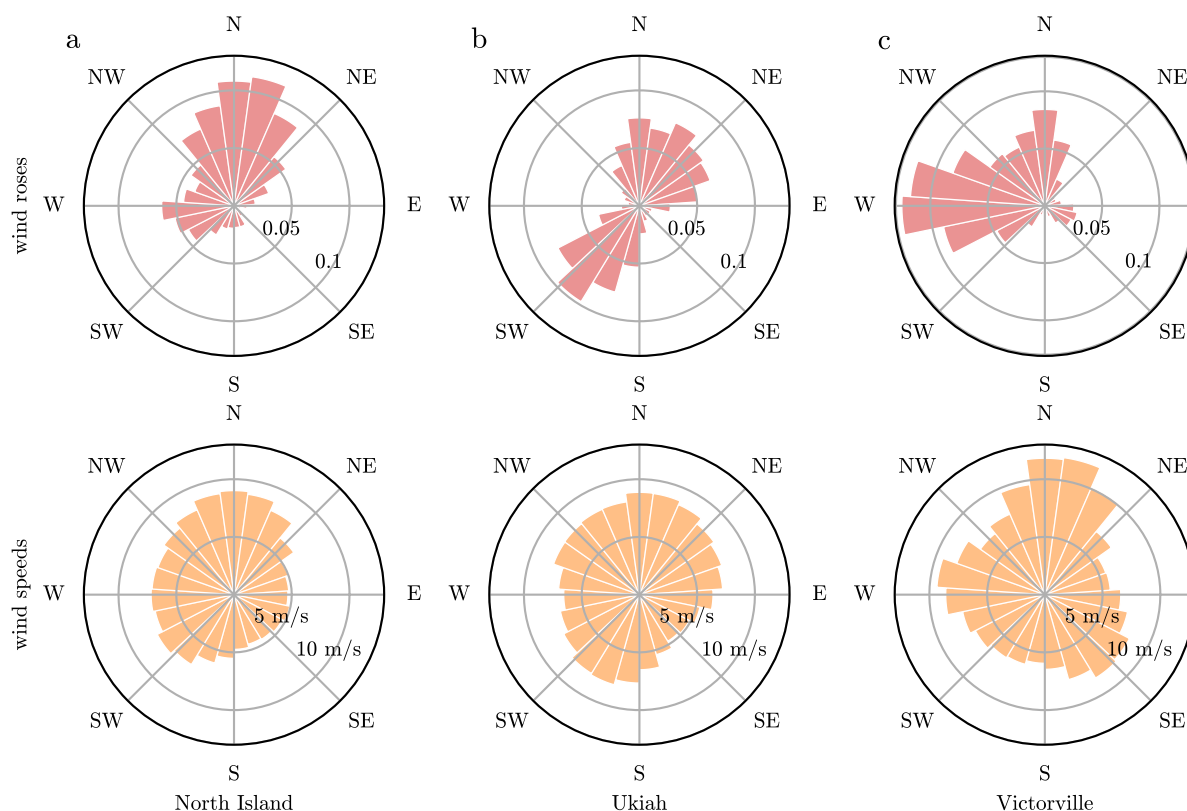



Figure 4. The three wind roses and associated average wind speeds used in this study. The wind resources are from (a) North Island, California, (b) Ukiah, California, and (c) Victorville, California.

125 3.4 Wind Resource

As the goal of this paper is to demonstrate the performance of our layout parameterization method in wind farm optimization for any scenario, we chose three different wind roses from cities in California, USA: North Island, Ukiah, and Victorville³. We divided the wind roses into 23 equal bins for each wind rose, with an associated directionally averaged wind speed, shown in Fig. 4. We have assumed that the wind speed distribution from each wind direction can be approximated with a Weibull distribution defined with the directionally averaged wind speeds (Fig. 5 and Eq. 6). Weibull distributions have been shown to be good representations of real wind speed data (Justus et al., 1978; Rehman et al., 1994; Seguro and Lambert, 2000).

³https://mesonet.agron.iastate.edu/sites/windrose.phtml?station=AAT&network=CA_ASOS



$$f(U, U_{\text{mean}}) = \frac{k}{U_{\text{mean}}} \left(\frac{U}{U_{\text{mean}}} \right)^{k-1} e^{-(U/U_{\text{mean}})^k} \quad (6)$$

Figure 5. Example Weibull distributions for two different average wind speeds. Each wind direction is associated with an average wind speed (shown in Fig. 4), which is used for the value U_{mean} .

In Eq. 6, f is the probability of wind for a given wind speed, U is any wind speed, and U_{mean} is the directionally averaged wind speed for the direction bin of interest. The shape parameter, k , is assumed to be equal to 2.0 for every wind direction, which is a realistic value for the Weibull distributions that represent real wind speed probability data (Rehman et al., 1994; Seguro and Lambert, 2000). For each wind direction, we have sampled the Weibull distribution at five equally spaced points. Five wind speed samples and 23 wind direction samples are chosen as the sampling amount required to converge to the true wind farm production for a given wind farm (Stanley and Ning, 2019).

4 Optimization

140 In this paper we compare how optimizing with BG wind farm layout parameterization compares to two common currently used parameterization methods. We have optimized wind farms using a simple grid parameterization (referred to as “grid optimization”), BG parameterization (“BG”), and by directly optimizing the location of each turbine independently (“direct optimization”). Examples of these layouts, along with the baseline layout that was used to compare results in Sect. 5.1, are shown in Fig. 6.

145 In each case, the objective function of the optimization was to maximize the annual energy production (AEP) of the wind farm, shown in Eq. 7.

$$\text{AEP} = 8760 \sum_{i=1}^{23} \sum_{j=1}^5 P(\phi_i, U(\phi_i)_j) f_i f_j \quad (7)$$

In this equation, 8760 is the number of hours in a year, P is the wind farm power production, ϕ is the wind direction, V is the free-stream wind speed, f_i is the wind direction probability, and f_j is the wind speed probability. The design variables were determined by the optimization method that was used. For the grid optimization, the design variables were the grid spacing in the x and y directions, dx and dy , the grid offset b , and the grid rotation θ for a total of four variables. The discrete variables




Figure 6. Example optimal layouts achieved with each parameterization method. These are 100 turbine layouts, with an average turbine spacing of four rotor diameters and the Princess Amalia wind farm boundary. They were optimized with the wind rose from North Island, California. (a) The baseline grid to which other methods were compared in Sect. 5.1. (b) An example optimized grid layout. (c) An example optimized boundary-grid layout. (d) An example layout that was optimized directly.

in the grid were determined with the same method described above to find the discrete variables in the grid portion of the BG parameterization, except dx and dy were equal while determining the grid format. We experimented with different values of dy during grid initialization and found that setting them equal provided the best results. The design variables for the BG optimization were the same as the grid optimization for the inner grid turbines, and an additional variable s defining the start location of the boundary turbines for a total of five design variables. For the direct optimization methods, the design variables were the x and y locations of each turbine in the wind farm for a total of 200 design variables. In each optimization, we applied turbine spacing constraints and boundary constraints. The turbine hub locations were constrained to not be within two rotor diameters of any other turbine hub. Additionally, the turbine hubs were constrained to be within the defined wind farm



160 boundary. This optimization is expressed in Eq. 8.

$$\begin{aligned} & \text{maximize} && \text{AEP} \\ & \text{w.r.t.} && dx, dy, b, \theta && (\text{grid}) \\ & && dx, dy, b, \theta, s && (\text{BG}) \\ & && x_i, y_i \ (i = 1, \dots, 100) && (\text{direct}) \\ & \text{subject to} && \text{boundary constraints} \\ & && \text{spacing constraints} \end{aligned} \tag{8}$$

We used the optimizer SNOPT, which is a gradient-based optimizer well suited for large-scale nonlinear problems such as the wind farm layout optimization problem (Gill et al., 2005). A challenge of gradient-based optimization is the tendency to converge to local solutions. In order to better search design space, we optimized the problem to convergence 100 times with randomly initialized design variables. This was done for each parameterization method, lending confidence that the best solution after optimizing the 100 random starts is near the global optimum. From the random starting points, we were also able to determine the spread of solutions obtained with each layout parameterization.

We used exact-analytic gradients in each optimization. The gradients for each portion of the model were obtained with an automatic differentiation source code transformation tool, Tapenade (Hascoet and Pascual, 2013). To combine the gradients to get the total derivative of the objective with respect to each of the design variables, we used the open-source optimization framework, OpenMDAO, which propagates the partial derivatives of each small section of the model and calculates the gradients of the entire system (Gray et al., 2010).

Using exact, rather than finite-difference, gradients is important in this study because the computational expense required for optimization problems with increasing design variables scales better with exact gradients (see Fig. 2). For the parameterized optimizations, the exact gradients were not as vital in terms of computational expense, but they were very important for the direct optimizations which had 200 design variables. In addition to reducing the function calls required to reach convergence, the exact gradients helped the optimizer converge to a better solution, avoiding many of the numerical difficulties that often plague the optimization process when using finite-difference gradients.

5 Results and Discussion

180 In this section we demonstrate how the optimal wind farms using BG parameterization compared to wind farms that have been optimized directly, or with a common grid parameterization. We will discuss the best results, the computation expense required to optimize, and the multimodality of the design space with each parameterization method.

5.1 Best Results

Figure 7 shows the best results of the 100 random starts for each parameterization method, compared to a simple baseline grid (Fig. 6a). In Fig. 7, subfigures a, b, and c show results for varied turbine spacing, wind roses, and boundary shapes, respectively.

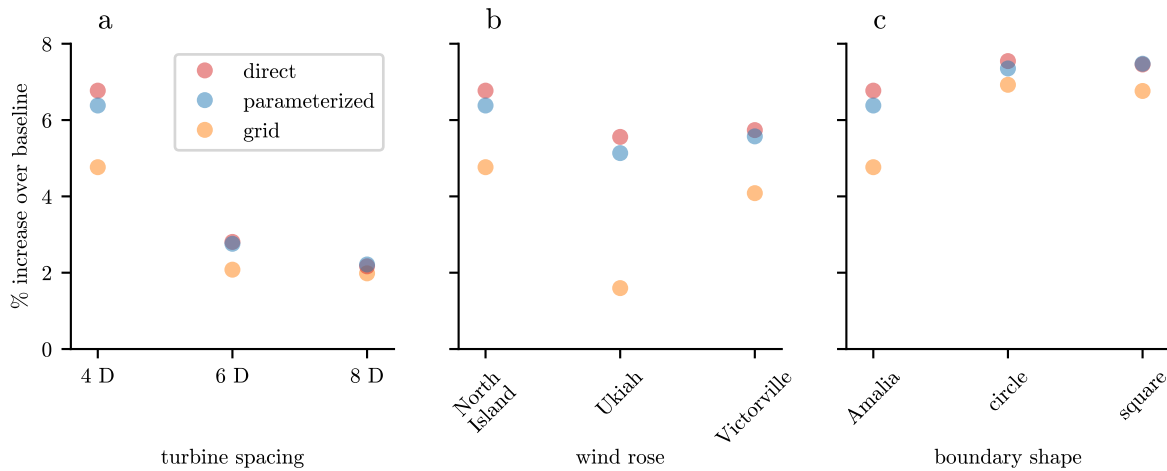


Figure 7. The best annual energy production achieved with 100 randomly initialized optimizations. Shown are the best results from the grid turbine parameterization (four design variables), our new boundary-grid parameterization method (five design variables), and by directly optimizing the location of each turbine (200 design variables). Results are shown as a percent increase over a baseline grid layout. (a) Varied average turbine spacing in the wind farm. (b) Varied wind rose. (c) Varied boundary shape.

Figure 7a shows the optimal results for wind farms with varied average turbine spacing, with the North Island wind rose and Princess Amalia wind farm boundary. For each wind farm, BG layout parameterization performs very well compared to the direct layout optimization results. The worst case, for the smallest wind farm with an average turbine spacing of four rotor diameters, is a mere 0.39% difference between the improvement from the direct optimization and the improvement from our parameterization. For the largest turbine spacing, an average spacing of eight rotor diameters, BG optimization actually performs slightly better than the direct optimization. For the smallest, most tightly packed wind farm, the optimized grid performs significantly better than the baseline, but underperforms by about 2% compared to the other parameterization methods. Even at an average turbine spacing of six rotor diameters, the direct and parameterized optimizations perform about 0.7% better than the grid optimization, which may or may not be significant depending on the uncertainty of the models used.

190 For the largest wind farm, the optimal grid performs almost identically to the other parameterization methods. For large wind farms where the turbines are spaced very far apart, wakes are mostly recovered by the time they reach other turbines in the wind farm. In these cases, even an optimized grid performs almost as well as the direct optimization.

Figure 7b shows results for optimized wind farms with different wind resources, with an average turbine spacing of four rotor diameters and the Princess Amalia wind farm boundary. The wind roses and the associated directionally averaged wind speeds are shown in Fig. 4. As with the varied turbine spacing results, the BG results are almost identical to the direct optimizations.

200 The largest difference occurred for the Ukiah wind rose, for which direct optimization performed 0.43% better than the BG



optimization. Additionally, for all wind resources, the BG optimizations perform significantly better than the simple grid optimizations. For the North Island and Victorville roses, the BG layout optimization performs about 1.5% better than the grid parameterization, where for the Ukiah rose the parameterized is 3.5% better.

205 Figure 7c shows the results for a varied wind farm boundary. The farms in this subfigure have an average turbine spacing of four rotor diameters and the North Island wind rose. Consistent with the previous results, the parameterized optimization performs superbly. The worst case difference is 0.39% compared to the direct optimization, which occurs for the most complex Amalia wind farm boundary. The parameterized optimization performs very slightly, 0.02%, better than the direct optimization for the square wind farm boundary. For the circle and square boundaries, the optimized grid parameterization performs very
210 well compared to the parameterized and direct optimizations. The grid optimization for the circle boundary is only 0.43% worse than the parameterized optimization and only 0.62% worse than the direct optimization. At least for the North Island wind rose, the grid parameterization may be sufficient for the simple, symmetric boundaries.

In terms of the best achievable wind farms with each parameterization method, our new BG method performs almost identically to optimizing the location of each wind turbine directly. Of all the cases we tested, the largest difference between the
215 BG and direct results was 0.43%, and in some cases the BG results were slightly better than the direct optimization. With only five design variables, we can create wind farms that perform the same as farms that have been designed with 200 variables. While the grid parameterization is able to achieve good results for some wind farms, it often performs much worse than our parameterization. One additional variable is a small price to pay for significant improvement in optimal wind farm design.

5.2 Computational Expense

220 The utility of any wind farm layout parameterization is not only measured by the ability to create high energy producing wind farms, but by the ability to do so quickly and reliably. Figures 8, 9, and 10 are histograms showing optimal results and the computational expense required for each of the 100 optimizations run for each wind farm and parameterization method. In each figure, Subfigures a-c show the normalized optimal AEP for each of the 100 runs, and Subfigures d-f show the number of function required to converge to a solution. The AEP results have each been normalized by the maximum AEP achieved by
225 the direct optimizations for the associated wind farm. Also note that the number of function calls are shown with a log scale.

In general, the grid and the BG optimal AEP results have a similar spread, with the BG results shifted up higher. Compared to the direct optimizations, the grid and BG optimizations have a larger spread in optimal solutions. This is a consequence of the discrete variables that are initialized at the start of each optimization run. The number of rows and columns, as well as their organization in the grid are determined by the randomly initialized rotation design variable, θ . Some of these grid formations
230 are more desirable than others, leading to higher AEP values. This spread in optimal solutions is not a significant issue, because the number of functions calls required for the grid and BG optimizations are an order of magnitude lower than that required by the direct optimization. This allows for many randomly initiated runs in a short amount of time. If it did become an issue, the spread could be reduced by predefining the discrete grid variables, or including them as design variables in a gradient-free formulation.



235 With regards to the function calls required to converge, the grid optimizations required about one third of the function calls
to converge compared to the BG optimizations, while the direct optimizations required about an order of magnitude more.
Function calls are an important measure of computational expense, as they are correlated with time and processing power
required to optimize. Here it is important to remember that our results were obtained with exact-analytic gradients, meaning
that one function call was required to obtain the wind farm AEP, as well as the gradients with respect to each of the design
240 variables. The same is true of the constraints, one function call gave both the constraint values and the gradients. Without
exact gradients, a finite-difference method would need to be used to calculate the gradients. At every optimization step, finite-
difference gradients require one (forward or backward difference) or two (central difference) additional function calls for
every design variable to approximate the gradients. Thus, if forward-difference gradients were used rather than exact, the grid
optimizations would need about four times as many function calls to reach a solution, the BG optimization would need about
245 five times as many function calls, and the direct optimization would need 200 times as many function calls to converge. This is
the best case scenario, as optimizations with finite-difference gradients often have trouble converging. Compared to gradient-
free optimization, the exact analytic gradients are vital. The direct optimization with a gradient-free technique would be near
impossible because of the massive required computational expense (Ning and Petch, 2016; Thomas and Ning, 2018).

5.3 Multimodality

250 One of the major difficulties of the wind farm layout optimization problem is the extreme multimodality of the design space
(Fig. 1). There can be thousands or even millions of local solutions, often varying drastically in their quality. Figure 11 shows
one dimensional sweeps across the design variables, for each of the three different parameterization methods discussed in this
paper. Because of the number of variables in this problem, it is difficult to fully represent the full design space graphically,
however this figure is a good indicator of the multimodality of the different design spaces. Figures 11a, 11b, and 11c show the
255 multimodality of the grid, BG, and direct layout parameterizations, respectively.

Parameterizing the design space with a grid and with the BG method (Figs. 11a and 11b) does not completely remove
the multimodality of the wind farm layout problem. However, it does result in a smoother response and fewer local minima
compared to the design space when each of the turbines are optimized directly. These function spaces can be explored easily
with a few random starting locations, or with a gradient-free optimization method. The design space when varying the location
260 of individual turbines (Figs. 1 and 11c) is much more noisy, filled with comparatively larger peaks and valleys in the design
space. These figures only show the design space with respect to the location of one turbine, which is defined with two variables.
The full space consists of the location of all 100 turbines, or 200 variables, for which the multimodality and overall noisiness of
the design space is exacerbated. Figures 11a and 11b do not show the function space with respect to the discrete grid variables.
Even so, considering each combination of the feasible grid variables is more desirable than the difficulty involved with the
265 200-dimensional function space of the direct layout definition.

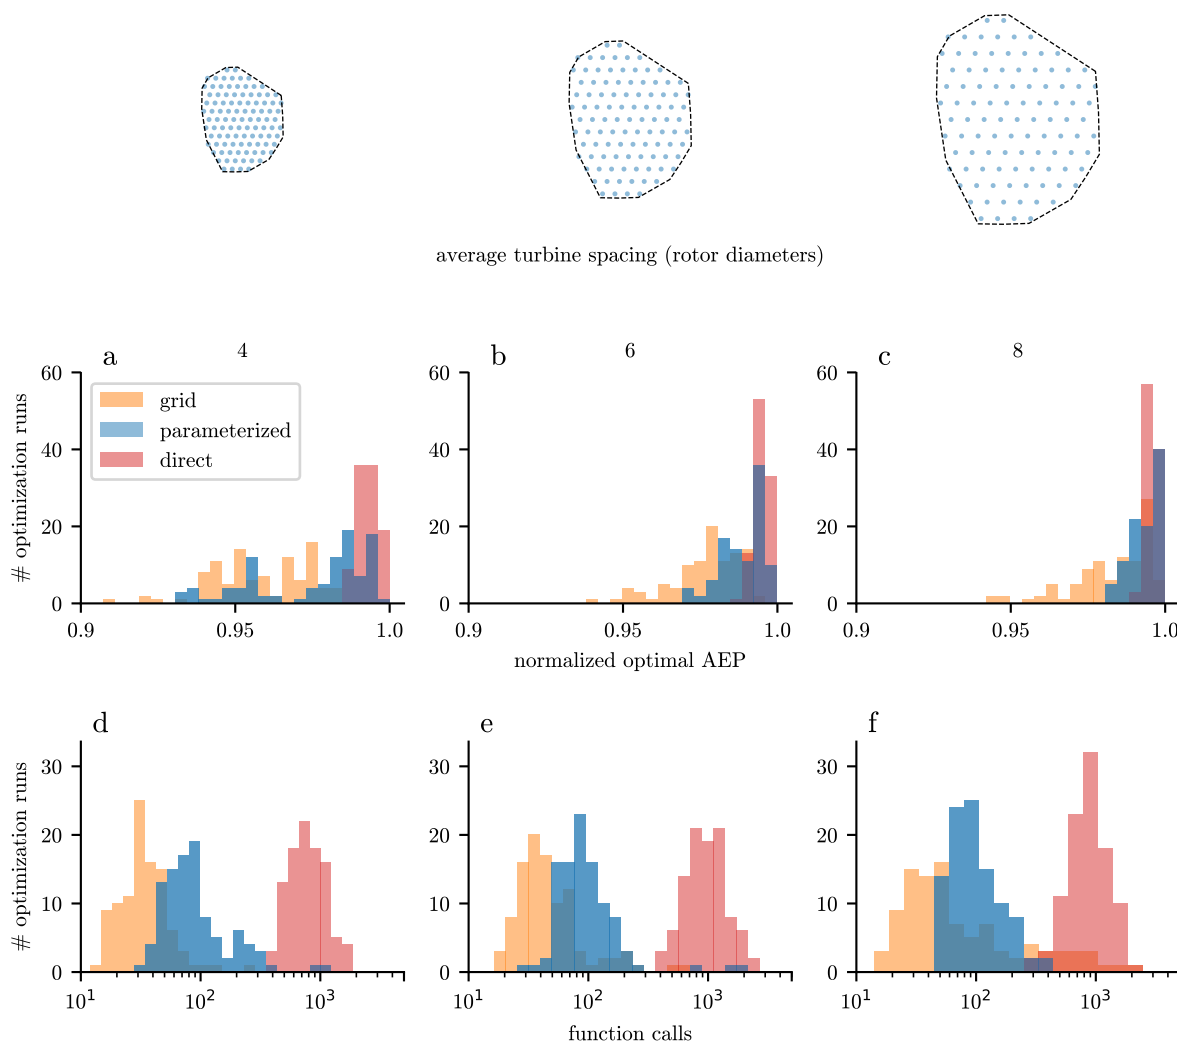


Figure 8. Results from 100 randomly initialized optimizations for wind farms with varied average turbine spacing and 100 wind turbines. The farm optimized had the Princess Amalia boundary, and the wind rose from North Island, California. Shown are results using the grid turbine parameterization, our new boundary-grid parameterization, and by direct optimization. The optimal annual energy production distribution achieved for each of the optimization runs, in wind farms with varied turbine spacing of 4, 6, and 8 rotor diameters for subfigures (a), (b), and (c), respectively. The number of function calls required to converge for each of the optimization runs, in wind farms with varied turbine spacing of 4, 6, and 8 rotor diameters for subfigures (d), (e), and (f), respectively.

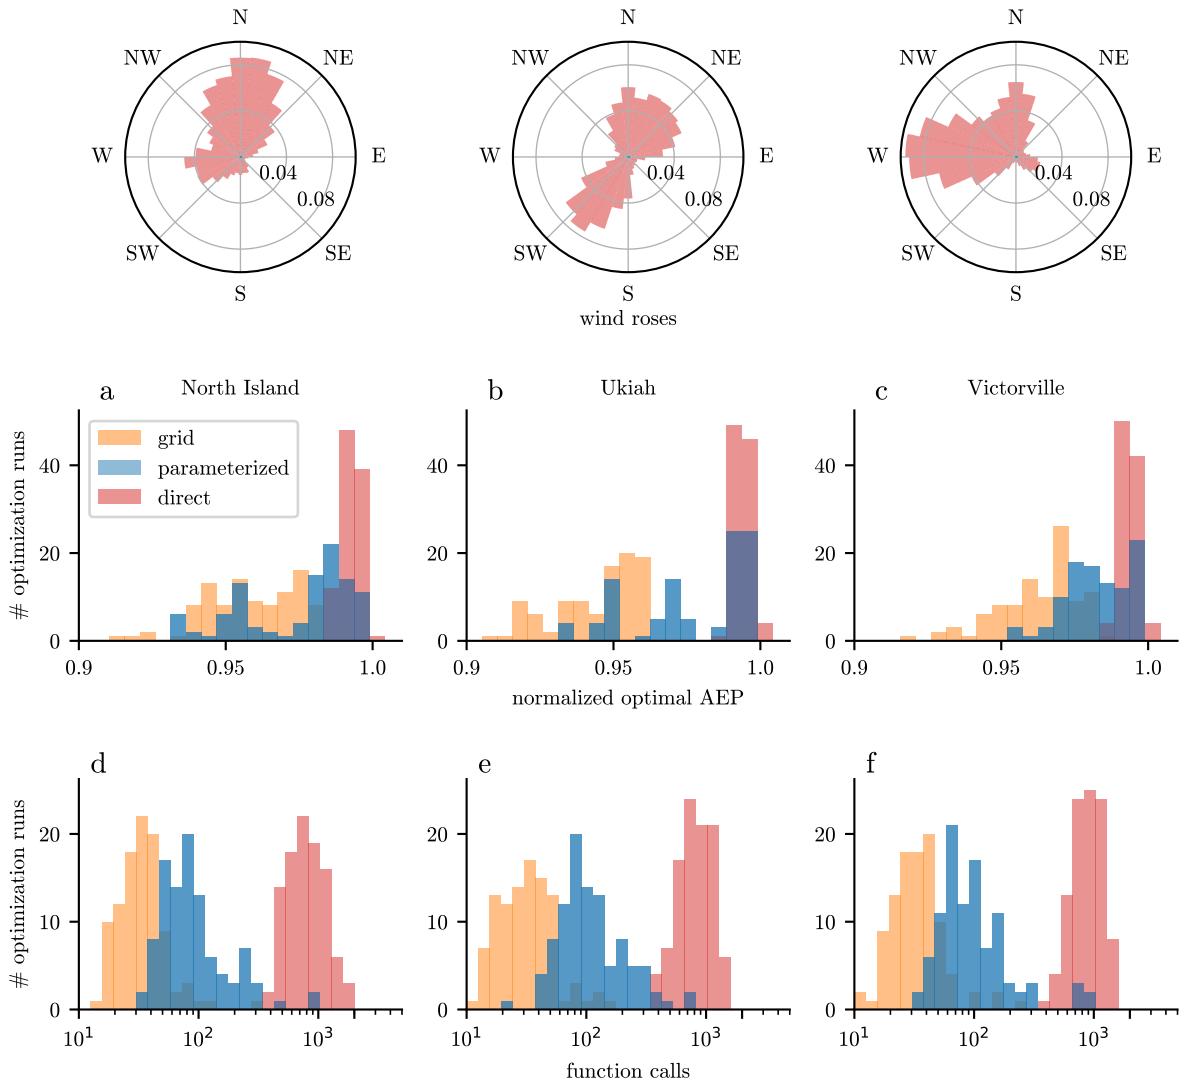


Figure 9. Results from 100 randomly initialize optimizations for wind farms with varied wind roses and 100 wind turbines. The farm optimized had the Princess Amalia boundary, and the average turbine spacing was four rotor diameters. Shown are results using the grid turbine parameterization, our new boundary-grid parameterization, and by direct optimization. The optimal annual energy production distribution achieved for each of the optimization runs, in wind farms with varied wind roses. Wind rose from (a) North Island, California, (b) Ukiah, California and (c) Victorville, California. The number of function calls required to converge for each of the optimization runs, in wind farms with varied wind roses. Wind Rose from (d) North Island, California, (e) Ukiah, California and (f) Victorville, California.

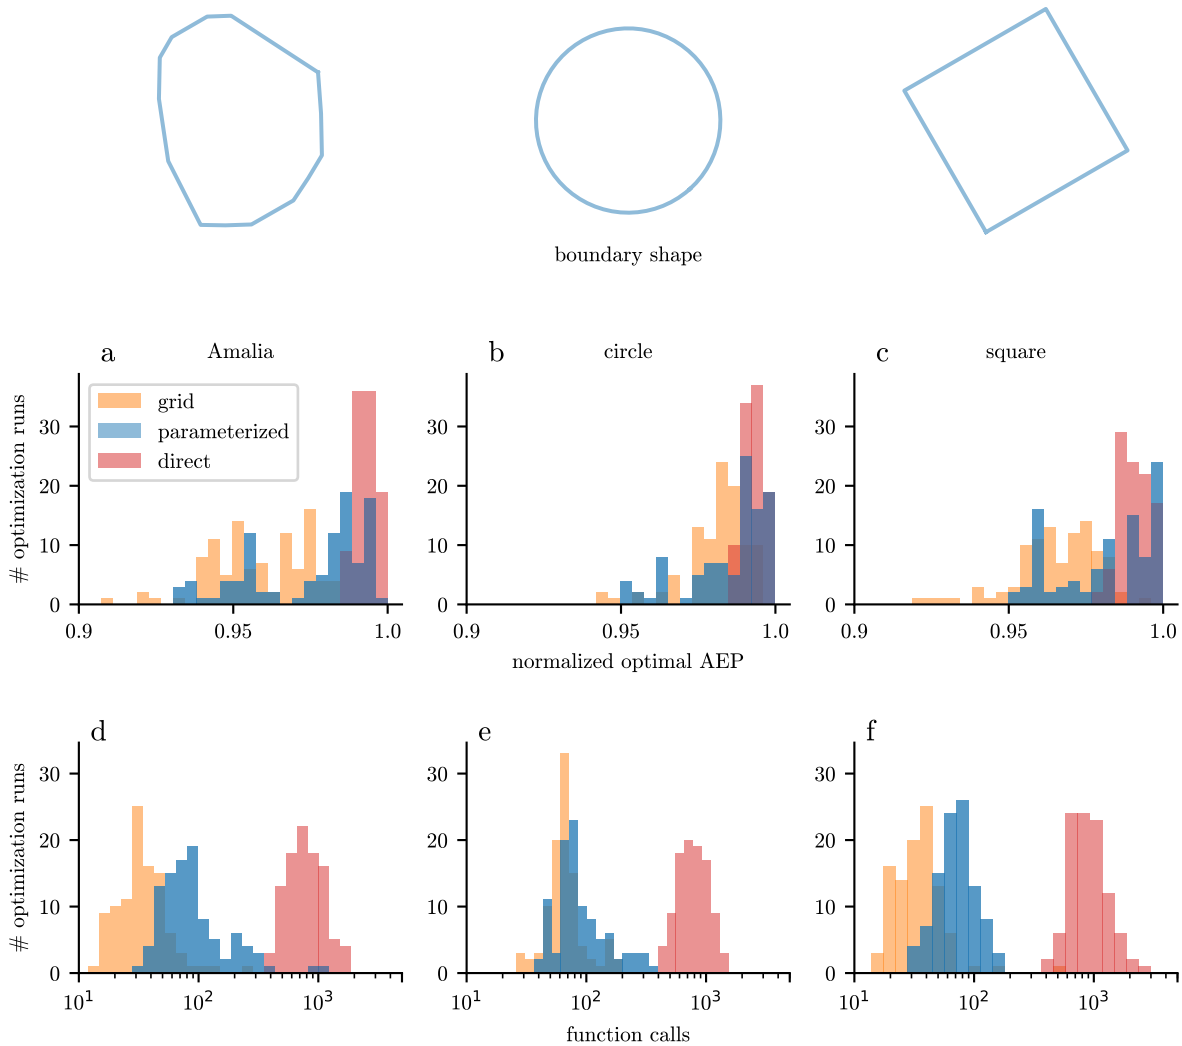


Figure 10. Results from 100 randomly initialize optimizations for wind farms with varied wind farm boundaries and 100 wind turbines. The average turbine spacing was four rotor diameters, and the wind rose was from North Island, California. Shown are results using the grid turbine parameterization, our new boundary-grid parameterization, and by direct optimization. The optimal annual energy production distribution achieved for each of the optimization runs, in wind farms with varied boundary shapes. (a) Princess Amalia wind farm boundary. (b) Circular wind farm. (c) Square wind farm. The number of function calls required to converge for each of the optimization runs, in wind farms with varied boundary shapes. (d) Princess Amalia wind farm boundary. (e) Circular wind farm. (f) Square wind farm.

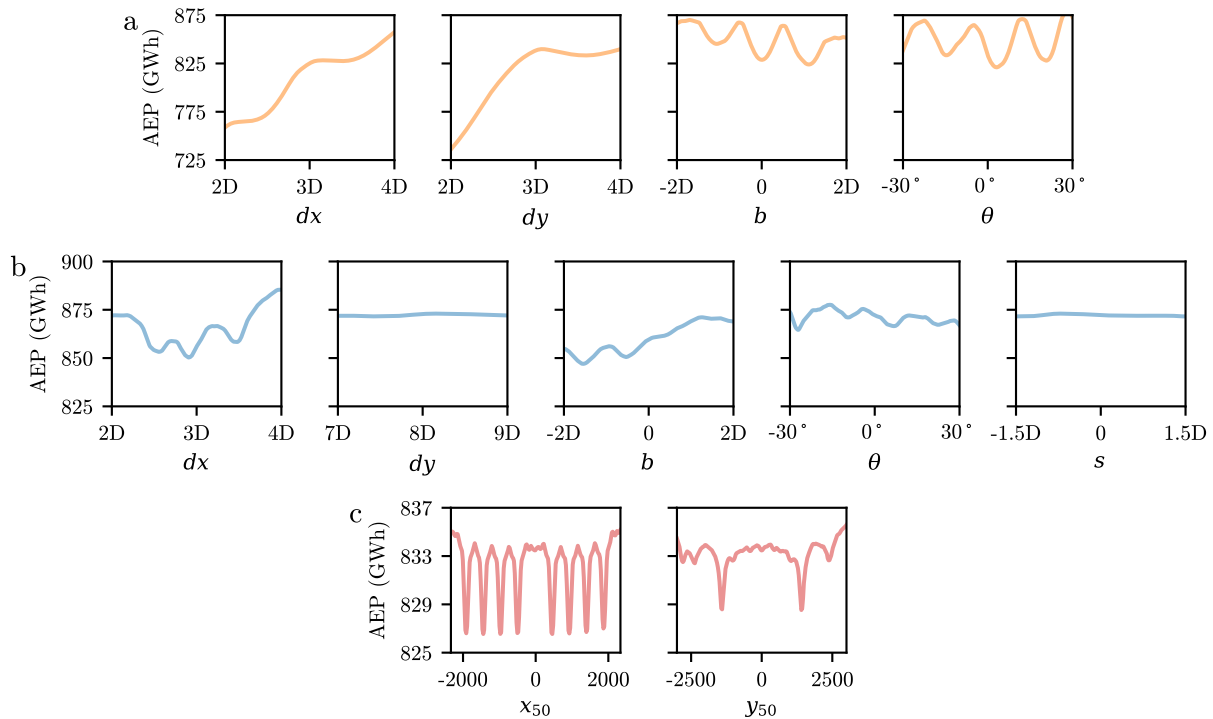


Figure 11. One dimensional sweeps across the design space of each parameterization method discussed in this paper. These figures show the multimodality of each of the design spaces. (a) The simple grid parameterization. (b) Our newly presented boundary-grid parameterization. (c) Moving the location of one wind turbine across the wind farm in x and y (refer to Fig. 1). With the direct turbine layout definition there are actually 200 variables. This figure shows the multimodality in just 2 of these variables, where the whole design space is much more complex.

6 Additional Benefits to BG Parameterization

BG parameterization requires few variables, produces wind farm layouts that perform similarly to ones that have been optimized directly with much lower computational expense, and reduces the multimodality of the design space. In addition, there are some innate design characteristics that are useful in wind farm design. First, the layouts produced are regular, aesthetically pleasing patterns. To the untrained eye, BG parameterization looks well designed compared to the seemingly random layouts that are often produced when every turbine location is optimized individually. This can play an important role in the public perception of large scale wind energy. Second, BG parameterization has clear roads or shipping lanes naturally built into the design. Roads and shipping lanes are requirements in wind farm design that are often neglected in research studies. Finally, the layouts produced have a clear cabling pattern. Wind farm cabling is an expensive and complex part of wind farm design. With



275 BG parameterization, cabling requirements can be clearly minimized by running cables across each of the rows, and around the boundary without the need for complex cabling algorithms.

7 Conclusions

In this paper, we have presented the new boundary-grid wind farm layout parameterization method. This method uses only five design variables, regardless of the number of wind turbines, but is capable of producing turbine layouts that perform just as well to layouts where the location of each wind turbine has been optimized directly. Of all the cases that we tested, the largest difference in AEP improvement over a common baseline between BG layout optimization and a direct layout optimization was 0.43%, and in a few cases BG parameterization even performed better. BG parameterization also performed better than a simple grid parameterization in all cases tested.

In addition to being able to match the optimal energy production of wind farms that were directly optimized, BG parameterization requires an order of magnitude fewer function calls to reach a solution. This is with exact-analytic gradients, which means if finite-difference gradients or a gradient-free optimization method were used instead, our parameterization method would require at least two to three orders of magnitude fewer function calls to optimize. BG parameterization also reduces the multimodality of the design space, simplifying the optimization process and making it easier to find a good solution.

The BG layout definition places a portion of the wind turbines around the boundary, spaced equally traversing the wind farm perimeter. The rest of the turbines are placed in a grid inside the farm boundaries. The wind farm layouts created have a regular, aesthetically pleasing pattern, naturally defined roads and shipping lanes, and an easily defined cabling pattern. BG parameterizations solves many of the problems that typically accompany wind farm layout optimization. It is a simple, easily implemented technique that can immediately be applied by researchers and wind farm developers, playing an important role in continued growth of wind energy.

295 Code and data availability

The code written for this paper is included here: DOI: 10.5281/zenodo.3261037. All dependencies, with the exception of the optimizer SNOPT are open source.

Author Contributions

APJS performed every part of this research and is the primary author. AN helped develop ideas and direction, provided feedback, and provided editing for the paper.

Competing Interests

The authors declare no competing interests.



References

- Baker, N. F., Stanley, A. P. J., Thomas, J. J., Ning, A., and Dykes, K.: Best Practices for Wake Model and Optimization Algorithm Selection
305 in Wind Farm Layout Optimization, in: AIAA Scitech 2019 Forum, San Diego, CA, <https://doi.org/10.2514/6.2019-0540>, 2019.
- Barthelmie, R. J., Frandsen, S. T., Nielsen, M., Pryor, S., Rethore, P.-E., and Jørgensen, H. E.: Modelling and measurements of power losses and turbulence intensity in wind turbine wakes at Middelgrunden offshore wind farm, *Wind Energy*, 10, 517–528, 2007.
- Barthelmie, R. J., Hansen, K., Frandsen, S. T., Rathmann, O., Schepers, J., Schlez, W., Phillips, J., Rados, K., Zervos, A., Politis, E., et al.: Modelling and measuring flow and wind turbine wakes in large wind farms offshore, *Wind Energy*, 12, 431–444, 2009.
- 310 Bastankhah, M. and Porté-Agel, F.: Experimental and theoretical study of wind turbine wakes in yawed conditions, *Journal of Fluid Mechanics*, 806, 506–541, 2016.
- Briggs, K.: Navigating the complexities of wake losses, *Norik American Windpower*, 10, 2013.
- Emami, A. and Nogreh, P.: New approach on optimization in placement of wind turbines within wind farm by genetic algorithms, *Renewable Energy*, 35, 1559–1564, 2010.
- 315 Feng, J. and Shen, W. Z.: Solving the wind farm layout optimization problem using random search algorithm, *Renewable Energy*, 78, 182–192, 2015.
- Gebraad, P. and Van Wingerden, J.: Maximum power-point tracking control for wind farms, *Wind Energy*, 18, 429–447, 2015.
- 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.
- 320 Gill, P. E., Murray, W., and Saunders, M. A.: SNOPT: An SQP algorithm for large-scale constrained optimization, *SIAM review*, 47, 99–131, 2005.
- González, J. S., García, Á. L. T., Payán, M. B., Santos, J. R., and Rodríguez, Á. G. G.: Optimal wind-turbine micro-siting of offshore wind farms: A grid-like layout approach, *Applied energy*, 200, 28–38, 2017.
- Gray, J., Moore, K., and Naylor, B.: OpenMDAO: An open source framework for multidisciplinary analysis and optimization, in: 13th
325 AIAA/ISSMO Multidisciplinary Analysis Optimization Conference, p. 9101, 2010.
- Guirguis, D., Romero, D. A., and Amon, C. H.: Toward efficient optimization of wind farm layouts: Utilizing exact gradient information, *Applied energy*, 179, 110–123, 2016.
- Hascoet, L. and Pascual, V.: The Tapenade automatic differentiation tool: Principles, model, and specification, *ACM Transactions on Mathematical Software (TOMS)*, 39, 20, 2013.
- 330 Herbert-Acero, J., Probst, O., Réthoré, P.-E., Larsen, G., and Castillo-Villar, K.: A review of methodological approaches for the design and optimization of wind farms, *Energies*, 7, 6930–7016, 2014.
- Ju, X. and Liu, F.: Wind farm layout optimization using self-informed genetic algorithm with information guided exploitation, *Applied Energy*, 248, 429–445, 2019.
- Justus, C., Hargraves, W., Mikhail, A., and Graber, D.: Methods for estimating wind speed frequency distributions, *Journal of applied meteorology*, 17, 350–353, 1978.
- 335 Kanev, S., Savenije, F., and Engels, W.: Active wake control: An approach to optimize the lifetime operation of wind farms, *Wind Energy*, 21, 488–501, 2018.
- Ning, A. and Petch, D.: Integrated Design of Downwind Land-Based Wind Turbines Using Analytic Gradients, *Wind Energy*, 19, 2137–2152, <https://doi.org/10.1002/we.1972>, 2016.



- 340 Parada, L., Herrera, C., Flores, P., and Parada, V.: Wind farm layout optimization using a Gaussian-based wake model, *Renewable energy*, 107, 531–541, 2017.
- Perez-Moreno, S. S., Dykes, K., Merz, K. O., and Zaaier, M. B.: Multidisciplinary design analysis and optimisation of a reference offshore wind plant, in: *Journal of Physics: Conference Series*, vol. 1037, p. 042004, IOP Publishing, 2018.
- Rehman, S., Halawani, T., and Husain, T.: Weibull parameters for wind speed distribution in Saudi Arabia, *Solar Energy*, 53, 473–479, 1994.
- 345 Seguro, J. and Lambert, T.: Modern estimation of the parameters of the Weibull wind speed distribution for wind energy analysis, *Journal of wind engineering and industrial aerodynamics*, 85, 75–84, 2000.
- Stanley, A. P. and Ning, A.: Coupled Wind Turbine Design and Layout Optimization with Non-Homogeneous Wind Turbines, *Wind Energy Science*, 4, 99–114, <https://doi.org/10.5194/wes-4-99-2019>, 2019.
- Stanley, A. P. J., Ning, A., and Dykes, K.: Optimization of Turbine Design in Wind Farms with Multiple Hub Heights, Using Exact Analytic
- 350 Gradients and Structural Constraints, *Wind Energy*, 22, 605–619, <https://doi.org/10.1002/we.2310>, 2019.
- Thomas, J. J. and Ning, A.: A Method for Reducing Multi-Modality in the Wind Farm Layout Optimization Problem, in: *Journal of Physics: Conference Series*, vol. 1037, *The Science of Making Torque from Wind*, Milano, Italy, <https://doi.org/10.1088/1742-6596/1037/4/042012>, 2018.
- U.S. Energy Information Administration: Short-term energy outlook, Department of Energy, pp. 1–49, 2019a.
- 355 U.S. Energy Information Administration: Annual energy outlook 2019 with projections to 2050, Department of Energy, pp. 1–83, 2019b.
- Van Dam, F., Gebraad, P., and van Wingerden, J.-W.: A maximum power point tracking approach for wind farm control, *Proceedings of The Science of Making Torque from Wind*, 2012.