



Objective and Algorithm Considerations When Optimizing the Number and Placement of Turbines in a Wind Power Plant

Andrew P. J. Stanley, Owen Roberts, Jennifer King, and Christopher J. Bay

National Renewable Energy Laboratory, National Wind Technology Center, Boulder, CO 80303 USA

Correspondence: Andrew P. J. Stanley (PJ.Stanley@nrel.gov)

Abstract. Optimizing turbine layout is a challenging problem that has been extensively researched in literature. However, optimizing the number of turbines within a given boundary has not been studied as extensively and is a difficult problem because it introduces discrete design variables and a discontinuous design space. An essential step in performing wind power plant layout optimization is to define the objective function, or value, that is used to express what is valuable to a wind power plant developer, such as annual energy production, cost of energy, or profit. In this paper, we demonstrate the importance of selecting the appropriate objective function when optimizing a wind power plant. We optimize several different wind power plants with different wind resources and boundary sizes. Results show that the optimal number of turbines varies drastically depending on the objective function. For a simple, one-dimensional, land-based scenario, we found that a wind power plant optimized for minimal cost of energy produced just 72% of the profit as the wind power plant optimized for maximum profit, which corresponded to a loss of about \$2 million each year. This paper also compares the performance of several different optimization algorithms, including a novel repeated-sweep algorithm that we developed. We found that the performance of each algorithm depended on the number of design variables in the problem as well as the objective function.

1 Introduction

Wind energy provides several advantages to the sustainable energy grid of the future. While producing power, wind turbines produce zero carbon dioxide or other air pollution, require no external fuel, and require little to no water. Additionally, wind plants have an energy payback time of less than a year and can produce energy in an economically efficient manner (Razdan and Garrett, 2017; Vestas, 2020). In fact, wind energy has been a central focus of research and development in past decades such that, currently, wind is one of the cheapest sources of energy available (Lazard, 2018). Because of its many benefits, but in large part due to the economic drivers, wind energy installations have grown throughout the world as has the relative share of energy produced by wind. In coming years, wind energy technology is projected to improve, and wind energy capacity and market penetration are to increase even further (U.S. Energy Information Administration, 2019).

~~Because of economies of scale,~~ utility-scale wind turbines are deployed in groups to minimize costs. This provides reduced cabling costs, easier construction and maintenance, and reduced land requirements. However, building turbines close together also introduces some challenges. One of these challenges is wake interaction between turbines. A wind turbine removes kinetic energy from the air around it and converts this energy to electricity, creating a wake of slow-moving and turbulent wind



behind it. When turbines are built close together, their wakes can reduce the amount of energy available in the wind, causing downstream turbines to produce less energy as a result. One way to reduce wake interactions between turbines is through wind power plant layout optimization. This optimization process involves determining the number of turbines to build in a wind power plant and their locations in order to reduce wake interactions and maximize performance. Finding an optimal wind power plant layout is a challenging, nonconvex problem with many interacting design variables. It is difficult to solve this problem without mathematical optimization tools because it often requires not-so-obvious trade-offs to reach a solution.

Appropriate methods of determining turbine locations within a wind power plant have been intensively studied, and researchers have demonstrated several methods that can be used to effectively optimize a wind power plant layout. The literature demonstrates a preference for gradient-free optimization methods applied to wind power plant layout optimization, and different studies showed success using genetic algorithms, particle swarm methods, and random search to determine improved wind power plant layouts (Hou et al., 2019). A common layout optimization method is to divide the wind power plant domain into a grid that defines possible turbine locations (at the center of the grid cells or at the intersections of the lines). One of the previously mentioned optimization methods is then used to determine at which of the predefined locations a turbine should be placed. In more recent years, some studies also showed success optimizing wind power plant layouts with gradient-based methods. This type of optimization requires a continuous design space and computationally or analytically provided gradients that increase the complexity of the problem formulation. However, the computational expense required for gradient-based optimization scales favorably with increasing numbers of design variables compared to gradient-free methods for which the computational expense scales very poorly.

Layout optimization studies are almost always performed assuming that the number of turbines in the wind power plant is previously known. Determining the optimal number of turbines in a wind power plant is a much more difficult problem to solve because it requires the optimization of at least one integer design variable or a discontinuous design space. Although it has not been discussed in the literature as much as layout optimization determining turbine placement, optimizing the number of turbines in a wind power plant has also been addressed in previous studies. Mosetti et al. (1994) first addressed this issue when they divided a wind power plant domain into 100 square cells as candidate turbine locations then used a genetic algorithm to determine the optimal number of turbines and at which of the potential locations they should be placed. Since this seminal paper was published, many other researchers proposed improvements to Mosetti's methodology and were able to find improved results, mostly by using new and better optimizers (Grady et al., 2005; Zergane et al., 2018; Ituarte-Villarreal and Espiritu, 2011; Moorthy and Deshmukh, 2015). Additionally, some applied a similar methodology to optimizing turbine number and layout at real geographical locations (Şişbot et al., 2010; Khanali et al., 2018). The vast majority of these more recent studies kept the same general optimization strategy, performed by dividing the wind power plant domain into a grid that defines potential turbine locations and using some optimizer to determine the best layouts.

Selecting the appropriate optimization methodology is a vital part of the wind power plant layout optimization process because it will determine the quality of the final solution as well as the required computational expense. In addition to the optimization algorithm, a critical step is to appropriately select the objective function. For wind power plant layout optimization studies, objectives that are often considered are annual energy production (AEP) or cost of energy (COE). With a fixed number



of turbines, the objective may not have much of an effect on the final solution. However, when the number of turbines is also being optimized, the objective function can have a profound effect on the final optimized layout. For one set of optimizations discussed in Sec. 6.4, the optimal number of turbines ranges from 15–54, and the annual costs range from \$6.75 million to \$21.96 million, depending on if the plant was optimized for AEP, COE, or profit.

65 In this paper, we studied two specific considerations in optimizing the number of turbines and their layout in a wind power plant. First, we studied how different objective functions alter the optimized number of wind turbines and their layout in a wind power plant. As far as we are aware, this is the first paper that starts to explore the effect of the objective on the final wind power plant design. The objective can greatly affect the final layout and performance of an optimized wind power plant. Second, we compared using different problem formulations and optimization algorithms in finding a solution. Previous papers have begun
 70 to research this topic but, in this paper, we specifically study how different algorithms perform depending on the objective and the size of the optimization problem. We compare a genetic algorithm and a greedy algorithm in a gridded wind power plant domain, two commonly used wind power plant optimization methods, as well as a genetic algorithm with the boundary-grid method and a new repeated-sweep algorithm in a gridded domain.

2 Wake Model

75 The wind speed downstream of a turbine is reduced because turbines extract energy from the flow and from the complex physics of the wakes they produce. In this paper, the desirability of the wind power plants we examined were dependent, to a large extent, on energy production. This energy production is a function of the wind speeds throughout the wind power plant. To calculate the wind speeds to be used in turbine power calculations, we used an analytic Gaussian wake model (Bastankhah and Porté-Agel, 2016; Abkar and Porté-Agel, 2015; Niayifar and Porté-Agel, 2016). The wake calculations were performed
 80 using FLOW Redirection and Induction in Steady State (FLORIS), which is a computationally inexpensive, controls-oriented tool to calculate the steady-state flow field in a wind power plant (NREL, 2020). We include a brief description of the Gaussian wake model in this paper but, for more details, refer to the original model paper (Bastankhah and Porté-Agel, 2016).

Using the Gaussian wake model, the velocity of the wake behind a turbine is computed with the following analytical expressions:

$$85 \quad \frac{u(x, y, z)}{U_\infty} = 1 - C e^{-(y-\delta)^2 / 2\sigma_y^2 - (z-z_h)^2 / 2\sigma_z^2} \quad (1)$$

$$C = 1 - \sqrt{1 - \frac{(\sigma_{y0}\sigma_{z0})C_T}{\sigma_y\sigma_z}} \quad (2)$$

where u is the velocity at a desired location (x, y, z) , C is the velocity deficit at the wake center, U_∞ is the freestream velocity, δ is the wake deflection, z_h is the hub height of the turbine, σ_y defines the wake width in the y direction, and σ_z defines the wake width in the z direction. Each of these parameters is defined with respect to each turbine. The subscript “0” refers to the initial values at the start of the far wake, which is dependent on ambient turbulence intensity, I_0 , and the thrust coefficient, C_T .



For additional details on near-wake calculations, refer to the original paper describing this model (Bastankhah and Porté-Agel, 2016). The velocity distributions σ_z and σ_y are defined as:

$$\frac{\sigma_z}{D} = k_z \frac{(x - x_0)}{D} + \frac{\sigma_{z0}}{D} \quad \text{where} \quad \frac{\sigma_{z0}}{D} = \frac{1}{2} \sqrt{\frac{u_R}{U_\infty + u_0}} \quad (3)$$

95

$$\frac{\sigma_y}{D} = k_y \frac{(x - x_0)}{D} + \frac{\sigma_{y0}}{D} \quad \text{where} \quad \frac{\sigma_{y0}}{D} = \frac{\sigma_{z0}}{D} \cos \gamma \quad (4)$$

where D is the rotor diameter, u_R is the velocity at the rotor, u_0 is the velocity behind the rotor, γ is the rotor yaw angle (which is assumed to be 0 in this paper), k_y defines the wake expansion in the lateral direction, and k_z defines the wake expansion in the vertical direction. For this study, k_y and k_z are set to be equal and the wake expands at the same rate in the lateral and vertical directions. The wake deficits were combined using the traditional sum of squares method (Katić et al., 1986). Equation 1 defines the wind speed, u , at a single desired point. To determine the average rotor wind speed used to calculate the power production of a wind turbine, we averaged the wind speeds sampled at nine locations across the swept rotor area, which is the default in FLORIS.

100

For this study, we used a 2.5-MW turbine definition. The turbine parameters are shown in Table 1, and the power and thrust coefficient curves, as well as the power curve, are shown in Fig. 1. As seen in the power curve, the rated wind speed is near

Table 1. Wind turbine parameters.

rated power	2.5 MW
rotor diameter	117.8 m
hub height	88 m




Figure 1. Left: C_P and C_T curves for the 2.5-MW turbine used in this study. Right: the power curve for this same turbine.

105

10 m/s.



3 Objective Functions

In this paper, we explored three different objective functions in our wind power plant optimizations: 1) AEP, 2) COE, and 3) annual profit. In this section, each objective function will be described in detail. We acknowledge that the models we used in this paper are simplistic. These simplified models are sufficient for this demonstration and investigation into varying results from different objectives; however, more detailed models can be easily included, depending on the use case and data available.

3.1 Annual Energy Production

AEP is a standard objective in wind power plant optimization (Pérez et al., 2013; Gebraad et al., 2017; Thomas and Ning, 2018). For problems where the value of energy produced by the wind power plant is fixed throughout its lifetime and independent of the time of day, and where the project cost remains constant or is not an important consideration, AEP is a reasonable objective. AEP optimization simply aims to maximize the energy production for some fixed cost. For example, AEP is a common objective for wind power plant layout optimization where the turbine number and design are fixed. Typically to calculate AEP, the wind directions and wind speeds are grouped into discrete bins in order to numerically calculate the integral:

$$AEP = 8,760 \sum_{i=1}^{n_d} \sum_{j=1}^{n_s} P_f(\phi_i, U(\phi_i)_j) f_i f_j \quad (5)$$

where 8,760 is the number of hours in a year, n_d is the number of wind direction bins, n_s is the number of wind speed bins per wind direction, P_f is the power production of the wind power plant, ϕ_i is the wind direction, U is the wind speed, and f_i and f_j are the frequency of wind associated with a given direction and speed.

The power of an individual turbine is calculated as follows:

$$P_t = \frac{1}{2} \rho A V^3 C_P(V) \quad (6)$$

where P_t is the power produced by a single turbine; ρ is the density of air, which we assumed is 1.225 kg/m^3 ; A is the rotor swept area of the wind turbine; C_P is the power coefficient of the turbine; and V is the effective wind speed across the swept area, which was calculated with the wake model discussed in Sec. 2. When there is a variable number of turbines, we expect that a wind power plant optimized for maximum AEP will have many turbines spaced close together, filling the available land. If there are no penalties for costs considered in the optimization, additional turbines will lead to an improved objective, even if they are extremely inefficient and operate with high wake interference.

3.2 Cost of Energy

In some optimization problems, AEP may not be an appropriate objective as it does not account for the added cost or complexity required to achieve gains in AEP. An example of this is wind turbine blade design, for which an increase in AEP comes at the cost of additional mass and, therefore, cost. In a situation like this, it may be more appropriate to perform multi-objective optimization or include both AEP and cost into a single objective. COE is another common metric used in wind power plant



design that captures both energy production and costs (Chen and MacDonald, 2014; Fleming et al., 2016; Stanley and Ning, 2019a). We calculated COE as a combination of costs divided by the AEP:

$$\text{COE} = \frac{\text{cost}}{\text{AEP}} \quad (7)$$

140

$$\text{cost} = \text{FCR}(\text{TCC} + \text{BOS}) + \text{O\&M} \quad (8)$$

where FCR is the annual fixed charge rate, which we assumed was 9.7% (Previsic, 2011); TCC is the turbine capital cost, which we assumed is \$829 per kW of plant capacity (Wiser et al., 2020); O&M is the operation and maintenance cost, which we assumed is \$44 per kW of plant capacity per year (Stehly and Beiter, 2020); and BOS is the balance of station cost. For this
 145 paper, we created a simple relation of BOS costs as a function of the installed wind power plant capacity from a set of higher fidelity BOS cost data (Key et al., 2020). This BOS cost function is shown in Fig. 2. As shown in the figure, the cost per kW decreases as the total capacity increases because of economies of scale. We expect a wind power plant optimized for COE to




Figure 2. The balance of station (BOS) cost model (per kW of installed wind power plant capacity) as a function of the total power plant capacity (Key et al., 2020).

have fewer turbines than one optimized for AEP. This objective heavily considers the additional costs from adding turbines to the wind plant. Extra turbines are only beneficial if the economies of scale from a cost perspective outweigh the losses from
 150 additional wake interference that is introduced.

3.3 Annual Profit

Another metric that may be used for an objective function is annual profit. Like COE, this objective can take into account both energy production and costs. Additionally, an objective of profit can consider more refined measures of the value of energy,



such as time-of-day pricing where the price of electricity varies depending on the time of day it is produced. This objective
155 would likely be of most interest to wind power plant developers, as opposed to AEP or COE previously discussed. For this
paper, we defined profit simply with a fixed power purchase agreement as:

$$\text{profit} = \text{AEP} * \text{PPA} - \text{cost} \quad (9)$$

where PPA is the power purchase agreement, which determines the monetary value of the energy produced. For this paper, we
assumed that the PPA is a constant, although we will vary this constant to study its effect during different optimizations. We
160 expect a wind power plant optimized to maximize profit to have fewer turbines than one optimized for maximum AEP, but
more turbines than one that is optimized for minimum COE. This objective still penalizes costs from adding more turbines, but
will find solutions with slightly suboptimal COE as long as the AEP gains lead to sufficiently increased revenue.

4 Design Variable Parameterizations

In this paper, in addition to the objective function, we explore different optimization techniques, and how they affect the final
165 solution and the computational expense required to find it. One important part of any optimization is how to parameterize
the design variables. In this section, we explain the two different parameterization methods we used in this paper: a gridded
domain, where the number of design variables increases as the grid refinement squared; and a boundary-grid method, where
the number of design variables remains constant at 11, regardless of the size of the domain or the number of turbines.

4.1 Gridded Domain Design Variables

170 The first set of design variables that we used in our optimization are similar to those initially used by Mosetti et al. (1994)
and involve dividing the domain into a square grid of potential turbine locations. In this problem formulation, each of the
grid points is a design variable, with the possible integer value of 1 (meaning a turbine exists in the associated position) or 0
(meaning the associated position is empty). Figure 3 shows this gridded domain for a square boundary with 8 row and column
grid discretizations. Each of the blue points represents a design variable and is a potential location for a wind turbine. The
175 computational expense required to optimize a problem generally scales poorly as the number of design variables increases.
So, the grid must be refined enough to sufficiently search the design space, but not so refined that the optimization becomes
computationally infeasible. Note that the number of design variables increases as the grid refinement squared, indicating that
the number of design variables can quickly become impossible to optimize if the grid becomes too refined.

4.2 Boundary-Grid Design Variables

180 The second parameterization that we optimized with was a modified version of the boundary-grid method. The boundary-grid
parameterization is a simple method to define the layout of turbines in a wind power plant with very few design variables
and still achieve layouts that perform just as well as wind power plants designed with more complex layout optimization
techniques. In essence, it consists of placing some of the turbines around the boundary and the rest regularly arranged in a grid




Figure 3. A square wind power plant that has been discretized with a square grid for wind turbine number and layout optimization.

(Stanley and Ning, 2019b). The boundary-grid parameterization has the huge benefit of keeping the same number of design variables regardless of the number of turbines being optimized. This means that the layout of large wind power plants with hundreds or even thousands of wind turbines could be optimized without prohibitively high computational expense. In its original formulation, the boundary-grid method was defined for use with a gradient-based optimizer. This required the user to predefine some of the discrete variables that could not be optimized with gradients. Because we used gradient-free optimization in this paper, we slightly reformulated the boundary-grid method allowing integer variables and a discontinuous design space. In total, there are 11 design variables that describe the location of every turbine and are shown in Fig. 4.

The turbines in the interior of the wind power plant are arranged in a grid that is defined with 9 variables. First is the grid border, B , shown in Fig. 4a. The grid border is a number between 0 and 1 that defines the fraction of the boundary that will contain the inner grid turbines. When $B = 1$, the grid border is exactly the same as the wind power plant boundary, and proportionally decreases in size until $B = 0$, meaning the grid border vanishes in the center of the plant. The rest of the variables that describe the interior grid turbines define a complete distorted square grid of turbine locations; however, only those that are inside of the grid border are used. Figure 4b represents the other 8 variables used to define the locations of the interior grid turbines. The grid height and width are represented by h and w , respectively. The center of the grid is shown as the point (cx, cy) , and the grid shear is shown in this figure as ϕ . The grid rotation (about the center (cx, cy)) is given by θ . Finally, the number of rows and columns are also design variables used in the optimization. Figure 4c shows how the design variables in Figs. 4a and 4b are combined to obtain the turbine locations. Turbines are placed at all of the grid intersection points that are inside of the grid border, shown by the blue dots. No turbines are placed at the grid intersection points outside of the grid border, indicated by the red dots in the figure.

The turbines around the boundary of the plant are equally spaced, traversing the perimeter of the plant. The boundary turbine locations are defined by two design variables represented in 4d. First, the number of turbines placed on the boundary is a design

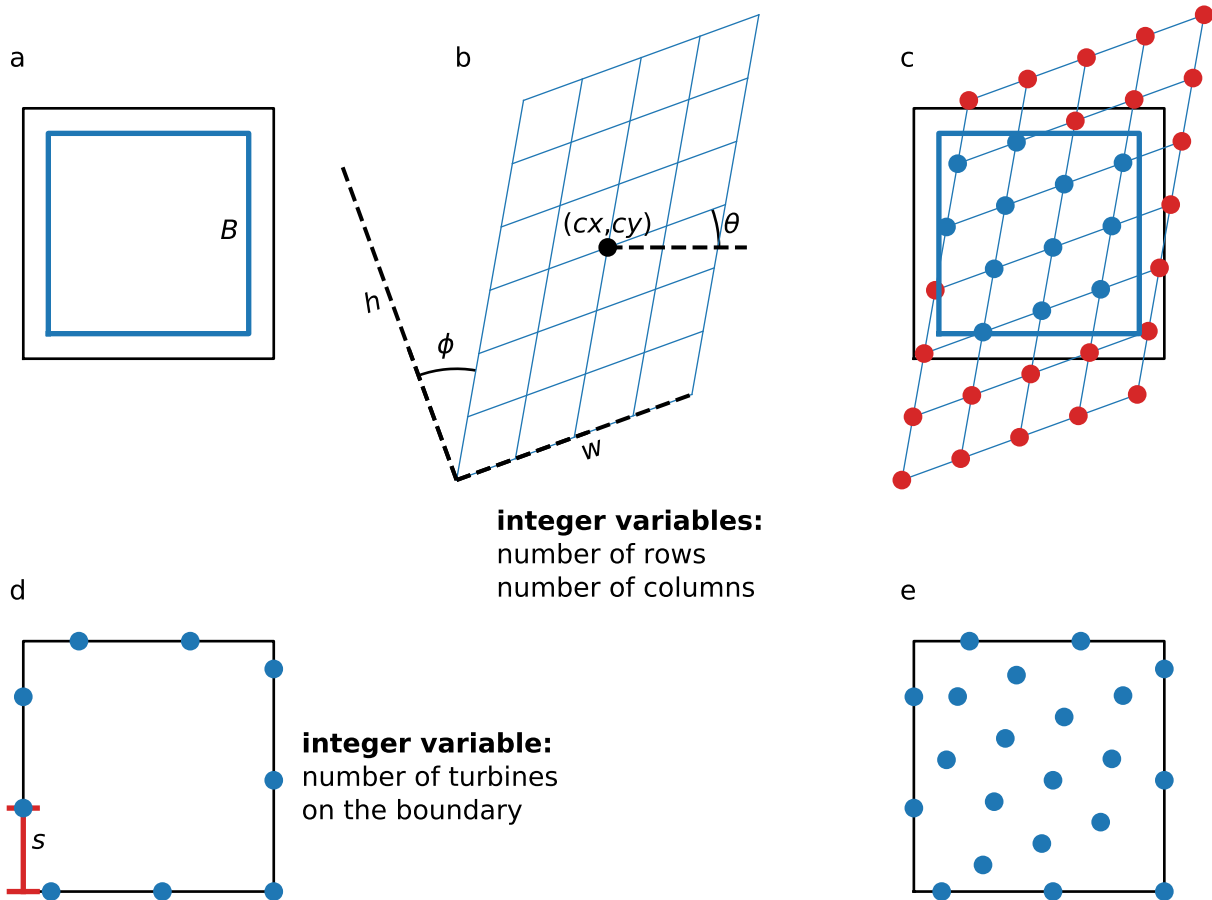


Figure 4. A description of the variables in the boundary-grid parameterization. 4a shows the grid border variable, B , and 4b shows the grid definition variables, h , w , (cx, cy) , ϕ , and θ , which represent the grid height, width, center, shear, and rotation, respectively. The number of grid rows and columns are also variables. 4c shows how all of the grid variables in 4a and 4b combine to define the interior grid turbine locations. 4d shows the boundary turbine variables, which are the boundary start location, s , and the number of boundary turbines. 4e shows the combined results of the boundary turbines and the interior grid turbines.

variable. Second, the starting location of the first boundary turbine, represented as s in Fig. 4d, is a design variable. The starting location is the distance from a constant anchor point at which the first boundary turbine is placed. Because the turbines are spaced equally around the wind power plant boundary, defining the location of this first boundary turbine implicitly defines the location of the rest of the boundary turbines. Figure 4e shows the final turbine locations defined by the boundary-grid parameterization variables shown in the rest of the figure. Notice that this is the combination of the boundary turbines in Fig. 4d and the blue inner grid turbines from Fig. 4c.



5 Optimization Algorithms

In this section, we discuss the details of the optimization algorithms we used in this paper. There are many algorithms that can be used to solve the wind power plant layout optimization problem, including determining the optimal number of wind turbines. In this paper, we chose to compare the performance of three gradient-free optimizers: a greedy algorithm, a genetic
215 algorithm, and a novel repeated-sweep algorithm.

5.1 Greedy Algorithm

The first optimization algorithm that we used was a greedy algorithm. Several researchers in the past implemented a greedy algorithm in performing wind power plant layout optimization, making this a good benchmark (Changshui et al., 2011; Song et al., 2015; Chen et al., 2016). We applied our greedy algorithm to the gridded plant parameterization. For this algorithm,
220 we started with one turbine placed in a random location within the plant domain. We then found the optimal location to place one additional turbine by evaluating the plant performance from placing the extra turbine at every potential turbine location in the grid. This process of adding one extra turbine was then repeated until adding an additional turbine did not cause an improvement in the objective. This algorithm is shown in Algorithm 1.

5.2 Genetic Algorithm

225 The second algorithm we used to optimize was a genetic algorithm. As with the greedy algorithm, genetic algorithms have also been a popular choice when performing wind power plant layout optimization studies with a discretized plant domain (Mosetti et al., 1994; Grady et al., 2005; Chen et al., 2013). For the results shown in this paper, we performed single-point crossover and used a mutation rate of 2%. For the gridded plant domain, adjacent bits in the chromosome were adjacent in the plant domain. This helped create offspring that did not violate spacing constraints, as entire sections of the wind power plant that
230 were traded during crossover would remain feasible (as long as they were feasible to begin with). After each generation, the entire population, consisting of parents and offspring, were ranked in order of performance. The better-performing half of the entire population was kept to act as parents for the next generation. Convergence was assumed after the best performance was within a tolerance of 10^{-3} for 25 generations. As the genetic algorithm was used for both the gridded parameterization and the boundary-grid method, continuous variables were binary encoded with 8 bits each. This means that for the boundary-grid
235 parameterization, the variables were encoded into 76 bits—3 bits for the integer number of rows, 3 bits for the integer number of columns, 6 bits for the integer number of turbines on the boundary, and 8 bits for each of the 8 continuous variables. A rule of thumb for genetic algorithms is to use a population size of 10 times the number of design variables (Martins and Ning, 2020). For the gridded plant domain, we followed this rule of thumb exactly because there was a large number of design variables. For the boundary-grid parameterization, we used a population of 100, which was slightly less than 10 times the number of design
240 variables. This gave us good results for our formulation. The genetic algorithm is represented is Algorithm 2.



Algorithm 1 Greedy Algorithm

Input:

number of bits: n

objective function: obj

```
1: initialize the best set of design variables  $A$ , an array of zeros of length  $n$ 
2: place a 1 in a random index of  $A$ 
3: initialize the best fitness value:  $f = \text{obj}(A)$ 
4: while not converged do
5:   set the iteration minimum to best fitness value:  $m = f$ 
6:   for  $i = 1 : \text{length}(n)$  do
7:     set temporary array  $B$  equal to  $A$ 
8:     switch index  $i$  of  $B$  (from 0 to 1 or 1 to 0)
9:     evaluate the objective function with temporary array  $B$ :  $g = \text{obj}(B)$ 
10:    if  $g < m$  then
11:      redefine the iteration minimum:  $m = g$ 
12:      define the iteration minimum array:  $C = B$ 
13:    end if
14:  end for
15:  if  $m = f$  then
16:    converged
17:  else
18:    redefine the best fitness value:  $f = m$ 
19:    redefine the best set of design variables:  $A = C$ 
20:  end if
21: end while
```

Output:

best fitness value: f

best set of design variables: A



Algorithm 2 Genetic Algorithm

Input:

number of bits: `n`
objective function: `obj`
population size: `pop_size`
maximum generations: `n_gens`
crossover rate: `crossover`
mutation rate: `mutation`
tolerance: `tol`
convergence generations: `conv_gens`

```
1: initialize the parent population parent_pop, a 2-D array of size pop_size by n
2: initialize the offspring population offspring_pop, the same size as parent_pop
3: initialize the generation counter, generation = 1
4: initialize the convergence criterion, difference = 1
5: while not converged and generation < n_gens do
6:   perform crossover to create offspring population
7:   mutate (leave the best individual unmutated)
8:   evaluate fitness of offspring
9:   rank the fitness of the total population
10:  evaluate difference, |best fitness of current generation – best fitness of generation conv_gens ago|
11:  if difference ≤ tol then
12:    converged
13:  else
14:    keep the best individuals between the current parent population and the offspring population as the parents for the next generation
15:  end if
16: end while
```

Output:

best fitness value
best design variables associated with the best fitness value



5.3 Proposed Method: Repeated Sweep Algorithm

The last optimizer we used was a novel repeated sweep algorithm. As far as we are aware, no method similar to our proposed repeated sweep algorithm has been used in past research for wind power plant layout optimization. Like the greedy algorithm, this optimizer was only used with the gridded plant parameterization, where each of the design variables is an integer, either 0 or 1. The repeated sweep algorithm consisted of three phases. First was a single search phase, followed by two trade phases.

In the search phase, one by one and in a random order, the value at each potential turbine location was switched from 1 to 0 or from 0 to 1. If the objective improved, the swapped value was kept; if not, the design variable retained its original value. This was done until every potential turbine location had been evaluated, and the value had been changed or retained. In both trade phases, each potential turbine location was again searched through one by one. However, in these phases, instead of exploring adding or removing turbines (like in the search phase), the potential turbine location traded values with the cell adjacent to it. In the way we formulated the problem, in the first trade phase, each position traded places with the cell to the right; in the second trade phase, each position traded places with the cell above it. As with the search phase, if a trade resulted in an improvement in the objective, the trade was kept. If not, the trade was rejected and the original locations were retained. Also, as with the search phase, the trades were done in a random order until a trade at each location had been evaluated. The three phases were repeated in order, search-trade-trade, until the objective function did not improve after a complete cycle of all three phases. The repeated sweep algorithm is shown in Algorithm 3.

5.4 Gradient-Based Optimization

The optimization algorithms discussed previously were gradient-free, which simultaneously optimize the number of turbines and their layout in a wind power plant. Another way to optimize turbine number and layout in a wind power plant is with gradient-based optimization. Gradient-based algorithms cannot optimize integer design variables or discontinuous design spaces—both of which are conditions that apply to the problem addressed in this paper. However, it is possible to repeat a gradient-based optimization multiple times with different numbers of wind turbines, then choose the overall best solution for the given objective. This process is computationally expensive for two main reasons. First, *a priori*, it is difficult to determine the approximate number of turbines that will be optimal. This means it would be necessary to repeat the optimization many times, using different numbers of wind turbines. Second, gradient-based optimizers are especially susceptible to converging to local minima in the design space. This problem is also prevalent in gradient-free optimization, but is more pronounced in gradient-based optimization. The problem can be mostly accounted for by repeating the optimization many times with different randomly initialized design variables, but this requires even more computation.

In this paper, our purpose was to compare some gradient-free methods that could be effectively used to solve for the optimal turbine number and placement in a wind power plant. However, for one small case discussed in Sec. 6.2 we also used gradient-based optimization in order to compare the results. For the optimizations in this section, the function calls were very fast, allowing us to quickly perform the hundreds of optimizations necessary to explore the design space. To perform the gradient-based optimization, we swept through all of the possible numbers of wind turbines that could fit in the wind power plant without



Algorithm 3 Repeated-Sweep Algorithm

Input:

number of rows and columns: n, m

objective function: obj

-
- 1: randomly initialize a feasible array of design variables A , a 2-D array of size n -by- m with 1s and 0s. The 1s in the matrix correspond to the physical locations of turbines in the wind power plant. Ensure that this initialized matrix satisfies all of the constraints.
 - 2: initialize the fitness: $f = \text{obj}(A)$
 - 3: **while** not converged **do**
 - 4: cycle to the next step: search, trade rows, trade columns. Loop back to the search step after completing trade columns. We started with the search step.
 - 5: **for** sweep through all of the indices of A in a random order **do**
 - 6: set temporary array B equal to A
 - 7: **if** search step **then**
 - 8: change the value of the current index of B from 0 to 1 or 1 to 0
 - 9: **else if** trade rows step **then**
 - 10: trade the value of the current index of B with the value on its right
 - 11: **else if** trade columns step **then**
 - 12: trade the value of the current index of B with the value above it
 - 13: **end if**
 - 14: **if** $\text{obj}(B) < f$ **then**
 - 15: keep the switched matrix $A = B$, and update the fitness value $f = \text{obj}(B)$
 - 16: **end if**
 - 17: **end for**
 - 18: **if** three consecutive steps are completed (in any order) without improving the fitness value **then**
 - 19: converged
 - 20: **end if**
 - 21: **end while**
-

Output:

best fitness value: f

best set of design variables: A



violating the spacing constraints, which was 2–18 turbines. For each number of turbines, we performed 50 optimizations with randomly initialized turbine locations. This gave us relatively high confidence that the solution we found with the gradient-based optimization was near the global solution. We used finite-difference gradients for these optimizations, which do not perform as well as analytic gradients, both in quality of the final solution and in computational expense. However, for the case in which we used the gradient-based optimizer, the wind rose is simple and the number of design variables is relatively small, meaning the finite-difference gradients perform sufficiently well. For the results in this paper, we used the open-source **SLSQP** optimizer available in SciPy (Virtanen et al., 2020).¹

5.5 Constraints

In our layout optimizations, we assumed there were only two constraints—a spacing constraint and a boundary constraint. The turbines were constrained to be at least two rotor diameters apart from each other. This minimum spacing constraint implicitly defined the maximum number of turbines that could be placed in the wind power plant. Additionally, turbines were constrained to remain inside a prescribed boundary.

6 Results

In this section, we discuss the results of our wind power plant simulations and optimizations. Included in this section are a simple, one-dimensional (1D) sweep of the different objective functions versus the number of wind turbines, then full two-dimensional (2D) wind power plant layout optimizations run for the different objectives and with the different optimization algorithms. The wind plants that we optimized and discuss in this section are a small wind power plant with a unidirectional wind rose, a large wind power plant with a unidirectional wind rose, and a large wind power plant with a full wind rose. Finally in this section, we present results from optimizing wind power plants for maximum profit with varying PPAs.

6.1 1D Example

First we discuss a simple, 1D example to demonstrate the effect different objective functions have on the optimal solution. For this example, we simulated a single row of wind turbines in line with the wind, which had a constant speed of 10 m/s. The length of this row was fixed at 25 km, and the turbines were equally spaced. For this scenario, we calculated the value of each objective as a function of the number of wind turbines in the simple wind power plant; results are shown in Fig. 5.

This figure is particularly interesting when you notice that the optimal number of turbines for each objective is very different. Obviously, a wind plant designed for maximum AEP will look very different than wind plants designed with other objectives in mind. When maximizing AEP, there is no penalty for the extra costs associated with building extra turbines. As long as adding another turbine produces more energy, it is superior—no matter how marginal the increase in energy and how large the increase in cost. It makes sense that the wind power plant optimized for AEP has the most wind turbines, 23 in this example, because more turbines are added until the wake effect from adding an extra turbine outweighs the additional power it provides.

¹<https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.minimize.html>




Figure 5. Different objectives as a function of the number of turbines in the wind power plant. For this example, the turbines are in line with the wind direction and are equally spaced in a wind power plant with a fixed length. From left to right, the objectives represented here are annual energy production (AEP), cost of energy (COE), and annual profit.

On the other hand, when COE is the objective, the optimal number of wind turbines is just nine, much lower than for the AEP objective. If the cost of the wind plant was modeled as a linear function of the number of turbines, the optimal COE solution would be just one turbine. A single turbine would have no wake interference from other turbines and would, therefore, produce energy for the lowest cost. However, there are some economies of scale involved with wind plant development, represented in our cost model by the decrease in BOS costs with increasing power capacity. This means there is some optimum greater than one where the wake interference is still relatively low and the costs per turbine in the wind plant are decreasing steeply with additional turbines.

Finally, a completely different solution is obtained when optimizing the wind power plant for maximum profit. While the COE objective optimizes the ratio between the value a wind plant produces and the cost, the profit objective optimizes the difference between the value a wind plant produces and the cost. At first, it may seem nonintuitive that the solution for optimal profit is different than the solution for optimal COE because minimized costs should be related to more profit. In Fig. 5, notice that the optimal COE solution is 9 turbines and produces a COE of about \$23/MWh. In this case, energy generation and, therefore, revenue generation, are limited because of the small number of turbines. For 18 wind turbines, a slightly higher COE is achieved of about \$25/MWh. From a COE perspective, this is suboptimal. However, the additional revenue produced from the extra turbines outweighs the increase in COE. The exact number of turbines for optimal profit depends on the monetary value of the energy, which is defined with the PPA. This means that the optimal solution is different depending on the PPA, represented by the different colors in the subfigure on the right.



As described previously, Fig. 5 shows different metrics as a function of the number of turbines in a 1D wind power plant. In addition to the number of turbines shown on the bottom axis, this figure also shows the turbine spacing in rotor diameters on the top axis. While this axis is useful in understanding the results from a more familiar perspective, one must be careful not to interpret these values incorrectly. The optimal turbine rotor diameter spacing reported in this figure is around 10 for AEP, 24 for COE, and between 12 and 18 for profit. These are very large turbine spacings for a wind plant and are likely infeasible because of land or cabling constraints. One must remember that these results were obtained from a very simple design space sweep with all wind turbines exactly in line with the wind. Because real wind power plants are built in two dimensions, with full wind-direction variability, it is actually possible to build turbines much closer together than is indicated in Fig. 5.

As demonstrated in Fig. 5, when determining the number of wind turbines to build in a wind power plant, an appropriate objective is essential to achieving a desirable solution. This is true of any optimization problem, but it is particularly important to remember for this application. COE is an extremely common objective function in wind plant design—and rightfully so. However, as demonstrated in this simple example, the optimal number of turbines to minimize COE is suboptimal if the aim is to maximize annual profit, which may (or may not) be what is most important to those designing the wind power plant. In this specific example, the optimal number of turbines for COE results in \$5.1 million of annual profit with a PPA of \$50/MWh, just 72% of the optimal profit of \$7.1 million. This significant difference in the optimal performance and wind plant design for different objectives has important implications for techno-economic considerations in wind plant design. Economic factors drastically change the optimal solution, which highlights the importance of having accurate cost models, and again identifying the correct objective for design and optimization.

Historically, capacity expansion models consider the price of power and land availability but assume a constant power density for a single turbine configuration and do not vary with PPA price or cost. This work illustrates not only the difference in capacity density for a COE minimization versus a profit maximization, but also the relationship of increasing energy prices on capacity in land-constrained scenarios. Aggressive carbon reduction scenarios or other renewable energy goals typically result in higher energy prices for renewables. These results in Fig. 5 show that for scenarios where energy prices increase, capacity densities may also increase. Similarly, where project costs increase because of smaller plant sizes and higher construction costs, projects may tolerate lower energy production from lower wind resource quality and higher wake losses in order to maintain a constant PPA price.

6.2 Small Plant with Unidirectional Wind Rose

With the 1D sweep of the design space complete to provide some intuition about the different objective functions, we will now discuss a simple layout optimization for a small wind power plant with a unidirectional wind rose. As stated before, we performed the optimization of each objective using a gridded domain, optimized with a greedy algorithm, a genetic algorithm, and a repeated sweep algorithm. We also optimized a boundary-grid layout parameterization with a genetic algorithm. Also, as mentioned before, for this small wind plant, we optimized the layouts using gradient-based optimization.

For this small wind power plant optimization, we assumed the domain was square with 800-m sides. The wind came from a direction of 300 degrees, or 30 degrees north of west. The wind speed was assumed constant at 10 m/s, which is close to the



355 rated wind speed for our turbine model. The PPA was assumed to be \$30/MWh. For the gridded design variables, the domain was discretized into a 10-by-10 grid. We ran each optimization method five times to convergence because the final solution is dependent on the randomly initialized population or design space. Because each of the optimization algorithms has some stochastic qualities, with enough time and randomly initialized starts, each optimization method will potentially be able to find a very good solution. However, we believe that five optimizations for each is enough to give a good idea of their performance relative to each other for each of the objective functions.

Results for the small wind power plant optimizations with a unidirectional wind rose are shown in Table 2 and Fig. 6. Table 2 shows the optimization results and the computational expense associated with each optimization method and for each objective function. The first column shows the objective function, and the second column shows the optimization method. The third column provides the optimized number of turbines in the wind plant in bold and the average turbine spacing in rotor diameters associated with that number of turbines. The fourth column provides the best solution from the 5 optimizations in bold, along with this same value normalized by the best solution out of all of the optimization methods for the given objective. In this column, the best solution is highlighted in blue and the worst solution is highlighted in red. Finally, the fifth and sixth columns provide the total time and the total number of function calls to run the 5 optimizations for each optimization method. The gray rows in this figure show the gradient-based optimization results. Notice one cell in a gray row is blue, indicating the gradient-based optimization found the best solution for the maximum profit objective. Figure 6 shows the flow field for the best layout for each objective function. These are the layouts corresponding to the blue highlighted cells in Table 2. These flow fields show a horizontal slice of the wind power plant at the turbine hub height. The black lines represent the wind turbines, the red areas represent faster freestream wind speed, and the blue areas represent a slower waked wind speed. We did not include a color legend because we only wish to demonstrate qualitative information with this figure; therefore, exact values are not important for this purpose.

6.2.1 Small Power Plant with Unidirectional Wind Rose: Different Objectives

First, we will discuss the differences between the optimal solutions for the different objective functions. For optimal AEP, the best solution has as many turbines as the optimizer can fit into the wind power plant without violating spacing constraints. As can be seen in the top subplot of Fig. 6, the optimal layout has turbines that are spaced very close together. Wakes are strong in the flow field, which contains several turbines that are fully or partially waked. For this objective, it doesn't matter if some turbines are greatly affected by wakes as long as their energy contribution is positive. Now, in this case, the optimal solution had the maximum number of turbines as could fit into the boundary. However, from Fig. 5 we see that even for the AEP objective, there is a point where adding additional turbines actually becomes detrimental. We also see from Fig. 5 that this could occur at a relatively large turbine spacing, between 9-10 rotor diameters. For the 1D sweep, the turbines are all exactly in line with the wind. Additionally, rather than having two or three turbines waked in line, there are many in line with each other. This indicates that, in large part, the AEP is determined by deep array effects. For the small wind plant layout optimization discussed in this section, there are at most three turbines in line with each other. In this case, adding turbines, even if they are



Table 2. Complete optimization results for the small wind power plant with unidirectional wind rose.

objective	optimization	num. turbines/ avg. spacing (D)	optimal value/ normalized	time (s)	function calls
AEP (GWh)	greedy grid	12/2.76	238/0.792	15	2,267
	genetic grid	16/2.26	293/0.973	1,867	105,872
	sweep grid	15/2.36	270/0.897	4	204
	genetic BG	16/2.26	301/1.000	1,022	50,519
	GB	17/2.17	299/0.993	12,082	651,265
COE (\$/MWh)	greedy grid	10/3.14	22.16/1.015	11	1,860
	genetic grid	11/2.93	21.84/1.000	1,443	106,755
	sweep grid	10/3.14	21.88/1.002	8	656
	genetic BG	9/3.40	22.16/1.014	319	26,604
	GB	11/2.93	21.93/1.004	12,071	655,392
annual profit (\$MM)	greedy grid	11/2.93	1.88/0.918	15	2,208
	genetic grid	12/2.76	1.99/0.971	1,306	95,687
	sweep grid	12/2.76	1.76/0.860	5	341
	genetic BG	14/2.48	1.85/0.905	532	36,510
	GB	13/2.61	2.04/1.000	13,130	702,650

fully waked, increases the AEP. If we were to repeat the optimization for a much larger domain we would see results similar to Fig. 5, where having too many turbines could actually be detrimental for AEP.

390 While the wind power plant optimized for AEP maximizes the number of turbines in the design space, the wind plant optimized for minimum COE looks very different. This wind plant has 11 turbines, as opposed to 16 for maximum AEP. The turbines are arranged such that waking is minimal, if it exists at all. For this objective, it appears that the optimizer maximizes the number of unwaked wind turbines. We can conclude that when the turbines are waked, the loss in energy production outweighs the benefits gained from economies of scale in the cost model. Therefore, additional turbines are good if they meet some minimum power production requirement, which is dictated by the cost model.

395 For the last objective, profit with a PPA of \$30/MWh, the optimized number of turbines is 13. This is between the optimal number of turbines for the COE and AEP objectives. When optimizing for profit, the solution appears to be a balance between minimizing COE and maximizing AEP. A turbine is allowed to be waked as long as the gains from the energy it produces outweigh the costs of adding the extra turbine. As will be discussed in more detail in Sec. 6.6, the point where adding an additional turbine is no longer profitable is determined by the PPA. A lower PPA will drive the solution for maximum profit toward the solution for minimum COE, while a higher PPA will drive it toward the solution for maximum AEP.

We want to emphasize that the results we show are not meant to demonstrate exact solutions or guidelines for determining the number of turbines in a wind power plant. Specific solutions will depend on wind resources, turbine parameters, boundary

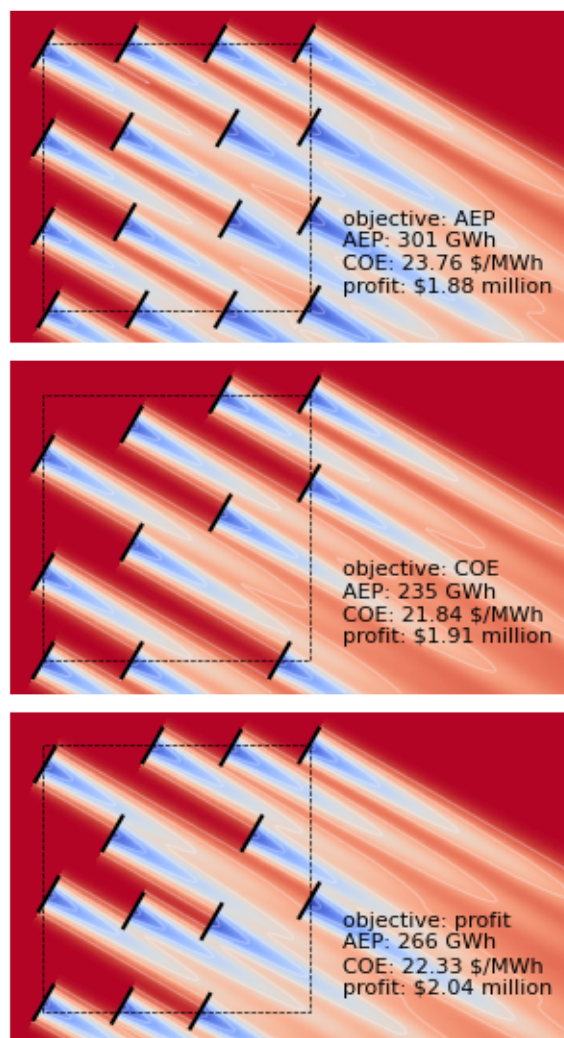


Figure 6. The optimal layouts for each objective for the small wind power plant with a unidirectional wind rose. From top to bottom, the associated objective functions are AEP, COE, and profit. The text within each figure provides the values for all three metrics for each wind plant.

shape and size, PPA, constraints, and other factors. Our purpose is to demonstrate that the optimal number of turbines and their
 405 layout are completely different depending on the objective. The true objective must be carefully formulated when optimizing
 the layout of a wind power plant. Figure 7 shows how the wind plants optimized for the three different objectives compare in
 other metrics. We demonstrated objectives of AEP, COE, and profit, and how they all produce different solutions. All of the
 wind plants optimized for a specific metric greatly underperform in the other metrics that we calculated. The one exception is
 the optimal profit solution, which also accomplishes a relatively low COE.




Figure 7. A comparison of the performance metrics of the three different wind power plants optimized for different objective functions. These results are shown for the small wind plant with a unidirectional wind rose.

While we included three specific objective functions in this paper, there are many other considerations that could be included in the objective. For example, it may be desirable to maximize the profit generated by each turbine in the wind power plant above some minimum value. This would keep the optimizer from adding a turbine that only provided a marginal return on investment. One could also optimize for profit or COE while constraining the AEP to be above some desired minimum value. When using mathematical optimization, the objective function must be designed to truly represent the desired performance because this will drive the final solution.

6.2.2 Small Power Plant with Unidirectional Wind Rose: Different Algorithms

In this section, we will discuss the performance of the solutions found with each optimization strategy and their computational expense. This information is presented in the last three columns of Table 2. As explained previously, the third to last column shows the optimized solution in bold, as well as the normalized value, to easily see how the optimized solutions compare to each other. The last two columns are different measures of the computational expense. The time column is straightforward and is the total wall time required to run the optimizations. For this paper, everything was run without parallelization on a laptop with a 2.4 GHz 8-core Intel processor. However, just the time as a measure of computational expense may be misleading. There is other overhead in the optimization time other than just objective function calls; therefore, we included a column for total objective function calls and run time, which together give a decent representation of the total computational expense of each algorithm.

For this small wind power plant with the unidirectional wind rose, the greedy and repeated sweep algorithms do not perform very well. Although they are by far the best when it comes to required computational resources, both algorithms find comparatively poor solutions. Both of these algorithms rely on placing turbines far apart to get the maximum benefit possible at each step of the optimization. Because the domain is small, this makes it difficult to add additional turbines without violating



430 the spacing constraint. Doing so would require adjusting the location of multiple turbines at once to make room, which is not something the algorithms can do.

The genetic algorithm with the gridded plant domain performs very well for each objective function. It even outperforms the gradient-based optimization for the COE objective and performs within 3% of the best solution found for the AEP and profit objectives. However, the computational expense was high, requiring the most time of any of the gradient-free algorithms and by
435 far the most function calls. However, because this problem is relatively small, the computational expense was not prohibitive.

The boundary-grid optimization solved with a genetic algorithm performs in the middle of the pack for the gradient-free algorithms. It performs very well with the AEP objective, but poorly with the COE and profit objective. This is because of the small wind plant area. The boundary-grid formulation forces turbines to be equally spaced around the boundary. With a unidirectional wind rose, this means that some turbines will always be in the back of the power plant relative to the incoming wind.
440 As discussed before, waked turbines were very detrimental for COE and, by extension, detrimental from a profit perspective as well.

Finally, the gradient-based optimizer, while sweeping across the number of wind turbines, performs well for each objective function. Because this algorithm allows the most freedom, permitting each turbine to be placed wherever the optimizer deems best, we always expected the gradient-based optimizations to perform well. In fact, we expected this optimizer to perform
445 the best of all for each objective, making it quite surprising that it is outperformed for both the AEP and COE objectives. Even though the gradient-free algorithms did not provide as much freedom, they were able to find the best solution for these objectives.

At this point, we would like to reiterate that the results shown in Table 2 are for a limited number of starting optimizations. They do not indicate that the solutions found are the best solution that each optimizer is capable of finding. These results simply
450 show the optimizer performance with a small sample size. That said, for a small number of discretizations, finding the optimal number and layout of wind turbines for various objectives can be achieved with a simple genetic algorithm.

6.3 Large Power Plant with Unidirectional Wind Rose

The performance and required computational expense of optimization algorithms can vary dramatically depending on the problem size. In this section, we will present another set of optimizations we ran for a larger domain. This wind power plant is
455 a square with 1.6-km sides. For the gridded domain, the plant is divided into a 20-by-20 square grid, which maintains the same spacing between grid points as in the small wind plant example. For a wind plant of this size, we did not run the gradient-based optimizer because of the large computational expense required to run the sweep across all of the possible optimal number of turbines. We only ran the optimizations for the four gradient-free methods we previously discussed. As in the previous section, we ran each optimization method 5 times to convergence for each objective. The wind resource and PPA for this wind power
460 plant were also assumed to be the same as for the small wind plant. The results for this wind power plant optimization are presented similarly to those for the small wind plant, with full results shown in Table 3, and the best layout for each objective shown in Fig. 8.



Table 3. Complete optimization results for the large wind power plant with unidirectional wind rose.

objective	optimization	num. turbines/ avg. spacing (D)	optimal value/ normalized	time (s)	function calls
AEP (GWh)	greedy grid	40 /2.55	663 /0.920	746	34,012
	genetic grid	23 /3.58	450 /0.624	2,539	44,095
	sweep grid	47 /2.32	710 /0.985	113	1,220
	genetic BG	48 /2.29	721 /1.000	3,310	48,824
COE (\$/MWh)	greedy grid	22 /3.68	20.51 /1.007	423	28,393
	genetic grid	16 /4.53	21.00 /1.031	2,656	63,104
	sweep grid	20 /3.91	20.37 /1.000	119	4,895
	genetic BG	18 /4.19	20.73 /1.018	1,002	35,814
annual profit (\$MM)	greedy grid	31 /2.97	4.98 /0.979	592	31,902
	genetic grid	20 /3.91	3.33 /0.655	2,826	47,977
	sweep grid	31 /2.97	4.93 /0.970	116	2,528
	genetic BG	33 /2.86	5.09 /1.000	2,126	52,869

6.3.1 Large Power Plant with Unidirectional Wind Rose: Different Objectives

The general trends that we observed from the smaller wind power plant optimizations hold true for this larger wind plant as well. The wind plant optimized for AEP has as many wind turbines in the plant as possible without violating the spacing constraints, leading to a large number of turbines. The wind plant optimized for COE has turbines that are minimally waked, leading to a very small number of turbines. The wind plant optimized for profit is a middle ground between the previous two objectives. One main difference between the results for this large wind plant and the small wind plant is in the optimal solution for minimum COE. For the large wind plant, some of the turbines in the layout optimized for COE are more waked than in the small wind plant. However, these waked turbines are far downstream of the waking turbine, meaning that the wind speed in the wake has recovered much of the way back to the freestream wind speed. In the small wind plant, there was not much space for the wakes to recover, meaning any amount of waking on a turbine was more detrimental to COE.

Figure 9 shows several metrics for the three wind power plants that were optimal for the different objective functions. Although the trends in this figure are similar to those in Fig. 7, the differences in metrics are more extreme for the different objective functions. While all of the metrics are interesting, there are a few specific observations that are worth pointing out. First, the COE for the wind plant optimized for profit is relatively low, which is impressive because COE was never directly minimized. Second, the profit for the wind plant optimized for AEP is very low. Even though the plant produces a lot of energy, the tight turbine spacing causes high wake losses and inefficient turbines; thus, the high costs are not offset by the high energy production. Third, the wake losses for the wind plant optimized for minimum COE are very low, less than 2%. This is impressively low and is only possible because of the unidirectional wind rose.

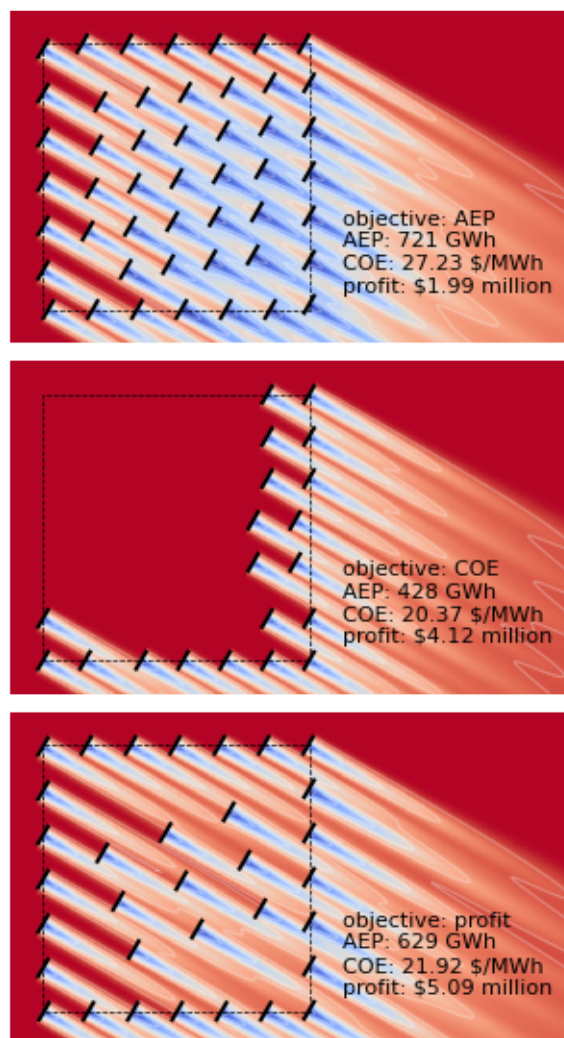


Figure 8. The optimal layouts for each objective for the large wind power plant with a unidirectional wind rose. From top to bottom, the associated objective functions are: AEP, COE, and profit. The text within each figure provides the value of all three metrics for each wind plant.

6.3.2 Large Power Plant with Unidirectional Wind Rose: Different Algorithms

While the general trends found for the solutions with different objectives were similar between the small and large wind power plants, the computational expense and performance of different algorithms were not. In the last three columns of Table 3, we see how well each optimization method performs. The most glaring difference is seen in the performance of the genetic algorithm with the gridded turbine domain. While with the small wind plant, this method provided the best or near-best results for each




Figure 9. A comparison of the performance metrics of the three different wind power plants optimized for different objective functions. These results are shown for the large wind plant with a unidirectional wind rose.

objective, in this larger domain it severely underperformed compared to the other optimization methods. While the genetic algorithm was easily able to handle the 100 design variables from the 10-by-10 grid of the small wind plant, it appears unsuited to find a good solution for the 20-by-20 grid of the large wind plant. This observation agrees with our previous intuition with genetic algorithms in that they tend to perform poorly as problem complexity increases.

490 Next, the greedy algorithm and the repeated sweep algorithm both perform relatively better in the larger wind plant than they did in the smaller wind plant. For the AEP and profit objectives, these algorithms do not perform the best, but produce respectable results. For the COE objective, there are three different algorithms that provide basically the same result with optimal solutions within 1% of each other.

Finally, the boundary-grid algorithm performs excellently for each objective function for this large wind power plant optimization. The organized structure forced by the boundary-grid method is able to fully take advantage of the unidirectional wind rose and create layouts with the appropriate number of turbines where waking is minimal. Even though it uses a genetic algorithm, the boundary-grid optimization performance does not suffer with the increased size of the wind plant. With this formulation, the number of design variables remains constant, independent of the number of turbines being modelled and the size of the domain.

500 6.4 Large Power Plant with Full Wind Rose

The final scenario in which we will examine the performance of each optimization algorithm is for the large wind power plant domain. Unlike the previous examples, the optimization results shown in this section include a full wind rose, discretized into 72 wind direction bins and shown in Fig. 10. The wind speeds are still assumed to be uniform at 10 m/s from each wind direction. Everything else for this optimization scenario is the same as in Sec. 6.3, including the wind plant dimensions, the grid discretizations, number of random starts, and PPA. The results for this section are shown in Table 4 and Fig. 11. Because
 505 the wind resource was unidirectional in the previous wind power plant layout figures, only one figure was needed to represent



the flow field associated with each layout. Because this section includes optimization results that were run with a full wind rose, we show two wind directions in Fig. 11, with each column representing the two dominant wind directions of west and south-southeast (shown in Fig. 10). This wind rose is similar to the one used in the first International Energy Agency (IEA) Task 37 wind plant layout optimization case study, with more finely discretized bins (Baker et al., 2019). The wind rose was specifically chosen to have a dominant wind direction in line with the upper and lower wind plant boundaries. Because the boundary-grid method is formulated to place turbines around the wind plant boundary, we wanted to see how it would perform with an unfavorable wind rose.




Figure 10. The full wind rose used for the optimizations in this study. The wind rose is divided into 72 bins with dominant wind directions from due west and from south-southeast.

6.4.1 Large Power Plant with Full Wind Rose: Different Objectives

Again, many of the general trends for the different objectives are the same as the results for the unidirectional wind rose cases. In this section, we will only focus on the key differences between the results for the large wind power plant with a unidirectional wind rose and the results for the large wind plant with the full wind rose. There are two observations we would like to point out. First, and most importantly, the optimal number of turbines for the COE and profit objectives is much lower for the full wind rose than for the unidirectional rose. For the unidirectional rose, the optimal numbers of turbines for COE and profit were 20 and 33, respectively. For the full wind rose, these numbers were 15 and 24, respectively. With the full wind rose, turbines will always be unfavorably waked for some wind directions. It is impossible to find a turbine layout in which all of the turbines are always unwaked. For the unidirectional rose, the turbines can be very closely spaced in directions not aligned with the wind



Table 4. Complete optimization results for the large wind power plant with full wind rose.

objective	optimization	num. turbines/ avg. spacing (D)	optimal value/ normalized	time (s)	function calls
AEP (GWh)	greedy grid	43 /2.44	582 /0.928	38,329	35,680
	genetic grid	26 /3.31	429 /0.685	37,208	48,739
	sweep grid	50 /2.24	604 /0.964	5,560	1,166
	genetic BG	54 /2.14	627 /1.000	157,191	41,175
COE (\$/MWh)	greedy grid	14 /4.95	22.61 /1.002	10,434	21,561
	genetic grid	11 /5.86	22.90 /1.015	28,386	46,823
	sweep grid	15 /4.73	22.57 /1.000	7,400	7,812
	genetic BG	12 /5.51	22.74 /1.007	25,260	37,071
annual profit (\$MM)	greedy grid	23 /3.58	2.68 /0.967	22,020	29,538
	genetic grid	19 /4.04	2.27 /0.821	34,824	47,266
	sweep grid	24 /3.48	2.77 /1.000	8,277	4,201
	genetic BG	24 /3.48	2.59 /0.936	60,995	40,959

and still avoid wakes and perform well in each objective. This luxury does not exist for the full wind rose. Turbines spaced close together will perform poorly at least some of the time. Therefore, fewer turbines are placed in the optimal wind plants so that the average spacing can be higher and reduce the wake effects. Second, notice that the optimal number of turbines for the AEP objective is higher for the full wind rose than for the unidirectional wind rose. This is simply an artifact of our methodology. For the AEP objective, the unidirectional wind rose would also benefit from having more turbines spaced closer together. However, the optimizer didn't find this solution from the five optimizations that we ran.

Similar to Fig. 7 and Fig. 9, Fig. 12 shows the different wind power plant performance metrics for each optimal wind plant with the different objectives. The general trends are similar to Fig. 9 with a few differences. First, and most apparent, the profit for the wind plant optimized for AEP is extremely low. In fact, this wind plant loses more than \$3 million each year because both the costs and wake losses are so high. There are at least two reasons for this extreme negative value for profit: 1) Our cost model has low economies of scale. At large numbers of turbines, the cost reductions for adding more turbines is low. And, 2) our minimum spacing constraints for this paper were extremely low, only two rotor diameters. At this close spacing, the wake losses are high. For wind plants where the minimum spacing is much larger (5 or more rotor diameters), a wind plant optimized for AEP would still perform well in the other objectives. The second difference in general trends is that the wake losses for these wind plants are much higher than for the optimal wind plants and the unidirectional wind rose. The reason for this was previously mentioned—with multiple wind directions, there will always be some wind directions for which a turbine operates within a wake.

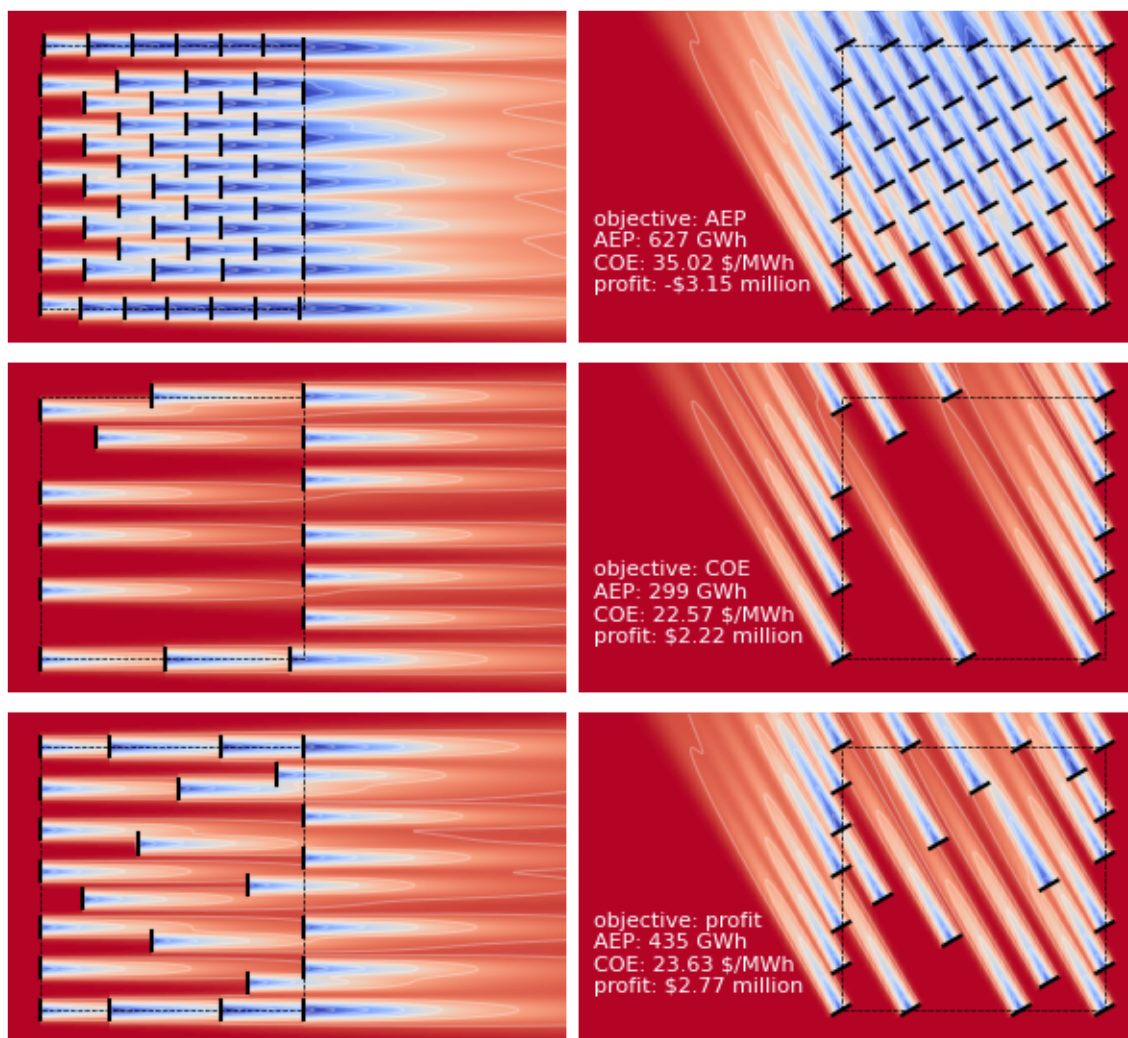


Figure 11. The optimal layouts for each objective for the large wind power plant with a full wind rose. From top to bottom, the associated objective functions are AEP, COE, and profit. The text within each figure provides the value of all three metrics for each wind plant. The two columns represent the flow field for each objective from the two dominant wind directions.

540 6.4.2 Large Power Plant with Full Wind Rose: Different Algorithms

There are three main differences we observe between the unidirectional wind rose algorithm performances and those for the full wind rose. First is the expected increase in computational time. The function calls for each algorithm and for each objective are very close for the different wind roses. However, the computational time for the full wind rose is significantly and expectedly higher. The computational expense of each function call for the full wind rose is about 72 times that of the unidirectional rose—




Figure 12. A comparison of the performance metrics of the three different wind power plants optimized for different objective functions. These results are shown for the large wind plant with a full wind rose.

72 being the number of wind direction bins. The reason the time doesn't scale exactly linearly with the number of function calls is the certain amount of infeasible layouts that are produced during the optimization that do not call the full wake model. These calls are very fast, but not instantaneous, and do add up over time.

The second observation is about the impressive performance of the boundary-grid problem formulation, even though the wind rose is specifically selected to put this method at a disadvantage. As expected, the boundary-grid method tremendously outperforms the other methods for the AEP objective because this formulation can add many turbines to the wind power plant with very small adjustments to the design variables. For the COE objective, the boundary-grid optimization performs well, within a percent of the best solution found. The boundary-grid optimization performs the worst for the profit objective, about 6.5% worse than the best solution found. For many wind plants, the boundary-grid method would perform much better because prominent wind directions will not always be aligned with the wind plant boundaries.

The third and final observation is about the performance of the repeated-sweep algorithm. For the AEP objective, this algorithm performs worse than the boundary-grid method, but still produces a wind power plant layout that performs favorably compared to the overall optimal, within 4%. The repeated-sweep algorithm is ineffective at placing as many turbines as possible into the wind plant, but still performs relatively well. For the COE and profit objectives, the repeated-sweep algorithm finds the best solution. For the COE objective, the greedy, repeated-sweep, and boundary-grid optimizations all perform similarly. However, for the profit objective, the repeated-sweep algorithm impressively outperforms all of the other algorithms. With a full wind rose and objectives that favor solutions where the turbines are minimally waked, this algorithm performs extremely



well. Because the best solutions have turbines that are spaced farther apart, the optimizer is able to search the design space without violating the turbine spacing constraints. In addition to finding superior optimal solutions, the repeated sweep algorithm required much less computational expense. For this problem size, the repeated-sweep algorithm performed the best overall. In short, the boundary-grid method performed relatively well even with an unfavorable wind rose and boundary. The repeated-sweep algorithm performed well for the COE and profit objectives and optimized very quickly compared to the other methods.

6.5 Overall Algorithm Performance

In this section, we will discuss the overall performance of each optimization algorithm for the small wind power plant with a unidirectional wind rose, the large wind plant with the unidirectional wind rose, and the large wind plant with the full wind rose. Figure 13 shows the overall performance of each algorithm and the computational expense for each objective and size of wind plant. Much of this information has been discussed in previous sections, but we will review each of the algorithms from an overall perspective.

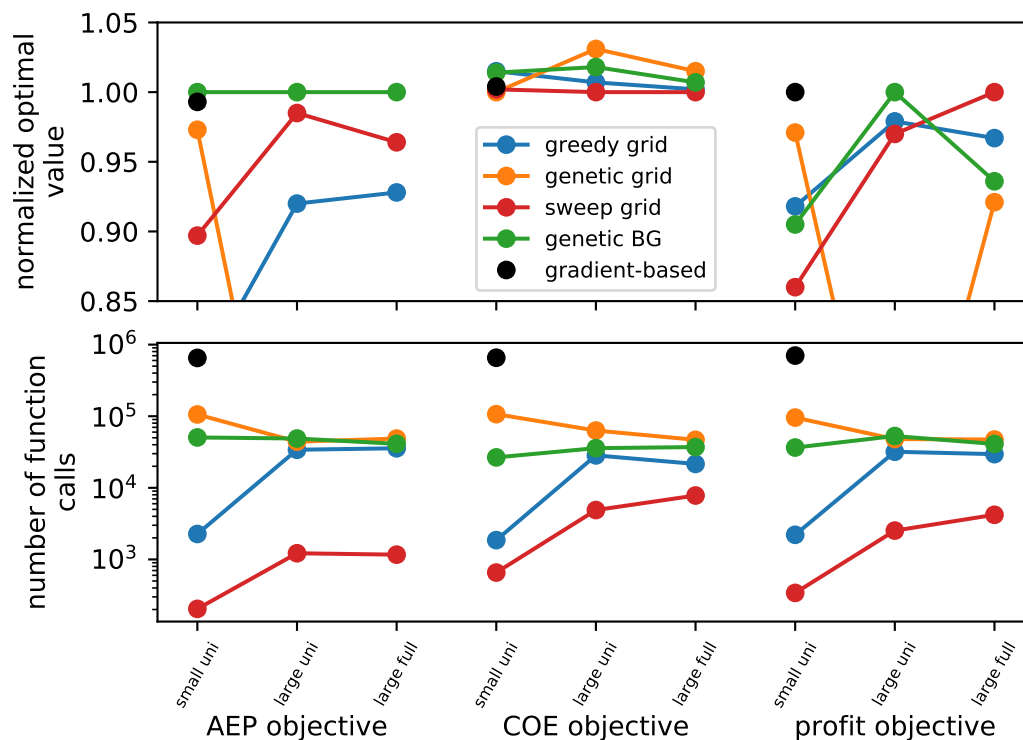


Figure 13. Overall performance representations of the different optimization algorithms. The different colors represent the different optimization methods. The top figure shows the best solution quality, and the bottom figure shows the number of function calls required.



The greedy algorithm performed poorly for the small wind power plant and better for the larger wind plant, although it never performed the best. The computational expense for this algorithm was very low for the small wind plant, but greatly increased for the large wind plant. We can conclude that the computational expense of the greedy algorithm scales poorly with increasing design variables.

The genetic algorithm with the gridded plant domain performed very well for the small wind plant and was close to the best solution for each objective. However, for the larger wind plants, this optimization method performed very poorly. We want to be absolutely clear that part of the reason this optimization method performed so poorly was because we kept the convergence criteria the same for the small and large wind plants. Individual parameters could be altered to get better performance with the gridded domain and the genetic algorithm; however, it is unlikely that it would perform as well as the other algorithms. Genetic algorithms typically have a hard time finding a good solution with large numbers of design variables.

Like the greedy algorithm, the repeated-sweep algorithm performed comparatively poorly for the small wind power plant, but much better for the large wind plants. In fact, this algorithm achieved the best, or close to the best, solution for all of the objectives for each of the large wind plant optimizations. In addition, the computational expense for this algorithm was by far the lowest for the problem sizes included in this paper. That said, similar to the greedy algorithm, the number of function calls required for the repeated-sweep algorithm increased greatly from the small wind plant to the large wind plants. Even though it is computationally efficient for the problem sizes in this paper, as the problem grows, the computational expenses will grow as well and eventually limit the size of wind plant that can be optimized.

The boundary-grid optimization consistently performed the best for the AEP objective. It generally performed poorly for the COE objective, and its performance for the profit objective varied. For the large wind power plant and unidirectional wind rose, the boundary-grid method found the best solution for the profit objective. However, with the full wind rose, the boundary-grid method performed the worst for the profit objective. This is largely because, for our specific case, the full wind rose had a dominant wind direction directly in line with the wind plant boundaries. This is not always the case; for most other cases, the boundary-grid method will perform comparatively well. The primary reason that the boundary-grid method is so effective is that the required function calls do not change with the problem size. While the greedy and repeated-sweep algorithms require many more function calls as the wind plant domain increases in size, the boundary-grid method remained constant between the small and large wind plants. At some point, the boundary-grid method will be more computationally efficient than any of the other algorithms because the number of design variables always remains constant.

6.6 Varied Power Purchase Agreement

The final set of results we present in this paper explores how the optimal profit solution is affected by the PPA. For this section, the problem formulation was identical to the previous section, including using the full wind rose shown in Fig. 10, the wind speed (10 m/s), the wind power plant size (1600-by-1600 m square), grid discretization (20-by-20), and the number of randomly initialized optimizations run (5). However, in this section, instead of assuming a PPA of \$30/MWh, we repeated the optimization while varying the PPA from \$25/MWh–\$100/MWh. Because the repeated-sweep algorithm performed the best for the profit objective in Sec. 6.4, combined with this algorithm's low computational expense, we performed these optimizations



only with the repeated-sweep algorithm and did not compare these with the performance of the other algorithms. The results for these optimizations are shown in Fig. 14. From top to bottom, this figure shows the optimal number of turbines, profit (which was the objective of the optimizations), AEP, and COE as a function of the PPA.




Figure 14. The effect of PPA on the optimal solution for wind power plants optimized for profit. From top to bottom are the number of turbines, profit, AEP, and COE as a function of the PPA.

610 As the value of the energy produced increases (represented by increasing the PPA), the optimal number of turbines in the wind plant also increases. With a low PPA, the optimal solution resembles the optimal COE solution from the previous section.



Wakes are avoided as much as possible, and the number of turbines in the wind plant is low. With a higher PPA, the solution begins to approach the optimal AEP solution. Gains in AEP can be worthwhile with higher PPA—even if they come at the expense of reducing the overall efficiency of each individual turbine.

615 Some typical PPAs in the United States are shown in Fig. 15, including the levelized PPA for various projects since 2010. For the data shown in this figure, the levelized PPA does not take into account any federal tax credit. The different colors represent projects in different parts of the country, which indicates that projects in the central states tend to have lower PPAs. A decade ago, the range of PPAs in the United States was quite large, from around \$40/MWh all the way up to \$120/MWh; however, more recent PPAs are much lower, closer to \$20/MWh–\$40/MWh.




Figure 15. The PPAs for wind power plants built in the United States since 2010. Different colors represent wind plants in different regions of the United States.

620 6.7 Empirical Considerations

The wind power plant layout optimizations in this paper assume a fixed land area in which the wind turbines must remain inside. This is typically referred to as a land-constrained site and empirically represents a scenario where terrain, resource availability, social and siting considerations, or other factors limit the amount of land that is available for wind turbine installations. These factors, combined with installed project cost, project size, power purchase prices, and others, all influence capacity density.

625 As shown in Fig. 14, as PPA price increases for a constant project area, so does the optimal number of turbines. At these high turbine densities, the AEP gain from adding additional turbines is minimal because of large wake losses. The minimal



gains of AEP can justify the extra costs in these scenarios because the value of the small amount of additional electricity is so high. The optimal solutions that we found for our results greatly depend on the assumptions and models we used.

Another driving factor that can determine the capacity density of a given wind plant is market competition, which was not considered in this analysis. Wind plant developers and owners have target profitability values that they wish to achieve. In any energy market, the competition between developers will result in a range of profitability values that results in limited variation in capacity for a specific location. Variations in PPA prices for actual projects are typically due to variations in resource and cost.

Although there are many sites in the United States that are capacity constrained (limited by the capacity of the transmission interconnection), the wind plants optimized for this paper assume land-constrained sites, which are also found throughout the United States. Examples of land-constrained sites in the United States with high turbine density are New England and San Gorgonio, California. The high turbine densities are driven by high PPAs in New England and exceptional wind resources and high-capacity transmission in San Gorgonio. Projects in New England are typically built on ridgelines or other terrain features that limit siting, while those in the San Gorgonio Pass have limited land availability because of terrain and the concentrated wind resource. Many of these wind plants have capacity densities between 5 MW/km² and 10 MW/km² compared to regional averages in the interior United States and the Great Lakes Region between 1.5 MW/km² and 3 MW/km².

Many of the optimal solutions presented in this paper place wind turbines very close together. As turbine scale and cost, energy prices, incentives, social considerations, land availability, resource, and transmission availability vary in the future, the power density of wind power plants in the United States will vary as well. Increasing turbine capacity while maintaining a constant specific power has been shown in some cases to increase potential capacity density. This can include setbacks from structures and other infrastructure as a function of tip height and constant minimum relative rotor spacing between turbines (Bons et al., 2019). Innovations such as wake steering may allow siting flexibility and may increase power densities and increase energy capture.

7 Conclusions

In this paper, we presented our research on wind power plant layout optimization, including optimizing the number of wind turbines. We specifically discussed the effect of different objective functions on the optimal solution as well as the pros and cons associated with using different problem formulations and optimizers to solve the problem. We explored optimizing several different wind plants for objectives of AEP, COE, and profit. The number of turbines in each optimal solution varied dramatically. The plants optimized for AEP tended to have the most wind turbines, those optimized for COE had the least, and those optimized for profit were somewhere in between. The purpose of this research is not to provide an optimal wind plant layout for a specific wind plant boundary and wind conditions, nor is it to suggest general rules of thumb for designing wind plants with different objective functions. The purpose is to clearly demonstrate that the solution from optimizing a wind plant can be heavily influenced by the objective function, particularly when considering the number of wind turbines as a design variable. Specific solutions and layouts will be determined by the models and problem parameters, but the objective must be carefully



660 chosen to represent the desired outcome because any mathematical optimizer will exploit this objective. From the models we
used in this paper, the optimal number of turbines for a square wind plant with a full wind rose was 15, 24, or 54, depending
on the objective of COE, profit, or AEP, respectively.

The other area that we discussed in this paper is problem formulation and algorithm selection for performing the optimization
of turbine number and wind power plant layout optimization. We also presented a very simple repeated-sweep algorithm that
665 performs well, especially for the larger design spaces and for a full wind rose. For a coarse wind plant discretization, in
this paper we used a 10-by-10 grid and found that a simple genetic algorithm performed extremely well in selecting turbine
number and location, even compared to a gradient-based optimizer. However, for a more finely discretized wind plant, the
genetic algorithm performed poorly. A greedy algorithm and the presented repeated-sweep algorithm performed well for the
COE and profit objectives, particularly with a full wind rose leading to turbines that are spaced farther apart on average. At least
670 for the wind plant sizes that we used in this paper, the computational expense of the greedy and repeated-sweep algorithms
scaled favorably with finer grid refinement. The boundary-grid problem formulation optimized with a simple genetic algorithm
performed well regardless of the size of the wind plant, but it performed comparatively better for the larger wind plant size.
Another benefit of the boundary-grid method is that the number of function calls required to optimize the plant stayed relatively
constant as the size of the wind plant domain changed. As long as the time required for the function calls is reasonable, or
675 can be made reasonable, optimizing with the boundary-grid method will produce a layout that performs well regardless of the
objective.

Code availability. All of the code that was used for this paper is publicly available. The repository for this paper, including the optimizers,
run scripts, objective functions, and figure generation can be found at: <https://github.com/pjstanle/stanley2020-turbine-number>. The FLORIS
wake model can be found at: <https://github.com/NREL/floris>.

680 *Author contributions.* APJS led this research, including developing and testing optimizers, running the optimizations, generating figures,
and writing the paper. OR helped develop ideas, provided data for the cost models, and wrote and provided feedback on portions of the paper.
JK and CB provided ideas and guidance for the research as well as editing and feedback on the writing.

Competing interests. The authors declare no competing interests.

Acknowledgements. This work was authored by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy,
685 LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by the U.S. Department
of Energy Office of Energy Efficiency and Renewable Energy Wind Energy Technologies Office. The views expressed in the article do
not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the



article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.



690 References

- Abkar, M. and Porté-Agel, F.: Influence of atmospheric stability on wind-turbine wakes: A large-eddy simulation study, *Physics of Fluids*, 27, <https://doi.org/10.1063/1.4913695>, 2015.
- Baker, N. F., Stanley, A. P. J., Thomas, J. J., Ning, A., and Dykes, K.: Best Practices for Wake Model and Optimization Algorithm Selection in Wind Farm Layout Optimization, in: *AIAA Scitech 2019 Forum*, San Diego, CA, <https://doi.org/10.2514/6.2019-0540>, 2019.
- 695 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.
- Bons, M., Klessmann, C., Lotz, B., and Tiedemann, S.: Scientific foundation of the advice on distance regulations for onshore wind energy, Tech. rep., Federal Ministry for Economy and Energy, 2019.
- Changshui, Z., Guangdong, H., and Jun, W.: A fast algorithm based on the submodular property for optimization of wind turbine positioning, *Renewable energy*, 36, 2951–2958, 2011.
- 700 Chen, K., Song, M., Zhang, X., and Wang, S.: Wind turbine layout optimization with multiple hub height wind turbines using greedy algorithm, *Renewable Energy*, 96, 676–686, 2016.
- Chen, L. and MacDonald, E.: A system-level cost-of-energy wind farm layout optimization with landowner modeling, *Energy Conversion and Management*, 77, 484–494, 2014.
- 705 Chen, Y., Li, H., Jin, K., and Song, Q.: Wind farm layout optimization using genetic algorithm with different hub height wind turbines, *Energy Conversion and Management*, 70, 56–65, 2013.
- Fleming, P., Ning, A., Gebraad, P., and Dykes, K.: Wind Plant System Engineering through Optimization of Layout and Yaw Control, *Wind Energy*, 19, 329–344, <https://doi.org/10.1002/we.1836>, 2016.
- 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, 2017.
- 710 Grady, S., Hussaini, M., and Abdullah, M. M.: Placement of wind turbines using genetic algorithms, *Renewable energy*, 30, 259–270, 2005.
- Hou, P., Zhu, J., Ma, K., Yang, G., Hu, W., and Chen, Z.: A review of offshore wind farm layout optimization and electrical system design methods, *Journal of Modern Power Systems and Clean Energy*, 7, 975–986, 2019.
- Ituarte-Villarreal, C. M. and Espiritu, J. F.: Optimization of wind turbine placement using a viral based optimization algorithm, *Procedia Computer Science*, 6, 469–474, 2011.
- 715 Katić, I., Højstrup, J., and Jensen, N. O.: A simple model for cluster efficiency, in: *EWEA Conference and Exhibition*, pp. 407–410, Rome, Italy, 1986.
- Key, A., Roberts, O., and Eberle, A.: Scaling trends for balance-of-system costs at land-based wind plants: opportunities for innovations in foundation and erection, *Applied Energy*, in review, 2020.
- 720 Khanali, M., Ahmadzadegan, S., Omid, M., Nasab, F. K., and Chau, K. W.: Optimizing layout of wind farm turbines using genetic algorithms in Tehran province, Iran, *International Journal of Energy and Environmental Engineering*, 9, 399–411, 2018.
- Lazard: Lazard's Levelized Cost of Energy Analysis–Version 12.0, accessed 8/29/2019: <https://www.lazard.com/media/450784/lazards-levelized-cost-of-energy-version-120-vfinal.pdf>, 2018.
- Martins, J. R. R. A. and Ning, A.: *Engineering Design Optimization*, Cambridge University Press, (approved), 2020.
- 725 Moorthy, C. B. and Deshmukh, M.: A new approach to optimise placement of wind turbines using particle swarm optimisation, *International Journal of Sustainable Energy*, 34, 396–405, 2015.



- Mosetti, G., Poloni, C., and Diviacco, B.: Optimization of wind turbine positioning in large windfarms by means of a genetic algorithm, *Journal of Wind Engineering and Industrial Aerodynamics*, 51, 105–116, 1994.
- Niayifar, A. and Porté-Agel, F.: Analytical modeling of wind farms: A new approach for power prediction, *Energies*, 9, 741, 2016.
- 730 NREL: FLORIS. Version 2.2.3, <https://github.com/NREL/floris>, 2020.
- Pérez, B., Mínguez, R., and Guanche, R.: Offshore wind farm layout optimization using mathematical programming techniques, *Renewable energy*, 53, 389–399, 2013.
- Previsic, M.: *Economic Methodology for the Evaluation of Emerging Renewable Technologies*, 2011.
- Razdan, P. and Garrett, P.: Life Cycle Assessment of electricity production from an onshore V136-3.45 MW Wind Plant, 2017.
- 735 Şişbot, S., Turgut, Ö., Tunç, M., and Çamdah, Ü.: Optimal positioning of wind turbines on Gökçeada using multi-objective genetic algorithm, *Wind Energy: An International Journal for Progress and Applications in Wind Power Conversion Technology*, 13, 297–306, 2010.
- Song, M., Chen, K., Zhang, X., and Wang, J.: The lazy greedy algorithm for power optimization of wind turbine positioning on complex terrain, *Energy*, 80, 567–574, 2015.
- 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>, 2019a.
- 740 Stanley, A. P. and Ning, A.: Massive simplification of the wind farm layout optimization problem, *Wind Energy Science*, 4, 663–676, 2019b.
- Stehly, T. J. and Beiter, P. C.: 2018 Cost of Wind Energy Review, Tech. rep., National Renewable Energy Lab.(NREL), Golden, CO (United States), 2020.
- 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.
- 745 U.S. Energy Information Administration: Annual energy outlook 2019 with projections to 2050, Department of Energy, pp. 1–83, 2019.
- Vestas Sustainability: Life Cycle Assessments, <https://www.vestas.com/en/about/sustainability#!lcareports>, accessed: 2020-11-20, 2020.
- 750 Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors: SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python, *Nature Methods*, 17, 261–272, <https://doi.org/10.1038/s41592-019-0686-2>, 2020.
- 755 Wiser, R. H., Bolinger, M., Hoen, B., Millstein, D., Rand, J., Barbose, G. L., Darghouth, N. R., Gorman, W., Jeong, S., Mills, A. D., et al.: Wind Technology Data and Trends: Land-Based Focus, 2020 Update, Tech. rep., Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States), 2020.
- Zergane, S., Smaili, A., and Masson, C.: Optimization of wind turbine placement in a wind farm using a new pseudo-random number generation method, *Renewable Energy*, 125, 166–171, 2018.