A Simplified, Efficient Approach to Hybrid Wind and Solar Plant Site Optimization

Charles Tripp[1], Darice Guittet[1], Jennifer King[1], and Aaron Barker[1]

[1]National Renewable Energy Labortory, 15013 Denver West Parkway , Golden, CO 80401

Correspondence: Charles Tripp (charles.tripp@nrel.gov)

Abstract. Wind plant layout optimization is a difficult, complex problem with a large number of variables and many local minima. Layout optimization only becomes more difficult with the addition of solar generation. In this paper, we propose a parameterized approach to wind and solar hybrid power plant layout optimization that greatly reduces problem dimensionality while guaranteeing that the generated layouts have a desirable regular structure. We argue that the evolution strategies class

5 of derivative-free optimization methods is well-suited to the parameterized hybrid layout problem, and we demonstrate how hard layout constraints (e.g. placement restrictions) can be transformed into soft constraints that are amenable to optimization using evolution strategies. Next, we present experimental results on four test sites, demonstrating the viability, reliability, and effectiveness of the parameterized ES approach for generating optimized hybrid plant layouts. Completing the tool kit for parameterized ES layout generation, we include a brief tutorial describing how the parameterized ES approach can be

10 inspected, understood, and debugged when applied to hybrid plant layouts.

1 Introduction

Deployment of integrated hybrid renewable energy systems (HRES) is expected to increase because of their potential to improve flexibility, resilience, and economics. A hybrid wind-solar plant offers benefits due to resource complementarity and shared permitting, siting, equipment, interconnection, transmission, and transaction costs, but it can be difficult to optimally

15 site given additional constraints (Gorman et al., 2020). The design considerations of the stand-alone wind and solar plant apply to the hybrid plant in addition to those imposed by their colocation, such as sizing and the effect of wind turbine shading on solar energy performance. The turbines' layout, wind conditions, and operations are key to the wind plant's annual energy production (AEP). Losses due to wake effects are a major factor, with (Clifton et al., 2016) showing it to be the loss factor with the highest maximum value and highest variation. Photovoltaic (PV) array design and effective irradiance are important

20 site considerations for solar plants. Irradiance reduction can be estimated by on-site surveys or 3D models and minimized by reducing the ground coverage ratio (GCR), using tracking, and by reducing internal and external shading. Mismatch losses due to different electrical properties in shaded and unshaded portions can be estimated with module current-voltage models and minimized by module design and power electronics (MacAlpine et al., 2013; Bendib et al., 2015; Hanson et al., 2014). Matching components and balance-of-plant equipment with expected operational conditions, such as aging and resource availability,

is an important step in HRES design that is not undertaken by this work but for which the approach and tools provided here could be adapted.

Optimizing a single technology is a challenging task. In particular, wind plant layout optimization has been addressed in recent literature to maximize the power output of a wind plant Herbert-Acero et al. (2014); Padrón et al. (2019); Chen and MacDonald (2014). The wind plant layout problem is difficult to solve due to the high-dimensional nature of the problem – many design variables – and there are many local minima that make using gradient-based optimization methods difficult. International Energy Agency Wind Task 37 has developed reference wind plants to be able to compare wind plant layouts across literature. One study developed a method to reduce the number of design variables to increase the computational efficiency of the wind plant layout optimization problem Stanley and Ning (2019) using a boundary grid method. Other layout optimizations have focused on gradient-based optimization algorithms using analytical gradients and approximations to the model to avoid local minima Thomas and Ning (2018); Stanley et al. (2019). This task becomes even more challenging when addressing multiple technologies at a single site.

Previous work has pointed to the difficulties associated with hybrid system optimization and sizing and has identified several research opportunities in this area (Dykes et al., 2019). Existing literature focuses on challenges with the system objectives, decision variables, and constraints associated with the system (Upadhyay and Sharma, 2014; Haghi et al., 2017). These studies on HRES optimization have covered traditional optimization methods encompassing dynamic programming, mixedinteger linear programming, artificial intelligence methods, hybrid methods, and specifically developed software tools (Musselman et al., 2019; Fischetti and Pisinger, 2018; Gebraad et al., 2017; Ning et al., 2019; Cutler et al., 2017). These efforts have predominantly focused on technology sizing or objectives such as reliability, resilience, or downtime optimization. Additionally, most systems studied in the existing literature are independent HRES rather than grid-connected systems.

In these independent systems, such as microgrids, energy production profile, resilience, and downtime prevention are indeed more important than the raw cost of energy production. At the commercial and utility scale, however, projects are extremely cost-sensitive, and developers will seek small optimizations (on the order of 1-3%) that provide an increase in plant profitability. This paper focuses on utility-scale wind and solar hybrid plants. Specifically, this work focuses on a simplified layout optimization method for hybrid wind-solar plants, optimizing hybrid plant layouts for AEP. The goal of this work is to create a well-performing solution in a computationally efficient manner without requiring model gradients.

1.1 Contributions

In this work, the focus on layout optimization for hybrid wind-solar plant AEP will be shown to reduce wake, turbine flicker, and GCR losses for a variety of locations and site shapes. We present four distinct contributions:

– We propose parameterization as an effective tool to reduce the complexity and dimensionality of the hybrid layout optimization problem.

- We argue that the evolution strategies class of derivative-free optimization methods are well-suited to the parameterized hybrid layout problem, and we demonstrate how hard layout constraints can be transformed through parameterization, projection, and finesse into soft constraints amenable to optimization with ES.

- We present experimental results on four test sites, functioning as a proof of concept that our parameterized ES approach is a viable and reliable method for generating optimized layouts with materially increased AEPs over baseline layouts.

- And we present a mini tutorial on inspecting, understanding, and debugging derivative-free optimization approaches in the context of the layout optimization problem.

The proposed approach shows a viable path for hybrid plant developers to generate spatially efficient, high-performing, and maintainable hybrid plant layouts while using modest computational resources.

2 Hybrid Plant Model

A component-based modeling approach was used to allow for system design with increased fidelity and flexibility than would be employed for typical planning, policy, or sizing optimization. Analysis of such a model's design trade-offs often use advanced optimization algorithms due to search spaces that might be nonconvex or ill-behaved, for which the derivatives might be difficult to derive or cumbersome to evaluate. The derivative-free approach is well-suited for such problems but requires fast objective functions because these approaches typically require the evaluation of many candidate solutions. The hybrid plant model developed in this work extends the wind and solar models of the System Advisor Model (SAM), a techno-economic tool that combines renewable energy technologies with financial models (Blair et al., 2018). SAM estimates AEP for a given system configuration using equivalent-circuit and reduced-order models, databases of component performance, and loss factors at multiple points along the simulation. The financial models can be coupled to the performance models for a variety of ownership structures and markets. The separate wind and solar models are discussed in the following sections. The shadow flicker model includes the shading interaction between the wind turbines and solar panels using a geometric representation of the turbines to generate time series of shaded portions of the site. The shaded portions experience some plane-of-array (POA) irradiance loss relative to the full POA at that time step. These POA values are used in PVMismatch's two-diode equivalent-circuit model to estimate PV string output losses relative to unshaded strings (Mikofski et al., 2018; Chaudhari et al., 2018). The shadow flicker model is simulated for a full year and, to enable fast lookup during the objective evaluation of hybrid AEP, flattened into a table of annual loss factors by location relative to the turbine.

2.1 Wind Plant Model

SAM's wind plant model (Freeman et al., 2014) simulates the performance of a wind plant from the wind resource, turbine specifications, and plant layout. The wind resource data, taken from the Wind Integration National Dataset (WIND) Toolkit (Draxl et al., 2015), is hourly temperature, pressure, wind speed, and wind direction at 80 meters. The turbine used was taken from SAM's turbine library and is 1.5 MW, 77 m in diameter, with max power output at 14 m/s, and it is modeled at a hub height

of 80m. SAM's implementation of the eddy viscosity wake model was used for its simplicity, and the turbulence coefficient was set to 0.1. For each evaluation of the objective function, the wind plant layout is recalculated for simulating the wake losses and wind AEP.

2.2 Solar Plant Model

SAM's simple PV plant model, PVWatts®, simulates solar generation using solar resource and high-level system design inputs, such as size, module type, array type, tilt, azimuth, GCR, and DC-to-AC ratio (Dobos, 2014). The solar resource, taken from the National Solar Radiation DataBase (NSRDB), is hourly global horizontal irradiance, diffuse horizontal irradiance, direct normal irradiance, wind speed, and temperature (Sengupta et al., 2018). PVWatts makes assumptions about the system design and layout rather than modeling specific components. For the layout problem, detailed PV design would not affect the placement of the solar array within the site, whereas effects due to tracking modes and GCR can still be estimated while enjoying faster simulations using PVWatts. For this hybrid layout optimization, we used a single-axis tracking system and fixed all parameters besides GCR and system size.

2.3 Shadow Flicker Model

The shadow flicker model uses the turbine dimensions, the site's latitude and longitude, and a PV module model to create a lookup table, or map, of annual loss multipliers by location relative to the turbine. The shadow of the turbine falling on the xy plane for every time step is calculated from the tower height and width; the turbine's blade length, width, and angle; and the sun elevation and azimuth. The number of blade angles to run per time step is an input to the model and is assumed independent of wind speed. The POA is assumed to be reduced by 0.9 uniformly within the turbine's shadow. To calculate the impact of the reduced irradiance on the PV power output, the model places the turbine among a grid of 10-module PV strings aligned vertically, where the module is the default PVMismatch configuration of 96-cell, 3-string modules with a bypass diode per string. For each step, the power output of each PV string is calculated using the full POA for unshaded modules and the reduced POA for shaded modules. We did not model temperature effects, such as heat transfer with ambient or resistive heating, due to the partial shading. Further, the PV array could experience shading from multiple turbines, and how these shadows overlay across a single string might result in power loss that has a nonlinear relationship with the number of shaded modules; however, the simulation time due to additional turbines and an expanded grid of PV strings would be far too long to use within an optimization loop.

Therefore, we investigated ways to reduce the complexity of the model while preserving the required PV power loss information by exploiting the periodicity inherent to each cell of the inner turbine grid, and comparing a map generated from a full simulation with all nearby turbines and one generated by superimposing the losses from a single turbine. Due to the regularity of the inner turbine grid the shadow cast upon a cell internal enough to the grid is representative of other internal cells, so rather than simulating the shadows from all the turbines, we used a grid of 4x4 turbines to examine the central grid cell. The output of the shadow model is a 2D map of POA loss due to turbine shadow. The output of the PV power loss model is a map of PV power output loss on a 10 module string basis due to the reduced effective irradiance and module mismatch.

120   Figure 1 shows the results of piecing together from that central grid cell the image of the 4x4 turbine grid shadows. In particular, Figure 1a shows clearly how long shadows were cut off due to the 4x4 limitation. As the resolution increases, as shown in Figure 1b and Figure 1c, the effect is less noticeable because the relative weight of the long shadows decrease; similarly, as the simulation length increases, the relative POA loss at sunrise and sunset hours decreases.

(a) 1 angle per hour   (b) 3 angles per hour   (c) 9 angles per hour

Figure 1. Heat map of POA loss for a 4x4 turbine grid generated by repeating the central grid cell with different simulation resolutions. (b) and (c) both run three time steps per hour, with 1 or 3 blade angles per time step. The simulation length for each was a few days.

The central grid cell's shadows are not representative of cells that are near the edge of the inner grid. To further generalize
125   the shadow flicker model, we compared the similarity of shadow and PV power loss maps created by the central grid cell and by adding a single turbine's shadow and PV power loss map for each turbine. Figure 2 shows the shadow and PV power loss maps for an area around a turbine at (0, 0) that is 8 turbine diameters to the north, west, and east and 4 turbine diameters to the south at the latitude $33.209°$, longitude $-108.283°$. Weighted throughout the entire year by POA loss and PV power output loss, respectively, the shadow losses range from the greatest shadow reduction of 40.6% to no shadow reduction (100%) far
130   from the turbine, with an average reduction of 99.95%; whereas the PV power losses range from a reduction of 90.5% to 100%, with an average of 99.97%. Above the POA loss, the additional mismatch losses were found to be minimal. PV power loss at each point is a little less than shadow loss due to averaging across the 10 modules. This result reflects the simple PV module assumptions we made and could change given additional details, such as flicker's effect on temperature and power electronics. Comparison of the central grid cell's shadow and PV power losses to those generated by composing the losses from a single

135 turbine showed good agreement, with an average normalized difference of 0.176% for shadow and 0.466% for PV power. For a given candidate's PV array dimensions and location, the aggregate flicker loss is the sum of losses from each turbine and is multiplied to the PV power output.

Figure 2. Heat map of POA loss (left) and PV power loss (right) for a single turbine during a year at $(33.209°, -108.283°)$, simulated at 15-minute intervals with 12 blade angles per step.

3 Optimization Methodology

3.1 Parameterizing Hybrid Plant Layouts

140 Allowing every dimension of a hybrid plant layout to be optimized as a free variable makes for an extremely high-dimensional optimization problem. The position and type of each turbine and solar module must be chosen along with the configuration of the solar module strings. Additionally, layouts generated in this way likely have irregular designs, which can be undesirable for construction, maintenance, cabling, and other purposes. Here, we propose a parameterization that draws inspiration from recent work done to simplify the layout optimization of wind plants (Stanley and Ning, 2019). We propose a parameterization

145 of hybrid plant layouts that significantly reduces problem dimensionality and constrains the solution space to practical, regular layouts. While projecting the design space into a low-dimensional representation necessarily eliminates many potential layouts, we find that many excellent solutions can be readily discovered within the parameterized search space.

Our parameterization, summarized in Table 1, comprises of 11 dimensions: 5 turbine placement parameters and 6 solar placement parameters. Where reasonable, we chose parameters with $(0, 1)$ or similar bounds and relatively smooth, uniform impacts

150 on site layout. Our proposed turbine placement parameters are similar to those in (Stanley and Ning, 2019), but they have been adapted to both the hybrid layout problem and the derivative-free optimization approach. Boundary spacing and boundary offset determine the placement of turbines along the site boundary. Boundary spacing determines the distance between turbines placed along the site boundary, and it is relative to the minimum turbine spacing, set at 200 m in our experiments. Boundary offset is relative to boundary spacing and determines the rotational phase offset along the boundary when placing boundary

Table 1. Optimization variables used in our problem formulation.

Parameter	Definition	Bounds		Prior	
		Min	Max	μ	σ
Boundary spacing	Relative spacing of turbines along boundary (minimum turbine spacing)$(1 +$ boundary spacing$) =$ spacing	0	100	5	5
Boundary offset	Boundary turbine placement offset as ratio of boundary spacing	0	1	.5	2
Grid angle	Interior turbine grid rotation	0	π	$\pi/2$	π
Grid aspect power	Logarithm of the interior turbine grid aspect ratio $e^{\text{grid aspect power}} = \frac{\text{column spacing}}{\text{row spacing}} =$ grid aspect ratio	-4	4	0	3
Row phase offset	Interior grid turbine row starting offset as a multiple of intra-row spacing	0	1	.2	.5
Solar x position	Relative east-west position of the solar region within the site	0	1	.5	.5
Solar y position	Relative north-south position of the solar region within the site	0	1	.5	.5
Solar aspect power	Logarithm of the solar region's aspect ratio. $e^{\text{solar aspect power}} = \frac{\text{east-west size}}{\text{north-south size}} =$ solar aspect ratio	-4	4	0	3
Solar GCR	Ground coverage ratio of the solar region	.2	.9	.5	.5
Solar southern buffer	Southern solar buffer zone relative to minimum turbine spacing southern buffer length $=$ (solar southern buffer)$(1 +$ minimum turbine spacing$)$	1	10	4	4
Solar x buffer	Eastern and western solar buffer zone relative to minimum turbine spacing Southern buffer length $=$ (solar southern buffer)$(1 +$ minimum turbine spacing$)$	1	10	4	4

155 turbines. Grid angle, grid aspect power, and row phase offset control the angle, aspect ratio (the ratio of intra-row to inter-row spacing), and the placement offset between rows. Perturbing the raw aspect ratio causes relatively small changes to the layout for values larger than 1, moderate changes for values near 1, and large changes for values near 0. We instead optimize over the logarithm of the grid aspect ratio because the log of the ratio gives a smooth, more linear response of the layout to perturbations in its value over its entire domain.

160 For solar placement, we consider layouts with a single, contiguous solar region that is rectangular in shape except when placed against a site boundary. The x- and y-position of the center of the solar region are specified by the solar x-position and solar y-position variables, which range from 0 (along the western and southern bounds of the site, respectively) to 1 (along the eastern and northern bounds). The aspect ratio of the solar region is determined by the solar aspect power variable, which like the grid aspect power variable is the logarithm of the aspect ratio of the east-west and north-south lengths of the solar region.

165 To allow the optimizer to minimize shading and flicker losses, we introduce two buffer zones around the solar region from which the turbines are excluded. The solar southern buffer and solar x buffer specify the additional size of the southern and east-west buffers beyond the minimum setback, 200 m in our case, as a multiple of the minimum setback itself. Finally, the GCR of the arrays within the solar region is included as a design variable.

This parameterization does not specify the size of the solar region nor the spacing of turbines within the inner grid. These two variables are instead determined by performing binary searches to find the least dense layouts that accommodate all nonboundary turbines and all solar modules up to the specified wind and solar capacity constraints. Using a binary search to walk along the constraint boundary increases the layout search efficiency by generating candidate layouts that accommodate the maximum allowed solar and wind capacities given their parameterization. Due to the possible nonconvexity of the site boundary, turbine spacing and solar region size are not generally guaranteed to have monotonic responses to the number of turbines or solar modules, potentially causing a binary search to return suboptimal values; however, we did not encounter any issues in using this approach. Convoluted nonconvex site boundaries might need to be simplified for this approach to work, or a binary search could be replaced with a more robust technique that could reliably handle such conditions.

3.2 Objective Design

In this proof of concept, we choose to simply maximize estimated AEP, subject to separate wind and solar nameplate capacity constraints of 75 MW and 50 MW, respectively. We use up to 50 1.5 MW turbines with a minimum spacing of 200 m between turbines and between turbines and solar modules. Other objectives are possible and could simply be dropped in, such as capacity factor, net present value, payback time, or carbon payback time. Further, more detailed objectives could be considered, including cabling, interconnect, maintenance costs, or budgets And other modifications are possible, such as eliminating capacity constraints and instead simply using net present value as the objective and/or allowing the algorithm to trade off between wind turbines and solar modules to find an optimal mix of the two. Similarly, the objective could be modified to generate layouts that improve existing sites by determining the best locations for additional turbines and/or solar modules.

3.2.1 Soft Constraints

Derivative-free optimization methods generate candidates from generative distributions that can be difficult to adapt to hard constraints, so instead we use two forms of soft constraints to guide candidate generation to feasible layouts. We penalize only infeasible solutions, leaving the AEP objective fully intact within the feasible region. When evaluating infeasible solutions, we project them onto the nearest feasible solution by clamping parameter values to their bounds. Our first penalty is a simple quadratic penalty for parameter values outside their constraint boundaries. A quadratic penalty allows the optimizer to stray somewhat beyond the boundary, but due to the quadratic nature of the penalty, the optimizer is neatly repelled from highly infeasible solutions. The second penalty penalizes layouts for which many parameterizations exist due to interference of the site boundary with the solar region's layout. We penalize layouts with excessive solar buffers that extend beyond the site boundary when a smaller solar buffer would result in the same layout. And we penalize layouts with solar aspect ratios that differ from the actual solar region's aspect ratio or that specify a center of the solar region that does not match the actual center of the solar region (as computed from its axis-aligned rectangular bounds). In these cases, we simply add quadratic penalties for these deviations from the ideal parameterization of a given layout, and we did not find it necessary to carefully tune the relative weights of each penalty to get good performance and generate reasonable candidates.

3.2.2 Objective Function

Combining the AEP estimate with the soft constraint penalties results in Equation 1, the objective function used in our experiments:

$$\underset{\boldsymbol{x}}{\text{maximize}} \quad f(\boldsymbol{x}) = P_{\text{wind}}(\boldsymbol{x}) + P_{\text{solar}}(\boldsymbol{x})$$

$$- \eta_0 \left\| \max(\mathbf{0}, \boldsymbol{x} - \boldsymbol{x}_{\max}) \right\|^2 - \eta_0 \left\| \max(\mathbf{0}, \boldsymbol{x}_{\min} - \boldsymbol{x}) \right\|^2 - \eta_1 \sum_{s \in S} \left\| S_s(\boldsymbol{x}) \right\|^2 \tag{1}$$

where:

$\boldsymbol{x}$ is a column vector comprising of the scalar values from Table 1

$P_{\text{wind}}(\boldsymbol{x})$ yields the AEP of the wind plant in the layout described by $\boldsymbol{x}$

$P_{\text{solar}}(\boldsymbol{x})$ yields the AEP of the solar array in the layout described by $\boldsymbol{x}$

η_0 and η_1 are soft constraint nuisance parameters, set to $\eta_0 = 0.1$ and $\eta_1 = 1.0$ in our experiments

$\boldsymbol{x}_{\max}$ and $\boldsymbol{x}_{\min}$ are column vectors comprising of the minimum and maximum values from Table 1

S is the set of layout-based soft constraints penalty functions as described in Section 3.2.1

and $S_s(\boldsymbol{x})$ returns the amount by which soft constraint s is violated by $\boldsymbol{x}$.

3.3 Optimization Methods

Algorithm 1: The evolution strategies stochastic optimization approach

---

$\mathcal{G} \leftarrow \mathcal{G}_0$ `// Initialize generative distribution`

$\mathbf{c}^* \leftarrow \varnothing$ `// Set current best solution to none`

while *optimizing* **do**

     Draw candidate population, $P = \{\mathbf{c}_1, ..., \mathbf{c}_\lambda\}$, from $\mathcal{G}$

     Obtain the objective value of each candidate $F(\mathbf{c}_i)$

     $\mathbf{c}^* \leftarrow \underset{\{\mathbf{c}^*\} \cup \{\mathbf{c}_1, ..., \mathbf{c}_\lambda\}}{\text{argmax}} F(\mathbf{c}_i)$ `// Update best candidate found`

     Update $\mathcal{G}$ using candidate evaluations $\{(\mathbf{c}_1, F(\mathbf{c}_1)), ..., (\mathbf{c}_\lambda, F(\mathbf{c}_\lambda))\}$

return $\mathbf{c}^*$

Algorithm 1 lists an outline of the evolution strategies approach to stochastic optimization. Evolution strategies is a good fit for the hybrid plant layout optimization problem due to the highly nonconvex objective function, the difficulty in obtaining derivatives, their potentially noninformative nature, and the ability to generate multiple good layouts for consideration. Some evolution strategy implementations simply return the mean or other measures of $\mathcal{G}$. We chose to instead return the best solution found, $\mathbf{c}^*$, which experimentally improved performance over returning the mean in every comparison we tested. We evaluated three ES-based approaches for optimizing hybrid plant layouts.

3.3.1 Random Search

Random search is a straightforward evolution strategy where $\mathcal{G}$ is fixed and never updated. Random search simply generates candidate layouts from a stationary distribution, keeping track of the best-performing layout found so far. When the search is terminated, the best layout is returned. In our experiments with random search, we chose $\mathcal{G} \sim \mathcal{N}(\mu, \mathrm{diag}(\sigma))$. Due to its static

225 nature, random search's efficiency depends on the performance of candidate solutions drawn from $\mathcal{G}$. Unfortunately, this is a difficult task because good prior distributions are those that have a high probability of generating good layouts, but if we knew what parameter values would yield good layouts, we would not need to search for them.

3.3.2 Cross-Entropy Method

The cross-entropy method (CEM) is an evolution strategy that originates from rare event simulation that has been adapted to

230 both discrete and continuous variable optimization problems (Y. Rubinstein, 1997; de Boer et al., 2005). We used the common multivariate Gaussian form of CEM optimization, outlined in Algorithm 3.3 of (de Boer et al., 2005), with $\mathcal{G}_0 \sim \mathcal{N}(\mu, \mathrm{diag}(\sigma))$, a population size of $\lambda = 200$, and a selection proportion of $\gamma = \frac{1}{3}$. These choices mean that 200 candidates were generated on each iteration of Algorithm 1, and $\mathcal{G}$ was set to the maximum likelihood multivariate Gaussian fit of the best-performing (highest valued) 67 of those candidate layouts. CEM is an effective strategy that does a good job of efficiently finding high-

235 performing layouts, but it can be prone to getting stuck in local maxima.

3.3.3 Covariance Matrix Adaptation Evolution Strategy

Covariance matrix adaptation evolution strategy (CMA-ES) (Hansen and Ostermeier, 1996, 1997) is a sophisticated evolution strategy that augments CEM's approach with several techniques to better avoid local minima and to update the covariance matrix of $\mathcal{G}$ in a way that is analogous to the approximation of the inverse Hessian matrix as in quasi-Newton methods, such as

240 BFGS (Fletcher, 2000), BFGS-L (Malouf, 2002; Andrew and Gao, 2007), and Broyden's method (Broyden, 1965). CMA-ES has been extended and enhanced over the years to increase the algorithm's recombination efficiency (Hansen and Ostermeier, 2001), improve the time complexity of the update step (Hansen et al., 2003), increase robustness in the face of multimodal objective functions (Hansen and Kern, 2004), and more. We applied the modern implementation of CMA-ES as described in Appendix C of *A Tutorial on the Cross-Entropy Method* (Hansen, 2016). As with CEM, we use a prior distribution of

245 $\mathcal{G}_0 \sim \mathcal{N}(\mu, \mathrm{diag}(\sigma))$, a population size of $\lambda = 200$, and a selection proportion of $\frac{1}{3}$.

4 Experimental Results

As a proof of concept, we present experimental results generated by applying the proposed hybrid layout optimization approach to the four distinct combinations of two site locations and two site boundaries. We choose two distinct locations in the continental United States with the highest and lowest Pearson correlation coefficient (Pearson and Henrici, 1896) between wind

250 and solar resource, using the resource databases mentioned above. The high-correlation location, in which wind and solar re-
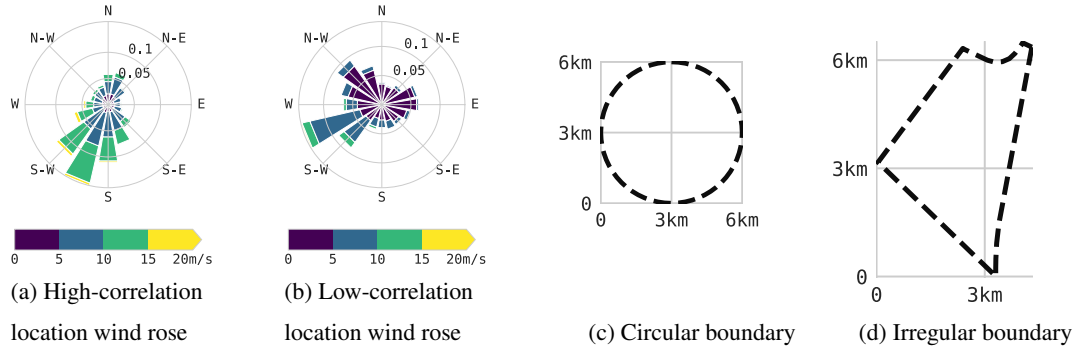
Figure 3. Wind roses and boundaries for the two locations and two boundaries used in our experiments.

sources tend to be present together with a correlation coefficient of 0.28, is located in California's Central Valley, directly south of Fresno and north of the City of Lemoore, at latitude $36.334°$, longitude $-119.769°$, and an elevation of $70m$. Given this moderate positive correlation coefficient, wind and solar resources in even the highest correlation location in the continental United States complement each other somewhat and, therefore likely to yield increased grid resilience and stability through in-
255   creased consistency in energy production. The high-correlation location has a predominant wind direction, as shown in Figure 3a. The low-correlation location is located in Southwest New Mexico, with a latitude of $33.209°$, longitude of $-108.283°$, and an elevation of $2000m$, and it has a resource correlation coefficient of $-.30$. This location presents wind and solar resources that are typically complementary and therefore present an excellent opportunity for hybrid power generation. As shown in Figure 3b, the low-correlation location's wind direction distribution is more dispersed, with lower typical wind speeds than found at the high-correlation location. In this section, we analyze experimental runs of our layout optimization approach as
260   applied to each of the four combinations of these two locations and each of the two site boundaries: a simple circular boundary with a 3 km radius, as shown in Figure 3c, and a nonconvex wedge-shaped irregular boundary, as shown in Figure 3d.

Table 2 summarizes the results of running each of the three evolution strategies optimization algorithms on each combination of the two site locations and two site boundaries. In all cases, CMA-ES achieved a higher mean performance than either CEM or
265   random search, whereas CEM took second place or tied with random search on each site. Taking a closer look, the optimization progress curves shown in Figure 4 indicate that CMA-ES typically achieves good performance with less variability than the other two methods, but CEM sometimes takes an early lead over CMA-ES which CMA-ES overcomes only after 20k–25k

Table 2. The mean performance gains over the baseline site for 10 runs of each optimization algorithm.

	High-Correlation Location		Low-Correlation Location	
Algorithm	Circular Boundary	Irregular Boundary	Circular Boundary	Irregular Boundary
CMA-ES	**3.97%**	**4.30%**	**4.49%**	**3.88%**
CEM	3.45%	4.08%	4.48%	3.86%
Random Search	3.45%	3.89%	4.37%	3.80%

(a) High-correlation site with a circular boundary.

(b) High-correlation site with an irregular boundary.

(c) Low-correlation site with a circular boundary.

(d) Low-correlation site with an irregular boundary.
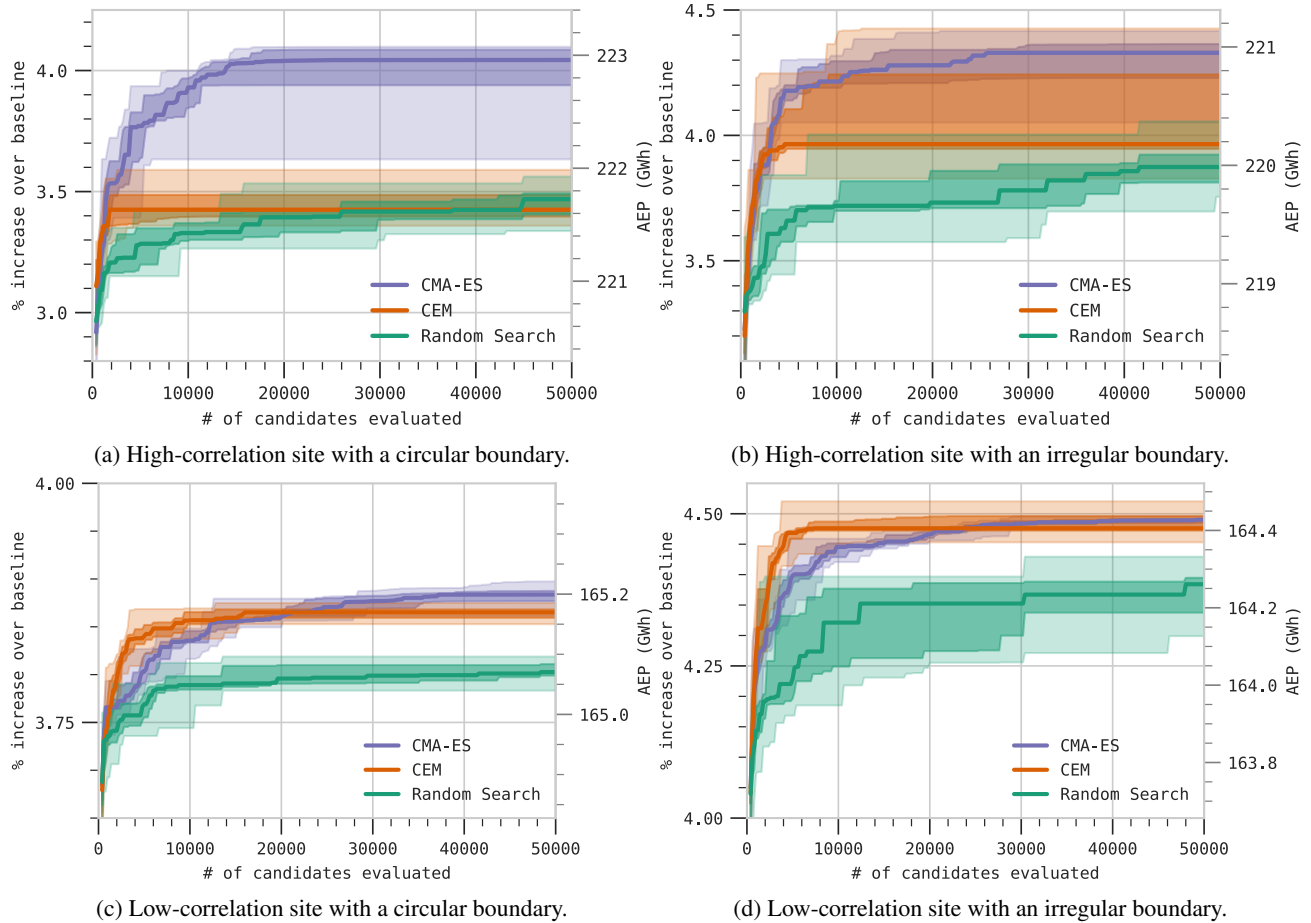
Figure 4. Optimization progress curves for each of the three evolution strategies optimization algorithms on each combination of the two site locations and the two site boundaries over the course of 10 optimization runs. Dark lines indicate median values as observed over 10 optimization runs. The dark fill around the median spans the 25th–75th percentile range, and the lighter fill spans the minimum to maximum range.

candidates have been evaluated. Random search has poor performance overall and tends to have a higher inter-run variability in the performance of its layouts; however, in the high-correlation location with the circular boundary, we see that CEM rapidly

270 becomes stuck in a local maxima, and random search can eventually outperform CEM in this case. In fact, we see that CEM rapidly finds a local maxima in all four cases, and this is likely why CMA-ES, which is more robust to local maxima, is able to eventually beat CEM in these tasks. Interestingly, we find that on every test site, most gains are achieved by the first or second iteration of each of the three algorithms. Among the 200 randomly generated initial candidates, there was always a site that increased the objective value by 2.8 to 4.0%. That is to say that simply drawing one random generation of candidates from

275 the prior distribution and choosing the best-performing layout from that set yielded the most gains to be had when optimizing layouts using this parameterization. It is possible that, in a more general sense, many layouts could be improved significantly

by simply generating a few hundred random perturbations of the layout parameters and choosing the best candidate found; however, in every case, all three optimization algorithms were also able to squeeze out additional performance beyond this initial improvement, with CMA-ES yielding the best overall results. We also see that CMA-ES and random search continue to

280   eek out additional gains between 40k and 50k evaluated candidates, suggesting that longer runs would likely yield additional gains. These results are not meant to be a definitive examination of which approach is best for the hybrid layout problem, but they are instead meant to show that there are viable evolution strategies-based approaches to solving the hybrid layout problem. It is possible that with careful tuning, for example, adjusting CEM's convergence parameters, these results would change somewhat; however, we found that CMA-ES was significantly easier to work with and easier to get running than other

285   techniques, and therefore we examine it in more detail later.

4.1   A Closer Look at the Generated Layouts

Figures 5–8 show a sampling of solution layouts generated by CMA-ES using our hybrid layout parameterization. In Figure 5, the high-correlation location and circular boundary generates a diversity of high-performing layouts. All these layouts pack all or all but one turbine into two inner grid rows, typically aligning turbine rows to an angle at a few degrees offset from the

290   prevailing wind direction. This arrangement minimizes mean wake losses in our eddy viscosity-based wake loss simulation, causing wakes to fall just to the side of downstream turbines under most wind conditions. We also see some solutions, such as the layout shown in Figure 5c, that align the grid closer to perpendicular to the prevailing wind direction. This configuration is also competitive, but the closer spacing between rows in the wind direction results in the southerly turbines incurring a bit more wake losses. Similarly, the solver finds a variety of good solar placements, many of which are nonintuitive, including

295   placements such as those shown in Figure 5c and Figure 5e, which place the solar region along the northern boundary of the site. Despite this northerly placement, the optimizer identified turbine placements that eliminate flicker losses.

Figure 6 shows solutions for the irregular boundary on the same high-correlation location. Unsurprisingly, these solutions share design characteristics with those using a circular boundary, but results differ in a few ways. The "taller" north-south aspect of the irregular boundary causes the optimizer to find solutions that align two and occasionally three rows of turbines with the

300   longer chords of the boundary, again offsetting turbine rows a bit from the prevailing wind direction. Unlike with the circular boundary, some solutions place a smattering of turbines along the site boundary, taking advantage of the additional breathing room afforded by this boundary. In most cases, the solar is packed into the southern tip of the site, eliminating flicker losses entirely; however, a few competitive layouts were found that place the solar region deep in the site's interior, an interesting trade-off that increases turbine spacing at the cost of some flicker and shading losses.

305   The solutions shown in Figure 7 are generated layouts for the low-correlation location and circular boundary. Here, we see that the more uniform and lower speed wind distribution results in very different solutions than at the high-correlation location. In response to a less concentrated wind direction distribution, the solver proposes layouts that space turbines evenly and place the solar region near the site center, giving turbines some additional separation. Similar results are shown in Figure 8 using the irregular boundary, which primarily differ in an increased utilization of boundary turbines, and placement of the solar

310   region into the northeastern corner of the site. These solutions are likely found because placing the solar in this corner actually
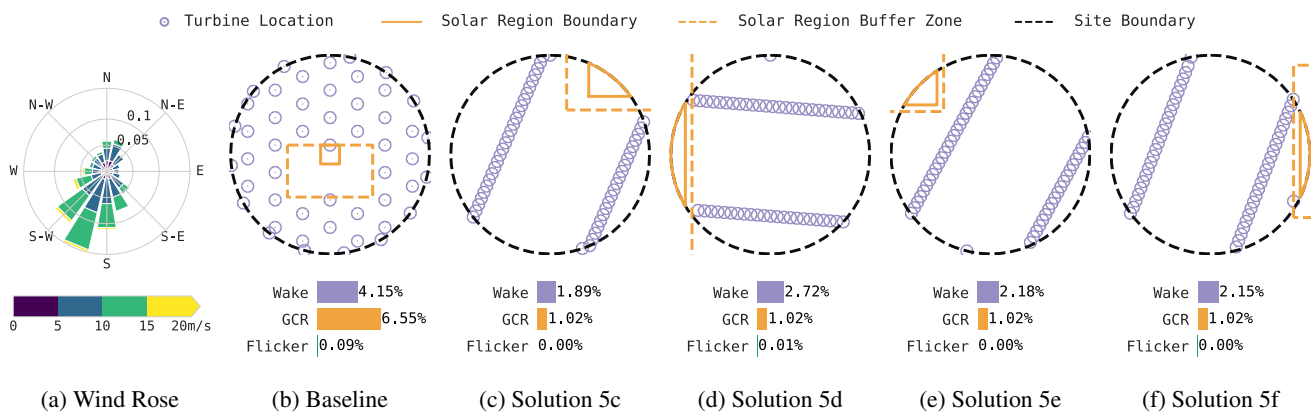
Figure 5. CMA-ES solutions for the high-correlation location and circular boundary. Turbine locations are marked in purple, the solar region is drawn with an orange solid line, and the surrounding solar buffer zone is marked with a dashed orange line.
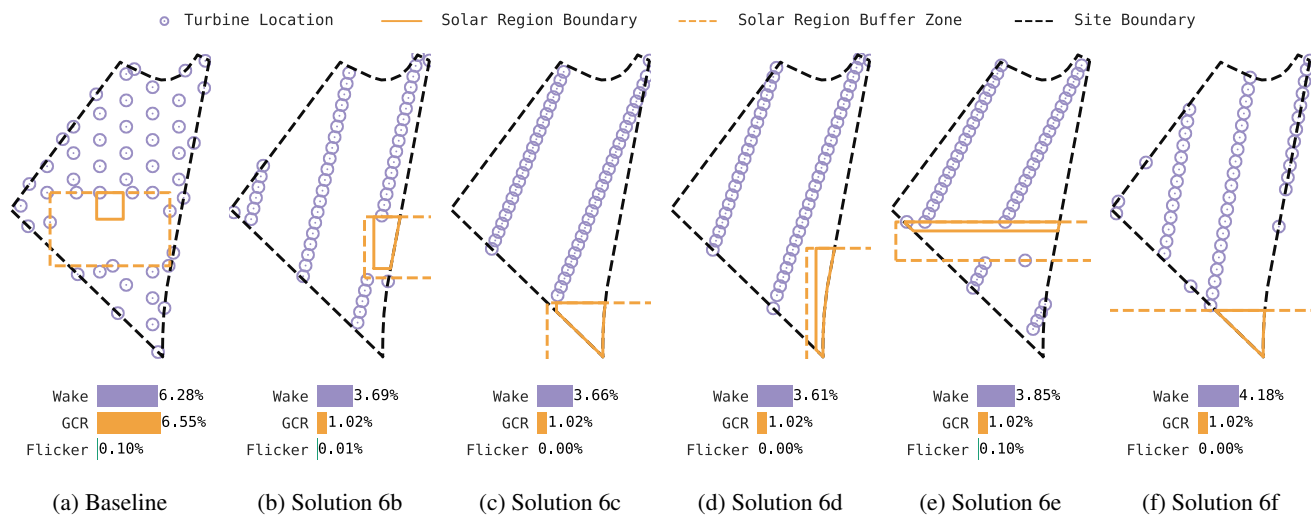
Figure 6. CMA-ES solutions for the high-correlation location and irregular boundary.
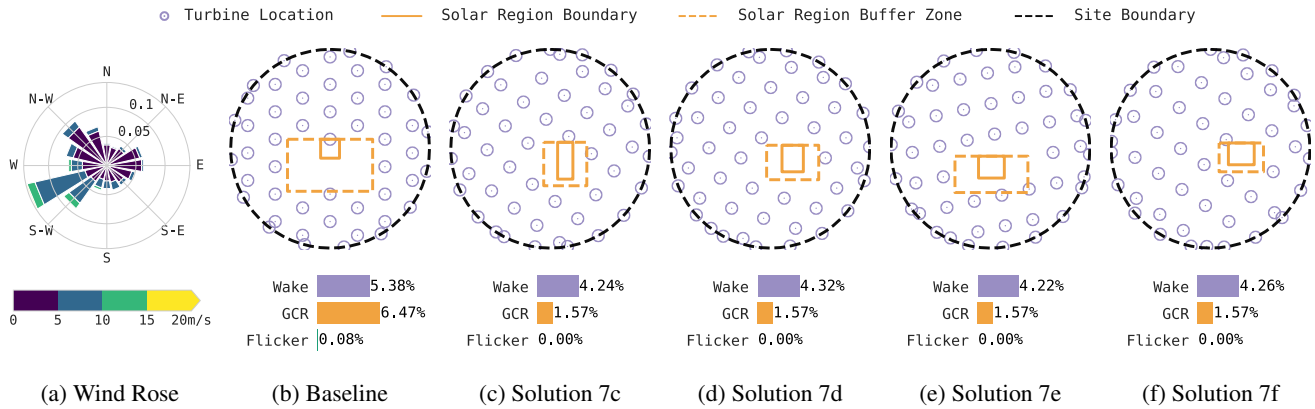
Figure 7. CMA-ES solutions for the low-correlation location and circular boundary.
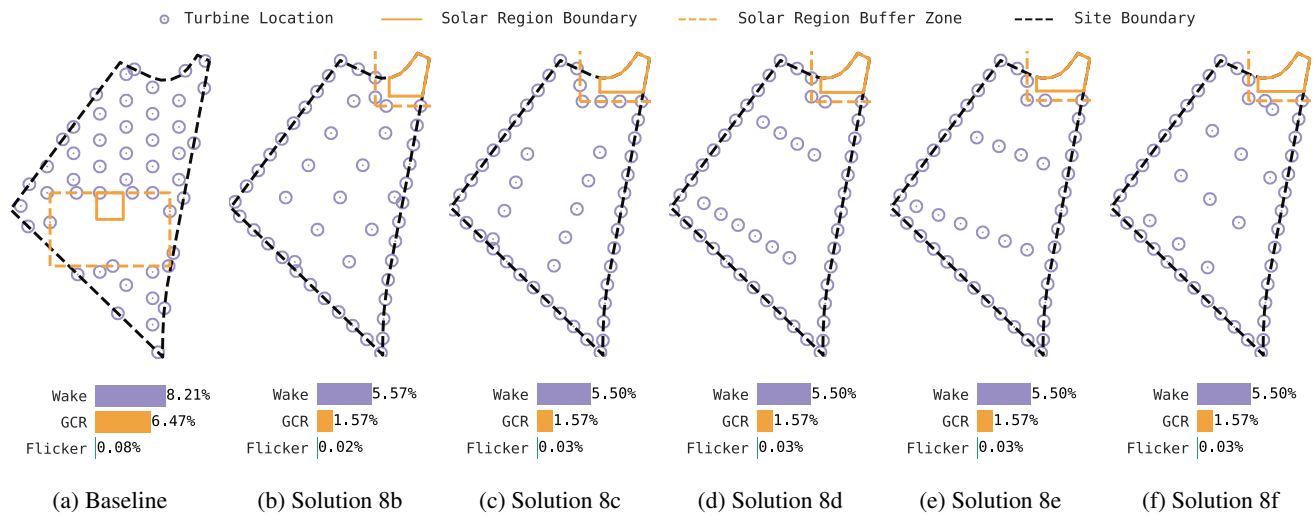
Figure 8. CMA-ES solutions for the low-correlation location and irregular boundary.

causes boundary turbines to avoid the corner and therefore achieve increased spacing. A further-refined parameterization might specially handle border turbine placement in sharp boundary peninsulas such as this one. The ability to generate multiple competitive alternative layouts is a distinct advantage of evolution strategies and other stochastic optimization approaches. Here, we see the creative power of these solution methods in finding a large diversity of viable candidate layouts, all of which yield high objective function scores. In choosing to lay out a hybrid site, one might use these methods to generate a number of good candidate sites, and then choose among them based on other important factors that are difficult to encode in such an objective function, such as ease of access, maintenance or cabling concerns, aesthetics, and more.

(a) Performance trajectories of current-best solutions.   (b) Losses for CMA-ES's solution trajectories.
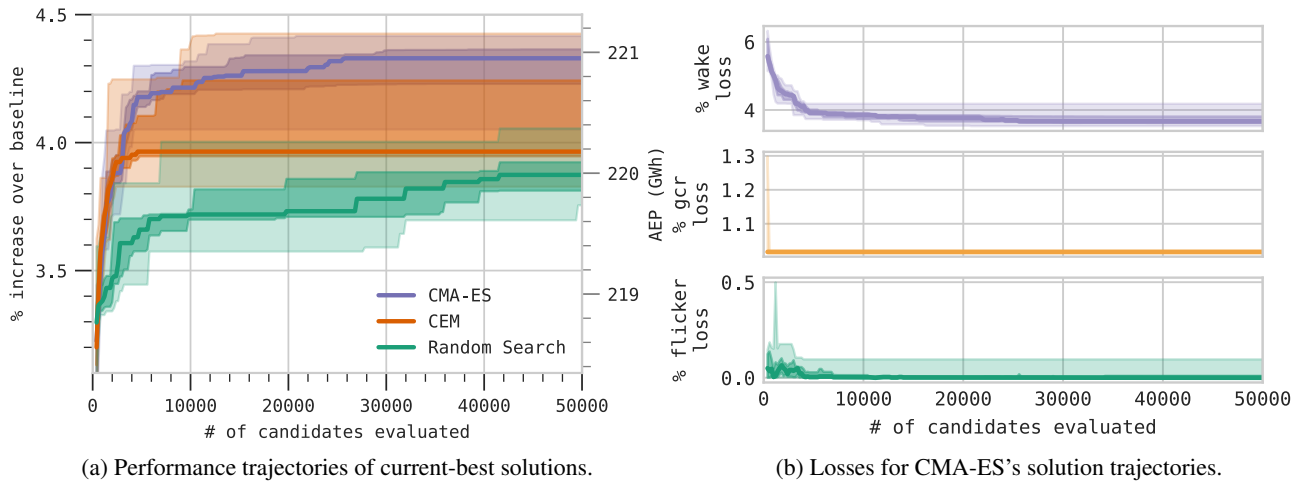
Figure 9. Optimization progress curves on the high-correlation site with an irregular boundary over 10 optimization runs. Dark lines indicate median values as observed over 10 optimization runs. The dark fill around the median spans the 25th–75th percentile range, and the lighter fill spans the minimum to maximum range.

4.2 Peering Into The Black Box: Interpreting and Debugging Derivative-Free Approaches

In this section, we make a case study of the application of CMA-ES to the high-correlation location and irregular boundary layout problem. We show how examining variable trajectories over the optimization run can give insight into the operation of the optimizer, and can help users understand and debug its performance. Graphing solution losses over an optimization run shown in Figure 9b indicates that the random search of the first iteration immediately finds a configuration that minimizes GCR losses; this corresponds to the GCR trajectory shown in Figure 10, where a low GCR is immediately found to minimize GCR losses. Over the next 200 to 800 evaluations, the optimizer concentrates on finding configurations that reduce or eliminate flicker losses. In most cases, the optimizer found solutions that completely eliminated flicker losses. Finally, the optimizer gradually whittles away wake losses. No configuration here can eliminate wake losses, but the optimizer adjusts the turbine grid position, angle, and aspect ratio, significantly reducing wake losses. During this time, we observed a small amount of variability in flicker losses as the optimizer found that it might be able to trade a bit of flicker loss to reduce wake losses, and we see that a handful of solution layouts, such as Figure 6e, make this trade-off.

Analyzing the optimization variables of CMA-ES's solution trajectories shown in Figure 10 indicate that most solutions use a small boundary offset, moderate boundary spacing, a northeasterly turbine grid angle, and tightly packed turbine rows. Solution solar configurations generally had moderate (near-square) aspects, although some such as the layout shown in Figure 6e are wide along the east-west (x) axis and narrow along the north-south (y) axis. Given ample space to place the solar capacity, low GCR's were universally preferred by CMA-ES, minimizing internal shading (GCR) losses. Similarly, the optimizer universally finds that large east and west buffer regions around the solar are not required to reduce flicker. From the x (east-west) and y

(a) Wind parameters

(b) Solar parameters

Figure 10. Solution trajectories using CMA-ES on the high-correlation site and irregular boundary. Dark lines indicate median values. The dark fill around the median spans the 25th–75th percentile range, and the lighter fill spans the minimum to maximum range observed over 10 optimization runs.

(north-south) solar position trajectories, we see that many solutions sensibly pack the solar region into the southeast corner of the site. Placing the solar at the southern end of the site eliminates turbine shading and flicker on the solar, but it can also pack turbines closer together into the northern portion of the site; however, other good placements are found by the optimizer, including as shown in Figure 6e, which places the solar closer to the middle of the site and uses a southern buffer to reduce

340 shading and flicker losses. This alternative arrangement is competitive because it allows the turbines to be spaced farther apart, helping to reduce wake losses. This trade-off would be more salient when using larger solar generation capacities and/or smaller wind capacities, causing the solar region to consume more space relative to wind turbines.

5 Conclusions

In this paper, we presented four distinct contributions. First, the idea of using parameterization to reduce the complexity and
dimensionality of the hybrid layout optimization problem. Parameterization can also restrict generated layouts to those with
certain desirable characteristics, such as a regular inner turbine grid and a consolidated solar patch. And we proposed a particu-
lar parameterization for the hybrid wind and solar plant problem. We proposed this parameterization as a starting point that can
be extended and adapted to meet the needs of different decision makers, site types, and objectives. Next, we proposed the use of
the evolution strategies class of derivative-free stochastic optimization algorithms to solve the parameterized layout problem.
Evolution strategies are well-suited to this problem due to their nonconvex nature, the difficulty of generating derivatives from
the objective function, high flexibility in the nature of the objective function, and evolution strategies' ability to generate a
variety of high-performing candidate solutions. We proposed an adaptation of hard siting constraints to soft constraints and
parameterized projections that make the problem amenable to derivative-free optimization. Then, we presented experimental
results on four test problems demonstrating the ability of this approach to generate good optimized layouts for a variety of site
types. Finally, we took a closer look at the optimization process and presented a brief tutorial on how to inspect, understand, and
debug the parameterized derivative-free approach to solving the hybrid layout optimization problem. Our proposed approach
shows a viable path for hybrid plant developers to easily generate efficient, maintainable, and aesthetically pleasing layouts
using modest computational resources.

Code availability. The experimental site optimization code used in this paper is available under the BSD-3-Clause open source licence on
GitHub as the NREL Hybrid Optimization and Performance Platform (HOPP) project, https://github.com/NREL/HOPP, Tripp et al. (2020).

Author contributions. Jennifer King supervised and managed the project in which this research was accomplished. Darice Guittet developed
and implemented the objective function, shading model, and flicker models. In doing so, she integrated the optimizer with the System
Advisor Model Blair et al. (2018). Aaron Barker's work to map hybrid power production potential was used to identify sites of interest for
our experimental runs. Charles Tripp programmed the optimizer, adapted the optimization problem to be compatible with derivative-free
optimization approaches, and designed and implemented the parameterization method presented in this paper. Jennifer King, Darice Guittet,
Charles Tripp, and Aaron Barker wrote, edited, and contributed to this manuscript.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. This work was authored [in part] by the National Renewable Energy Laboratory, operated by Alliance for Sustainable
Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by the U.S. Depart-

370 ment 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.

References

375   Andrew, G. and Gao, J.: Scalable training of L 1-regularized log-linear models, in: Proceedings of the 24th international conference on Machine learning, pp. 33–40, 2007.

Bendib, B., Belmili, H., and Krim, F.: A survey of the most used MPPT methods: Conventional and advanced algorithms applied for photovoltaic systems, Renewable and Sustainable Energy Reviews, 45, 637–648, https://doi.org/10.1016/j.rser.2015.02.00, https://ideas.repec.org/a/eee/rensus/v45y2015icp637-648.html, 2015.

380   Blair, N. J., DiOrio, N. A., Freeman, J. M., Gilman, P., Janzou, S., Neises, T. W., and Wagner, M. J.: System Advisor Model (SAM) General Description (Version 2017.9.5), https://doi.org/10.2172/1440404, 2018.

Broyden, C. G.: A class of methods for solving nonlinear simultaneous equations, Mathematics of computation, 19, 577–593, 1965.

Chaudhari, C., Kimball, G. M., Hickey, R., and Bourne, B.: Quantification of System-Level Mismatch Losses using PVMismatch, in: 2018 IEEE 7th World Conference on Photovoltaic Energy Conversion (WCPEC) (A Joint Conference of 45th IEEE PVSC, 28th PVSEC 34th

385   EU PVSEC), pp. 3626–3629, 2018.

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.

Clifton, A., Smith, A., and Fields, M.: Wind Plant Preconstruction Energy Estimates. Current Practice and Opportunities, http://www.nrel.gov/docs/fy16osti/64735.pdf, 2016.

390   Cutler, D., Olis, D., Elgqvist, E., Li, X., Laws, N., DiOrio, N., Walker, A., and Anderson, K.: REopt: A platform for energy system integration and optimization, National Renewable Energy Laboratory, Tech. Rep. NREL/TP-7A40-70022, 2017.

de Boer, P.-T., Kroese, D. P., Mannor, S., and Rubinstein, R. Y.: A Tutorial on the Cross-Entropy Method, Annals of Operations Research, 134, 19–67, https://doi.org/10.1007/s10479-005-5724-z, https://doi.org/10.1007/s10479-005-5724-z, 2005.

Dobos, A. P.: PVWatts Version 5 Manual, https://doi.org/10.2172/1158421, 2014.

395   Draxl, C., Hodge, B., Clifton, A., and McCaa, J.: Overview and Meteorological Validation of the Wind Integration National Dataset Toolkit (Technical Report, NREL/TP-5000-61740), 2015.

Dykes, K., King, J., and Diorio, N.: Research Opportunities in the Physical Design Optimization of Hybrid Power Plants, 4th International Hybrid Power Systems Workshop, https://www.nrel.gov/docs/fy20osti/74115.pdf., 2019.

Fischetti, M. and Pisinger, D.: Optimizing wind farm cable routing considering power losses, European Journal of Operational

400   Research, 270, 917 – 930, https://doi.org/https://doi.org/10.1016/j.ejor.2017.07.061, http://www.sciencedirect.com/science/article/pii/S037722171730704X, 2018.

Fletcher, R.: Practical methods of optimization, John Wiley & Sons, 2000.

Freeman, J., Jorgenson, J., Gilman, P., and Ferguson, T.: Reference Manual for the System Advisor Model's Wind Power Performance Model, 2014.

405   Gebraad, P., Thomas, J. J., Ning, A., Fleming, P., and Dykes, K.: Maximization of the annual energy production of wind power plants by optimization of layout and yaw-based wake control, Wind Energy, 20, 97–107, https://doi.org/10.1002/we.1993, https://onlinelibrary.wiley.com/doi/abs/10.1002/we.1993, 2017.

Gorman, W., Mills, A. D., Bolinger, M., Wiser, R. H., Singhal, N. G., Ela, E., and O'Shaughnessy, E.: Motivations and options for deploying hybrid generator-plus-battery projects within the bulk power system, The Electricity Journal, 33,

410   https://doi.org/https://doi.org/10.1016/j.tej.2020.106739, https://www.sciencedirect.com/science/article/pii/S1040619020300312, 2020.

Haghi, E., Farshidian, B., and Saboohi, Y.: Developing a model for optimal sizing of a small hydropower/PV hybrid system for electrification, 2017 5th IEEE International Conference on Smart Energy Grid Engineering, SEGE 2017, pp. 170–176, https://doi.org/10.1109/SEGE.2017.8052794, 2017.

Hansen, N.: The CMA evolution strategy: A tutorial, arXiv preprint arXiv:1604.00772, 2016.

415 Hansen, N. and Kern, S.: Evaluating the CMA evolution strategy on multimodal test functions, in: International Conference on Parallel Problem Solving from Nature, pp. 282–291, Springer, 2004.

Hansen, N. and Ostermeier, A.: Adapting arbitrary normal mutation distributions in evolution strategies: The covariance matrix adaptation, in: Proceedings of IEEE international conference on evolutionary computation, pp. 312–317, IEEE, 1996.

Hansen, N. and Ostermeier, A.: Convergence properties of evolution strategies with the derandomized covariance matrix adaptation: The (=,
420 Eufit, 97, 650–654, 1997.

Hansen, N. and Ostermeier, A.: Completely derandomized self-adaptation in evolution strategies, Evolutionary computation, 9, 159–195, 2001.

Hansen, N., Müller, S. D., and Koumoutsakos, P.: Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES), Evolutionary computation, 11, 1–18, 2003.

425 Hanson, A. J., Deline, C. A., MacAlpine, S. M., Stauth, J. T., and Sullivan, C. R.: Partial-Shading Assessment of Photovoltaic Installations via Module-Level Monitoring, IEEE Journal of Photovoltaics, 4, 1618–1624, 2014.

Herbert-Acero, J. F., Probst, O., Réthoré, P.-E., Larsen, G. C., and Castillo-Villar, K. K.: A review of methodological approaches for the design and optimization of wind farms, Energies, 7, 6930–7016, 2014.

MacAlpine, S. M., Erickson, R. W., and Brandemuehl, M. J.: Characterization of Power Optimizer Potential to Increase Energy Capture in
430 Photovoltaic Systems Operating Under Nonuniform Conditions, IEEE Transactions on Power Electronics, 28, 2936–2945, 2013.

Malouf, R.: A comparison of algorithms for maximum entropy parameter estimation, in: COLING-02: The 6th Conference on Natural Language Learning 2002 (CoNLL-2002), 2002.

Mikofski, M., Meyers, B., and Chaudhari, C.: PVMismatch Project: https://github.com/SunPower/PVMismatch, 2018.

Musselman, A., Thomas, V. M., Boland, N., and Nazzal, D.: Optimizing wind farm siting to reduce power system impacts of wind variability,
435 Wind Energy, 22, 894–907, https://doi.org/10.1002/we.2328, https://onlinelibrary.wiley.com/doi/abs/10.1002/we.2328, 2019.

Ning, A., Dykes, K., and Quick, J.: Systems engineering and optimization of wind turbines and power plants, pp. 235–292, https://doi.org/10.1049/PBPO125G_ch7, 2019.

Padrón, A. S., Thomas, J., Stanley, A. P., Alonso, J. J., and Ning, A.: Polynomial chaos to efficiently compute the annual energy production in wind farm layout optimization, Wind Energy Science, 4, 211–231, 2019.

440 Pearson, K. and Henrici, O. M. F. E.: VII. Mathematical contributions to the theory of evolution.—III. Regression, heredity, and panmixia, Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character, 187, 253–318, https://doi.org/10.1098/rsta.1896.0007, https://royalsocietypublishing.org/doi/abs/10.1098/rsta.1896.0007, 1896.

Sengupta, M., Xie, Y., Lopez, A., Habte, A., Maclaurin, G., and Shelby, J.: The National Solar Radiation Data Base (NSRDB), Renewable and Sustainable Energy Reviews, 89, https://doi.org/10.1016/j.rser.2018.03.003, 2018.

445 Stanley, A. P., Ning, A., and Dykes, K.: Optimization of turbine design in wind farms with multiple hub heights, using exact analytic gradients and structural constraints, Wind Energy, 22, 605–619, 2019.

Stanley, A. P. J. and Ning, A.: Massive simplification of the wind farm layout optimization problem, Wind Energy Science, 4, 663–676, https://doi.org/10.5194/wes-4-663-2019, https://wes.copernicus.org/articles/4/663/2019/, 2019.

Thomas, J. J. and Ning, A.: A method for reducing multi-modality in the wind farm layout optimization problem, in: Journal of Physics:
450   Conference Series, vol. 1037, p. 042012, 2018.

Tripp, C., Guittet, D., King, J., Barker, A., and Hamilton, B.: Hybrid Optimization and Performance Platform (HOPP),
://github.com/NREL/HOPP, https://doi.org/10.11578/dc.20210326.1, 2020.

Upadhyay, S. and Sharma, M. P.: A review on configurations, control and sizing methodologies of hybrid energy systems, Renewable and
Sustainable Energy Reviews, 38, 47–63, https://doi.org/10.1016/j.rser.2014.05.057, http://dx.doi.org/10.1016/j.rser.2014.05.057, 2014.

455 Y. Rubinstein, R.: Optimization of computer simulation models with rare events, European Journal of Operational Research, 99, 89–112,
https://doi.org/10.1016/S0377-2217(96)00385-2, 1997.