WIND
ENERGY
SCIENCE

A simplified, efficient approach to hybrid wind and solar plant site optimization

Charles Tripp, Darice Guittet, Jennifer King, and Aaron Barker

National Renewable Energy Laboratory (NREL), 15013 Denver West Parkway,
Golden, CO 80401, United States of America

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. Thus far, hybrid power plant optimization research has focused on system sizing. We go beyond sizing and present a practical approach to optimizing the physical layout of a wind–solar hybrid power plant. We argue that the evolution strategy class 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 evolution strategy approach for generating optimized hybrid plant layouts. Completing the tool kit for parameterized layout generation, we include a brief tutorial describing how the parameterized evolutionary approach can be inspected, understood, and debugged when applied to hybrid plant layouts.

1 Introduction

Hybrid power plants (HPPs) utilize multiple electrical generation methods to take advantage of each method's benefits while mitigating drawbacks of each individual method. Deployment of integrated hybrid renewable energy systems (HRESs) is expected to increase because of their potential to improve flexibility, resilience, and economics. The diversity of generation resource, the potential for complementary overbuild, and the importance of forecasting and control of HPP may provide value stacking and risk mitigation for the plant owner, as well as increased dispatch efficiency for the bulk grid. While solar photovoltaic (PV) with battery storage is the most common type of HPP, an increasingly prevalent hybrid combination is the combination of wind and solar. Wind–solar hybrid plants benefit from resource complementarity as well as shared permitting, siting, equipment, interconnection, transmission, and transaction costs. However, it can be difficult to optimally site wind–solar plants due to their higher combined complexity (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 array design and effective irradiance are important 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 alone is a challenging task. In particular, wind plant layout optimization has been addressed in recent literature to maximize the power output, minimize levelized cost of energy, or maximize expected profit (Herbert-Acero et al., 2014; Chen and MacDonald, 2014; Padrón et al., 2019; Nagpal et al., 2021; Croonenbroeck and Hennecke, 2021). The wind plant layout problem is difficult to solve due to the high-dimensional nature of the problem and the abundance of local minima. 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 HPP 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, mixed integer 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). HPP optimization efforts have predominantly focused on technology sizing or objectives such as reliability, resilience, or downtime optimization, making most problems amenable to mixed integer linear programming. In this paper, we go beyond sizing and present an approach to optimize the physical layout of a wind–solar HPP.

Additionally, many systems studied in the existing literature are independent HRESs 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.

This paper is divided into three subsequent sections and one Appendix. Section 2 describes the HPP model used, which considers flicker and shading losses caused by turbines shading solar panels in addition to single-source plant modeling factors such as turbine wake losses. Section 3 describes the layout optimization methodology and includes two distinct contributions. First, in Sect. 3.1 we propose parameterization as an effective tool to reduce the complexity and dimensionality of the hybrid layout optimization problem. Second, in Sect. 3.3 we argue that the evolution strategy (ES) 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. Next, in Sect. 4 we contribute 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 and reduce wake, flicker, and GCR losses over baseline layouts. Next, Sect. 4.1 through 4.3 interpret the results of our optimization and suggest general design principles for wind–solar HPPs under various conditions. In Appendix A we provide an approach to inspect, interpret, and debug derivative-free optimization approaches in the context of the layout optimization problem. And finally, we have made the source code implementation used to generate the optimized layouts presented here freely available, https://github.com/NREL/HOPP (last access: 20 March 2022).

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. In this work, we provide a proof of concept that stochastic optimization of low-dimensional parameterized layouts is an effective method for producing efficient hybrid plant layouts. By interpreting optimized layouts under various scenarios, we reveal possible general principles for wind–solar layout optimization.

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 reduced-order models, databases of component performance, and loss factors at multiple points along the simulation. The PV models in SAM use solar resource data containing irradiance and meteorological time series to estimate the energy production based on the type of PV modules, how many of and in what configuration these modules are electrically connected into strings, the type of PV inverters, and the geometric layout of the rows or trackers used for adjusting the angle of the panels to follow the sun, among other factors. The wind model in SAM estimates the energy production of a plant by simulating a single turbine based on its power curve, computing wake effects due to the turbine layout, and applying various losses. The financial models can be coupled to the performance models for a variety of ownership structures and markets, allowing financial metrics such as net present value (NPV) to be used as objectives over the energy-based objective in this paper. 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. At that time step, the shaded areas of the PV module experience some loss of plane-of-array (POA) irradiance, which quantifies the solar power per square meter on the surface of the array and depends on the sun position, the array orientation, ground surface reflectivity, and shading. The effect of shading on the solar panels also includes mismatch effects, which are caused by solar cells and modules experiencing different conditions resulting in negative impacts to the output of the entire PV module and string of modules, including power dissipation and heating. The shading-adjusted POA values are used in PVMismatch's two-diode equivalent-circuit model, which represents the current–voltage characteristic of a PV cell under variations in temperature and irradiance, to estimate PV performance and losses relative to unshaded strings (Mikofski et al., 2018; Chaudhari et al., 2018). This 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. Figure 1 illustrates an example layout of a simulated wind–solar hybrid plant.

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)
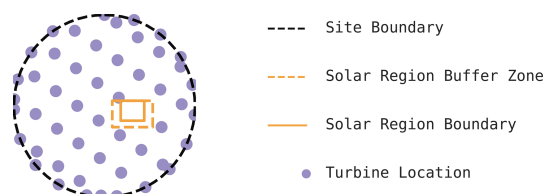
Figure 1. An example layout of a wind–solar hybrid power plant inside a circular site boundary.

Toolkit (Draxl et al., 2015), are hourly temperature, pressure, wind speed, and wind direction at 80 m. The turbine used here was SAM's default selection from the turbine library and is 1.5 MW, 77 m in diameter, with max power output at $14 \, \text{m s}^{-1}$, and it is modeled at a hub height of 80 m. SAM's reduced-order wake model options include a simple model (deficit factor-based), the park model (WAsP), and the eddy viscosity wake model. The Eddy Viscosity model was selected for its relative robustness and accuracy, and the default turbulence coefficient of 0.1 was used. 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 simplifying assumptions about the system and array design rather than modeling specific components. For the layout problem here, this is appropriate because the detailed system, array and electrical parameters available in SAM's detailed PV model would not affect the placement of the solar array within the site, whereas effects due to tracking modes and GCR can still be estimated with the faster simulations of PVWatts. For this hybrid layout optimization, we used a single-axis tracking system, as is found for most utility-scale systems, 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. To calculate the shape of the shadow on the ground from a wind tur-

bine with a given radius, we assume a tower height of $2.5R$, tower width of R; and blade width of $R/16$, following Mamia and Appelbaum (2016). Geometrically, the blade angle is ϕ ° from the z axis, and the turbine yaw angle is θ ° from north, increasing in the clockwise direction. The three blade shadows are calculated from their positions along the parametric equation of a general ellipse, which represents the shadow of the swept area deformed by the yaw and the sun position. The number of blade angles to run per time step is an input to the model and is not calculated from wind speed. The position of the output shadow polygon on the ground is relative to the turbine located at (0, 0). The code for these calculations are found in the shadow_flicker.py source file in the HOPP repository, https://github.com/NREL/HOPP/blob/master/hybrid/layout/shadow_flicker.py (last access: 20 March 2022). The plane-of-array (POA) irradiance 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, three-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 4×4 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.

Figure 2 shows the results of piecing together from that central grid cell the image of the 4×4 turbine grid shadows, on a small demonstration case of only a few days and a few blade angles per hour. In particular, Fig. 2a shows clearly how long shadows were cut off due to the 4×4 limitation. As the resolution increases, as shown in Fig. 2b and c, the effect is less noticeable because the relative weight of the long shadows decreases; similarly, as the simulation length increases, the relative POA loss at sunrise and sunset hours decreases.
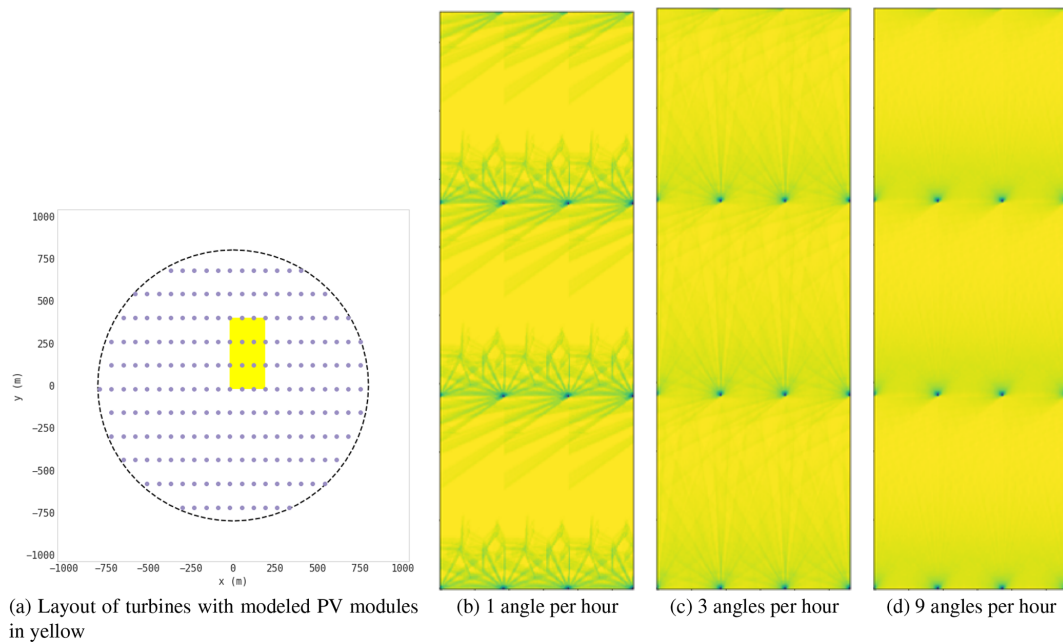
The central grid cell's shadows are not representative of cells that are near the edge of the inner grid. To further generalize the shadow flicker model, we compared the similarity of shadow and PV power loss maps created by the central grid cell and by adding a single turbine's shadow and PV power loss map for each turbine. Figure 3 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 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 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.

3 Optimization methodology

3.1 Parameterizing hybrid plant layouts

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 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.

(a) Layout of turbines with modeled PV modules in yellow   (b) 1 angle per hour   (c) 3 angles per hour   (d) 9 angles per hour

Figure 2. Heat maps **(b–d)** are of POA loss for a 4×4 turbine grid **(a)** generated by repeating the central grid cell with different simulation resolutions; **(b)** and **(c)** both run three time steps per hour, with one or three blade angles per time step. The simulation length for each was a few days.
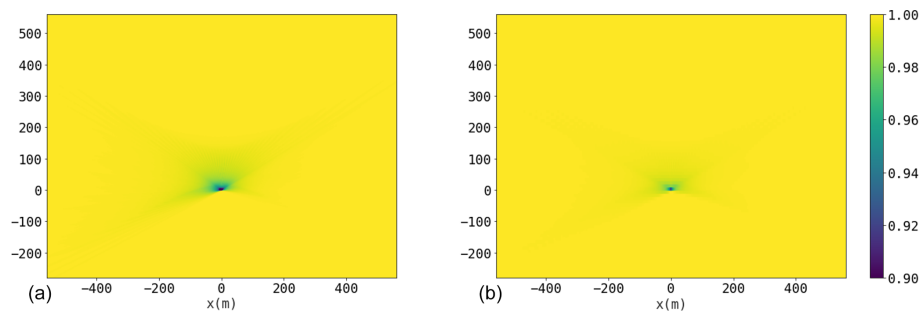
Figure 3. Heat maps of POA loss **(a)** and PV power loss **(b)** for a year at $(33.209°, -108.283°)$, simulated at 15 min intervals with 12 blade angles per step for a 77 m diameter turbine at $(0, 0)$.

Our parameterization, summarized in Table 1, comprises of 11 dimensions: five turbine placement parameters and six solar placement parameters. The exact implementation mapping parameter values to turbine and solar locations, including placement constraint handling can be found beginning on line 176 of the hybrid_parameterization.py source file in the HOPP repository, https://github.com/NREL/HOPP/blob/master/examples/ optimization/layout_opt/hybrid_parametrization.py#L176 (last access: 20 March 2022). Where reasonable, we chose parameters with (0, 1) or similar bounds and relatively smooth, uniform impacts on site layout. 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 relative to the minimum turbine spacing, 200 m in our ex-

periments. Boundary offset determines the rotational phase offset along the boundary when placing boundary 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 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. Therefore, we optimize over the logarithm of the grid aspect ratio, which yields a more linear response. These turbine placement parameters are analogous to those in Stanley and Ning (2019), and we direct the reader there for further description of and justification for these parameters.

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 center point of the

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	0.5	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	0.2	0.5
Solar x position	Relative east–west position of the solar region within the site	0	1	0.5	0.5
Solar y position Solar aspect power	Relative north–south position of the solar region within the site 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	0 -4	1 4	0.5 0	0.5 3
Solar GCR	Ground coverage ratio of the solar region	0.2	0.9	0.5	0.5
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

solar region is 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 is equal to the logarithm of the aspect ratio of the east–west and north–south lengths of the solar region. To allow the optimizer to minimize shading and flicker losses, we define two buffer zones around the solar region from which turbines are excluded. The solar southern buffer and solar x buffer specify the size of the southern and east–west buffers beyond the minimum setback, 200 m in our case, as a multiple of the minimum setback. Finally, the GCR of the arrays within the solar region is included as a design variable. Allowing flexible solar placement beyond the southern boundary of a site (where no shading or flicker losses would be incurred) enables the generation of layouts with interior or solar regions, which may also have little or no shading or flicker losses, but which allow for greater turbine separation and therefore lower wake losses. In fact, many of the optimized layouts discussed in Sect. 4.1 make this trade-off.

This parameterization does not specify the size of the solar region or 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 and 50 MW, respectively. As confirmed in Table 4, these capacity constraints were chosen to yield similar solar and wind AEPs of approximately 110 GWh at the high-correlation location. 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, including capacity factor, net present value, payback time, or carbon payback time. One objective of particular interest for hybrid plants is maximizing utilization of a limited grid interconnect, which can be similarly optimized with this approach in Sect. 4.3.

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. The exact implementation details can be found beginning on line 337 of the hybrid_parameterization.py source file in our repository, https://github.com/NREL/HOPP/blob/7ffb8c58d164ea32f2e0267dbe1869ac6fac9201/examples/optimization/layout_opt/hybrid_parametrization.py#L337 (last access: 20 March 2022).

3.2.2 Objective function

Combining the AEP estimate with the soft constraint penalties results in Eq. (1), the objective function used in our experiments

$$
\begin{aligned}
\underset{\boldsymbol{x}}{\text{maximize}} \; f(\boldsymbol{x}) = \; & P_{\text{wind}}(\boldsymbol{x}) + P_{\text{solar}}(\boldsymbol{x}) \\
& - \eta_0 |\max(0, \boldsymbol{x} - \boldsymbol{x}_{\text{max}})| \\
& - \eta_0 |\max(0, \boldsymbol{x}_{\text{min}} - \boldsymbol{x})| \\
& - \eta_1 \sum_{s \in S} |S_{\text{s}}(\boldsymbol{x})|,
\end{aligned} \tag{1}
$$

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; $\boldsymbol{x}_{\text{max}}$ and $\boldsymbol{x}_{\text{min}}$ are column vectors comprising of the minimum and maximum values from Table 1; S is the set of layout-based soft constraint penalty functions as described in Sect. 3.2.1; and $S_{\text{s}}(\boldsymbol{x})$ returns the amount by which soft constraint s is violated by $\boldsymbol{x}$.

3.3 Optimization methods

Algorithm 1 lists an outline of the evolution strategy (ES) approach to stochastic optimization. Evolution strategies are 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 (RS) is a straightforward evolution strategy where $\mathcal{G}$ is fixed and never updated. RS simply generates candidate layouts from a stationary distribution, keeping track of the best-performing layout found so far; thus, RS provides a simple interpretable baseline. When the search is terminated, the best layout is returned. In our experiments with RS, we chose $\mathcal{G} \sim \mathcal{N}(\mu, \text{diag}(\sigma))$. Due to its static nature, RS'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 both discrete and continuous variable optimization problems (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 $\mathcal{G}_0 \sim \mathcal{N}(\mu, \text{diag}(\sigma))$, 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-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

---

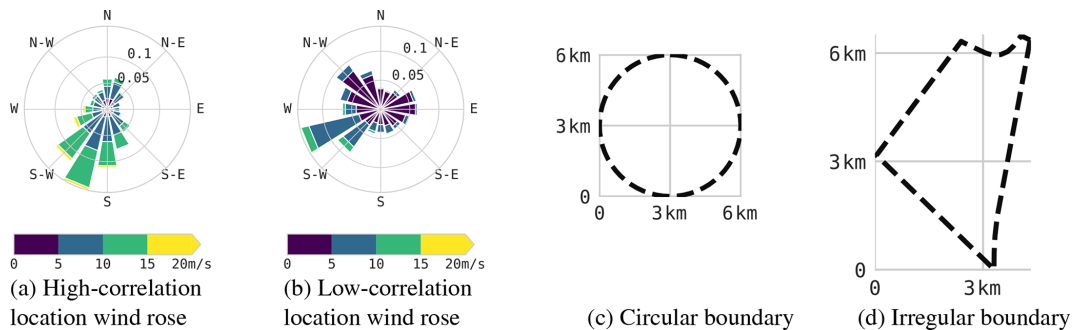**Algorithm 1:** An evolution strategies framework for stochastic optimization.

$\mathcal{G} \leftarrow \mathcal{G}_0$ // 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 $F(\mathbf{c}_i)$
    $\mathbf{c}^* \leftarrow \underset{\{\mathbf{c}^*\} \cup \{\mathbf{c}_1, ..., \mathbf{c}_\lambda\}}{\arg\max} F(\mathbf{c}_i)$ // 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 $\mathbf{c}^*$

(a) High-correlation location wind rose   (b) Low-correlation location wind rose   (c) Circular boundary   (d) Irregular boundary

Figure 4. Wind roses and boundaries for the two locations and two boundaries used in our experiments.

in quasi-Newton methods, such as the Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithm (Fletcher, 2000), the limited memory BFGS algorithm (BFGS-L) (Malouf, 2002; Andrew and Gao, 2007), and the earlier Broyden 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 $\mathcal{G}_0 \sim \mathcal{N}(\boldsymbol{\mu}, \text{diag}(\boldsymbol{\sigma}))$, 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 outlined in Table 2 in the continental United States having the highest and lowest Pearson correlation coefficient (Pearson and Henrici, 1896) between wind and solar resource, using the resource databases mentioned above. We chose to use

the Pearson correlation coefficient because it is the most popular criteria for analyzing the relationship between wind and solar resource (Jurasz et al., 2020; Iwanowski, 2018; Zhang et al., 2013), and unlike Spearman's rank correlation coefficient (Spearman, 1904), it is well-suited to the continuous-valued time-series data used to compare locations. The high-correlation location, in which wind and solar resources 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 70 m. 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 are therefore likely to yield increased grid resilience and stability through increased consistency in energy production. The high-correlation location has a predominant wind direction, as shown in Fig. 4. The low-correlation location is located in southwest New Mexico, with a latitude of 33.209°, longitude of $-108.283°$ and an elevation of 2000 m, and it has a resource correlation coefficient of -0.30. This location presents wind and solar resources that are typically complementary and therefore present an excellent opportunity for hybrid power generation. As shown in Fig. 4b, the low-correlation location's wind direction distribution is more dispersed, with lower typical wind speeds than found at the

Table 2. Comparison of the two test locations.

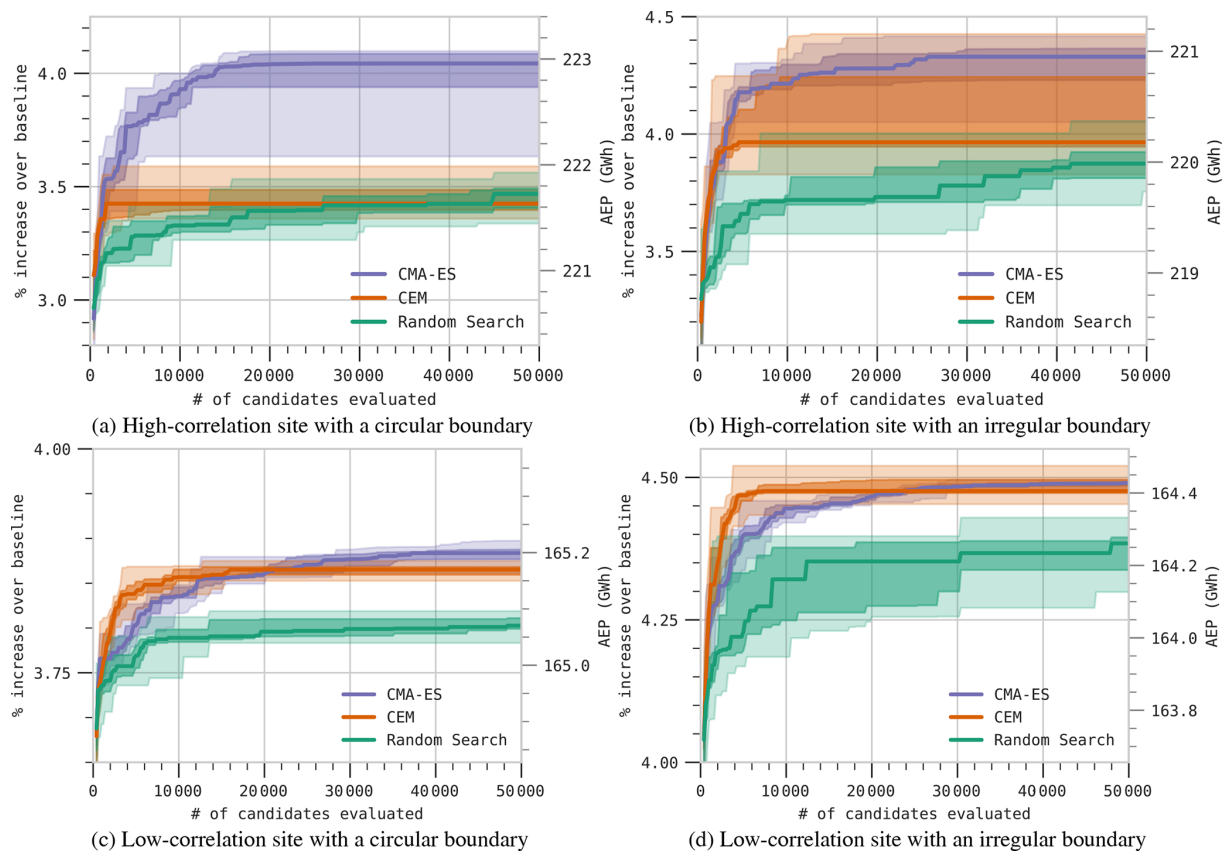| Site	Region	Latitude	Longitude	Elevation	Pearson correlation coefficient
High correlation	Central Valley of California	36.334°	−119.769°	70 m	0.28
Low correlation	Southwest New Mexico	33.209°	−108.283°	2000 m	−0.30

Figure 5. Optimization progress curves for each of the three evolution strategy 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.

high-correlation location. In this section, we analyze experimental runs of our layout optimization approach as 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 Fig. 4c, and a nonconvex wedge-shaped irregular boundary, as shown in Fig. 4d.

Table 3 summarizes the results of running each of the three evolution strategy 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 RS, whereas CEM took second place or tied with RS on each site. Taking a closer look, the optimization progress curves shown in Fig. 5 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 20 000–25 000 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 becomes stuck in a local maximum, and RS can eventually outperform CEM in this case. In fact, we see that CEM rapidly finds a local maximum 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
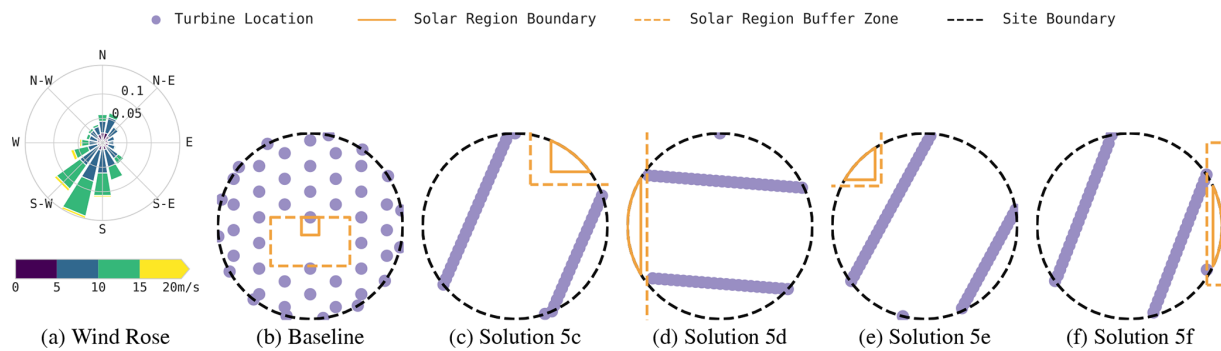
Figure 6. CMA-ES solutions for the high-correlation location and circular boundary. Turbine locations are marked in purple, the solar region is drawn with an orange solid line, and the surrounding solar buffer zone is marked with a dashed orange line.

**Table 3.** Mean performance gains over the baseline site for 10 runs of each optimization algorithm.

Algorithm	High-correlation location		Low-correlation location	
	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 %
RS	3.45 %	3.89 %	4.37 %	3.80 %

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 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 RS continue to eke out additional gains between 40 000 and 50 000 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-strategy-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 techniques, and therefore we examine it in more detail later.

4.1 A closer look at the generated layouts

Figures 6–9 show a sampling of solution layouts generated by CMA-ES using our hybrid layout parameterization. Each layout's performance statistics are listed in Table 4. In Fig. 6, 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 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 Fig. 6c, 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 placements such as those shown in Figs. 6c and e, 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 7 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 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
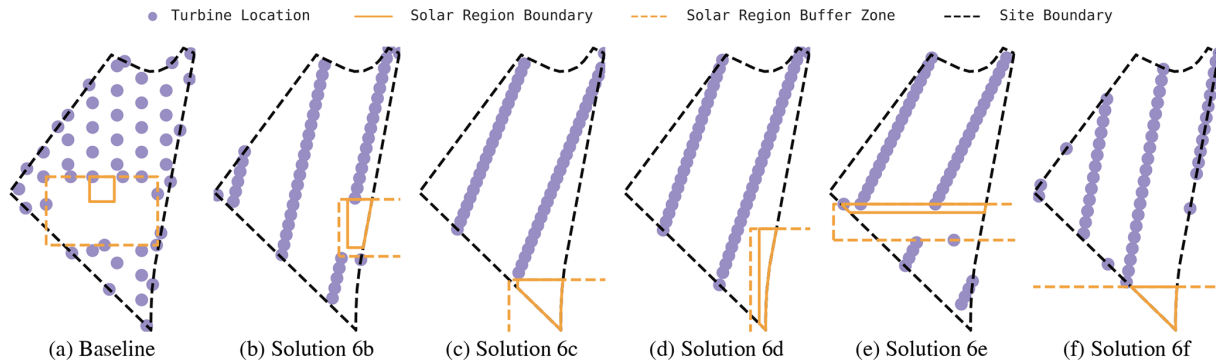
Figure 7. CMA-ES solutions for the high-correlation location and irregular boundary.
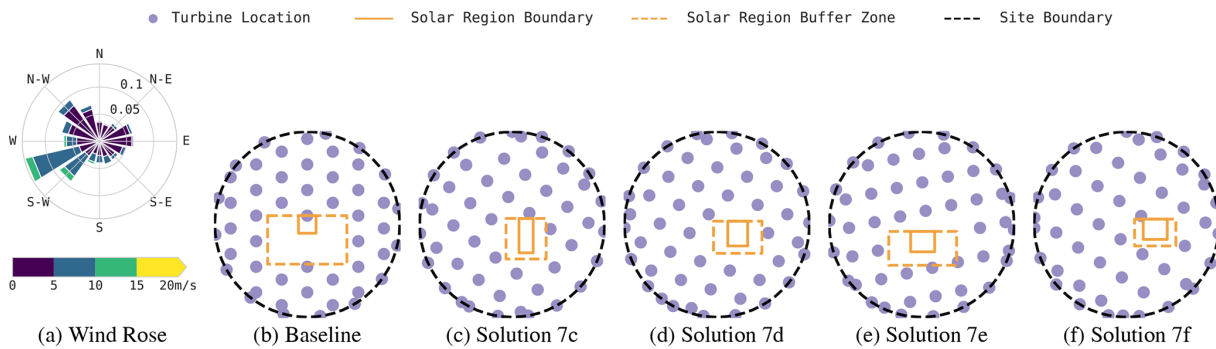
Figure 8. CMA-ES solutions for the low-correlation location and circular boundary.

increases turbine spacing at the cost of some flicker and shading losses.

The solutions shown in Fig. 8 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 Fig. 9 using the irregular boundary, which primarily differ, in an increased utilization of boundary turbines, and placement of the solar region into the northeastern corner of the site. These solutions are likely found because placing the solar in this corner actually 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.

Table 4 reveals that solutions to the high-correlation scenarios have approximately an order of magnitude greater spread in AEP than the low-correlation solutions, and this difference stems almost entirely from differences in wind-generated production. Curiously, the high-correlation sites produce only approximately twice as much wind energy as the low-correlation sites, not nearly enough to explain the much larger difference in AEP. It is possible that higher resource correlation presents a more challenging optimization objective, partly due to the greater impact of flicker losses on solar AEP. It is more important to avoid panel flicker when solar generation is high, and under high-correlation conditions solar generation is high when wind generation is also high, causing shading turbines to inflict greater flicker losses on solar AEP. This proposition is supported somewhat by the overall greater flicker losses seen in high-correlation solutions, but more investigation is needed to fully understand the cause of the variability.

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.
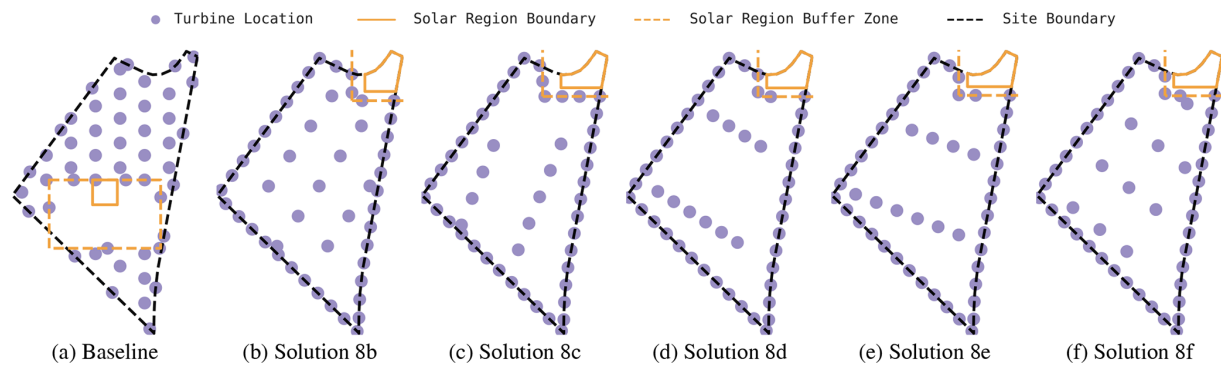
Figure 9. CMA-ES solutions for the low-correlation location and irregular boundary.

Table 4. CMA-ES layout performance statistics for each solution in Figs. 6–9.

| Location	Site	Solution	AEP (GWh)			Losses		
			Total	Solar	Wind	Wake	GCR	Flicker
High correlation	Circular	Baseline	214.29	101.36	112.93	4.15 %	6.55 %	0.09 %
High correlation	Circular	5c	223.06	107.46	115.60	1.89 %	1.02 %	0.00 %
High correlation	Circular	5d	222.08	107.46	114.62	2.72 %	1.02 %	0.01 %
High correlation	Circular	5e	222.73	107.46	115.26	2.18 %	1.02 %	0.00 %
High correlation	Circular	5f	222.76	107.46	115.30	2.15 %	1.02 %	0.00 %
High correlation	Irregular	Baseline	211.78	101.36	110.43	6.28 %	6.55 %	0.10 %
High correlation	Irregular	6b	220.94	107.45	113.48	3.69 %	1.02 %	0.01 %
High correlation	Irregular	6c	220.98	107.46	113.52	3.66 %	1.02 %	0.00 %
High correlation	Irregular	6d	221.04	107.46	113.58	3.61 %	1.02 %	0.00 %
High correlation	Irregular	6e	220.65	107.36	113.29	3.85 %	1.02 %	0.10 %
High correlation	Irregular	6f	220.36	107.46	112.90	4.18 %	1.02 %	0.00 %
Low correlation	Circular	Baseline	159.02	103.57	55.45	5.38 %	6.47 %	0.08 %
Low correlation	Circular	7c	165.21	109.09	56.12	4.24 %	1.57 %	0.00 %
Low correlation	Circular	7d	165.16	109.09	56.07	4.32 %	1.57 %	0.00 %
Low correlation	Circular	7e	165.22	109.09	56.13	4.22 %	1.57 %	0.00 %
Low correlation	Circular	7f	165.20	109.09	56.10	4.26 %	1.57 %	0.00 %
Low correlation	Irregular	Baseline	157.36	103.58	53.79	8.21 %	6.47 %	0.08 %
Low correlation	Irregular	8b	164.40	109.07	55.34	5.57 %	1.57 %	0.02 %
Low correlation	Irregular	8c	164.43	109.06	55.37	5.50 %	1.57 %	0.03 %
Low correlation	Irregular	8d	164.43	109.06	55.37	5.50 %	1.57 %	0.03 %
Low correlation	Irregular	8e	164.44	109.06	55.38	5.50 %	1.57 %	0.03 %
Low correlation	Irregular	8f	164.44	109.06	55.38	5.50 %	1.57 %	0.03 %

4.2 Layouts for varying capacity mixes

Figure 10 shows solutions for various solar to wind generation capacity proportions while holding total capacity equal to 125 MW. For solar-heavy specifications, turbines are placed where they will never shade the solar region and are also spread out to minimize GCR losses, with reducing wake losses being only a secondary concern. Figure 10b is a surprising layout which uses the solar region to position the turbines along two rows in a way which also yields low 2.27 % wake losses for this location. As solar capacity is decreased and wind capacity is increased, the solar region naturally shrinks and is gradually placed to allow for reduced wake losses, with solar losses taking a back seat. Figure 10f shows a primarily wind-based layout with solar stuffed in-between two turbine rows almost as an afterthought. However, even in this case flicker losses are only 0.1 %, and the panels are rarely shaded. These solutions suggest that solar-focused HPPs such as Fig. 10a and b can maintain high solar production by placing panels near the southern boundary of the site, or in a location away from turbine shading, while leaving ample space for low wake loss turbine placement. Sites with balanced production, such as in Fig. 10b–d, are likely well-served by placing solar production along one of the site
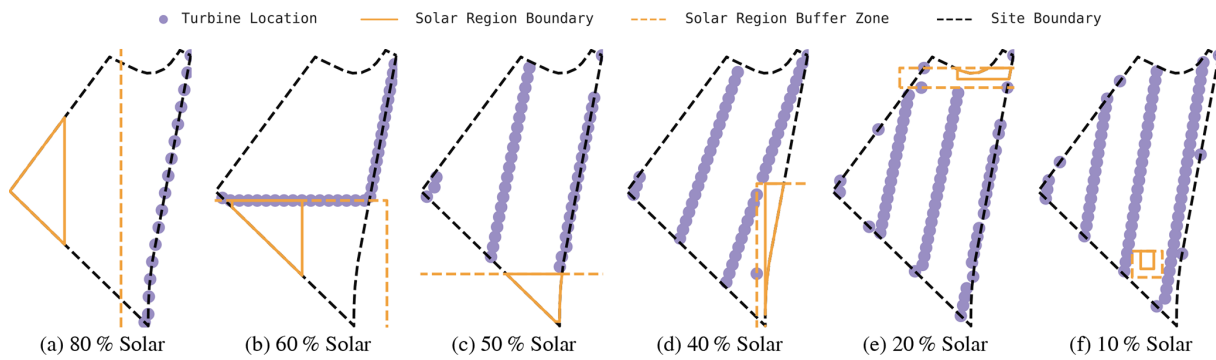
Figure 10. Solutions generated for the high-correlation location and irregular boundary for a range of solar and wind capacity mixes, holding total nameplate capacity at 125 MW.
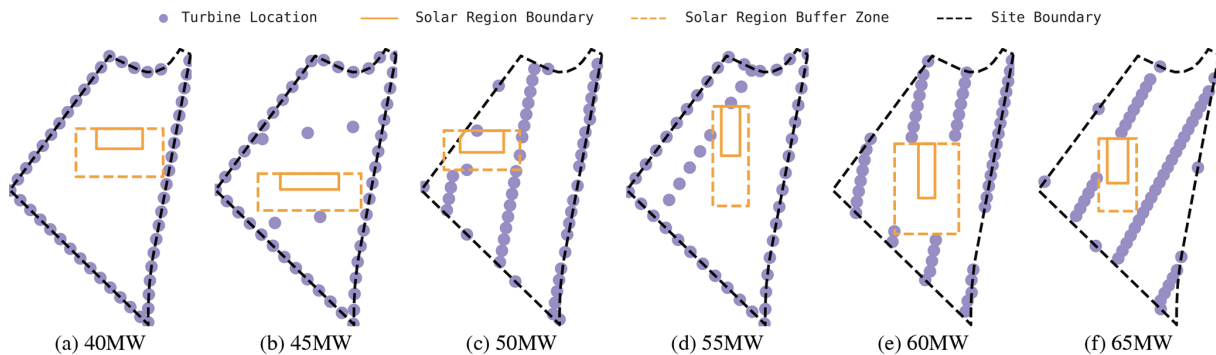
Figure 11. Solutions generated for the high-correlation location and irregular boundary for a range of interconnect capacities, maximizing mean interconnect utilization instead of AEP.

boundaries, particularly the southern boundary, leaving panels largely unshaded while still providing large contiguous spaces for turbine placement. Furthermore, Fig. 10e and f suggest that sites using principally wind generation can often find sufficient space between turbine rows to place modest solar generation zones without incurring significant shading or flicker losses and without moving turbines from their ideal placement. Future work could utilize this layout optimization strategy to, considering the physical layout, identify the mix of solar and wind generation that optimizes figures of interest such as levelized cost of energy or net present value for a particular site.

4.3 Optimizing alternate objectives: interconnect utilization

To evaluate the flexibility of the parameterized layout optimization approach, we generated the layouts shown in Fig. 11, maximizing interconnect utilization instead of AEP for a range of interconnect capacities for high-correlation location and irregular boundary. Because low interconnect capacities do not realize the benefit of peak energy production, the effect of losses during peak production times is unimportant. Therefore, as the interconnect capacity increases, turbines are shifted from the boundary to the interior grid, re-

flecting the increased importance of minimizing wake losses when energy production is high. As the interconnect capacity rises above peak production levels, the optimized layouts for 60 and 65 MW become similar to those found by maximizing AEP in Fig. 7. These results suggest that in highly interconnect-constrained scenarios, such as Fig. 11a, hybrid sites may be best served by placing all turbines along the boundary and solar in a central, unshaded location. This layout type is conducive to maintaining some level of energy production, even in atypical conditions, including rare wind speeds or directions that would cause pathological loss cases under layouts optimized for common conditions. While wake losses will typically be greater under average conditions than with denser rows of turbines, boundary placement, at least with the irregular site, is more robust to atypical conditions. For sites with few straight-line boundaries, such as a rectangular site, it is likely that setting some turbines back from the boundary by varying amounts would produce similar levels of robustness. For sites with moderate interconnect constraints, such as in Fig. 11c and d, interior solar placement is still a successful strategy as long as turbines are excluded from shading positions, and a moderate portion of turbines are placed in interior rows. Such layouts maintain reasonable robustness to atypical conditions while taking some

advantage of common conditions. Future directions include further incorporation of interconnect design parameters into more complex objective functions that would strengthen and deepen these design guidelines.

5   Conclusions

HPP optimization research has focused on system sizing. In this work, we deepen HPP optimization by presenting a practical approach to optimizing not just component sizes but also the physical layout of a wind–solar HPP. Furthermore, this framework can be refined and extended to optimize additional design parameters and achieve more detailed objectives as desired.

The proposed HPP layout optimization approach consists of four distinct contributions. First, we presented a model for estimating shading and flicker losses incurred due to turbine shading of solar panels, a critical piece for enabling wind–solar layout optimization. Second, we proposed utilizing a parametric approach to layout optimization for HPPs in order to reduce the dimensionality of the layout problem and to make it more amenable to non-convex optimization techniques. Third, using a specific parameterization for wind–solar layout optimization, we demonstrated the viability of this approach by using ES-based optimizers to generate high-performance layouts. Finally, we analyzed the optimized layouts under a number of scenarios to propose potential general layout guidelines for wind–solar layout optimization.

Future work includes expanding the parameterization to include additional design parameters such as wind and solar capacity mix, turbine type, and site size and shape; adding more detailed objective functions such as NPV and internal rate of return; and accounting for land use restrictions and costs. Additionally, one could formulate more efficient and capable optimization algorithms, including non-evolutionary approaches. The objective function could be improved to account for factors such as cabling, interconnect, maintenance costs, land use restrictions, and budgets. Other improvements are also possible, including eliminating capacity constraints and allowing the algorithm to trade between wind turbines and solar modules. Similarly, the objective could be modified to generate layouts that improve existing sites by determining the best locations for additional turbines and solar modules. The specific parameterization presented serves as a starting point that can be extended and adapted to meet the needs of different decision makers, site types, and objectives. This approach opens a viable path for hybrid plant developers to easily generate efficient, maintainable, and aesthetically pleasing layouts using modest computational resources.

Appendix A: Peering into the black box: interpreting and debugging derivative-free approaches

In this appendix, 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 Fig. A1b indicates that the RS of the first iteration immediately finds a configuration that minimizes GCR losses; this corresponds to the GCR trajectory shown in Fig. A2, 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 Fig. 7e, make this trade-off.

Analyzing the optimization variables of CMA-ES's solution trajectories shown in Fig. A2 indicates 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 Fig. 7e 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 GCRs 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 (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 Fig. 7e, which places the solar closer to the middle of the site and uses a southern buffer to reduce 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.

(a) Performance trajectories of current-best solutions

(b) Losses for CMA-ES's solution trajectories
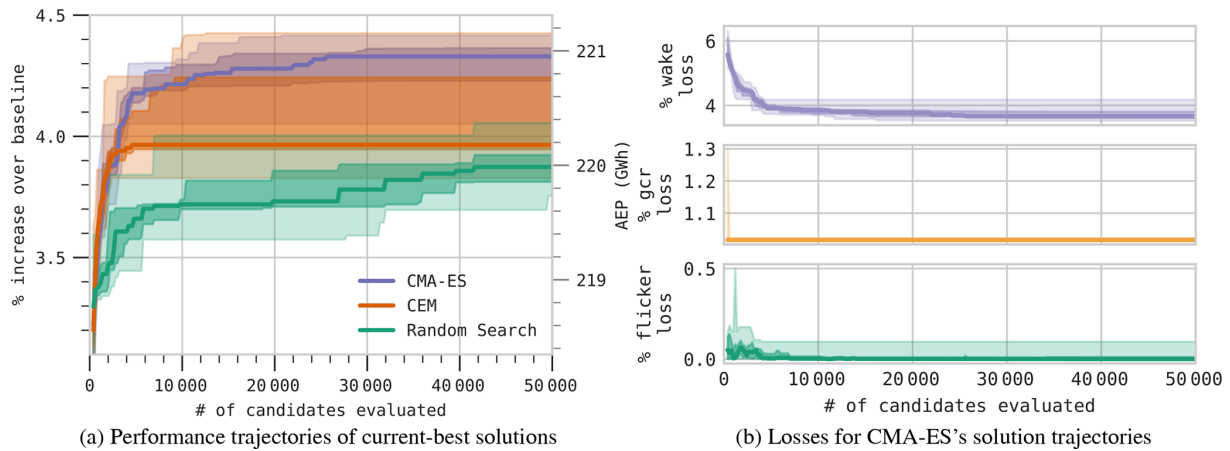
Figure A1. Optimization progress curves on the high-correlation site with an irregular boundary over 10 optimization runs. Dark lines indicate median values as observed over 10 optimization runs. The dark fill around the median spans the 25th–75th percentile range, and the lighter fill spans the minimum to maximum range.

(a) Wind parameters

(b) Solar parameters
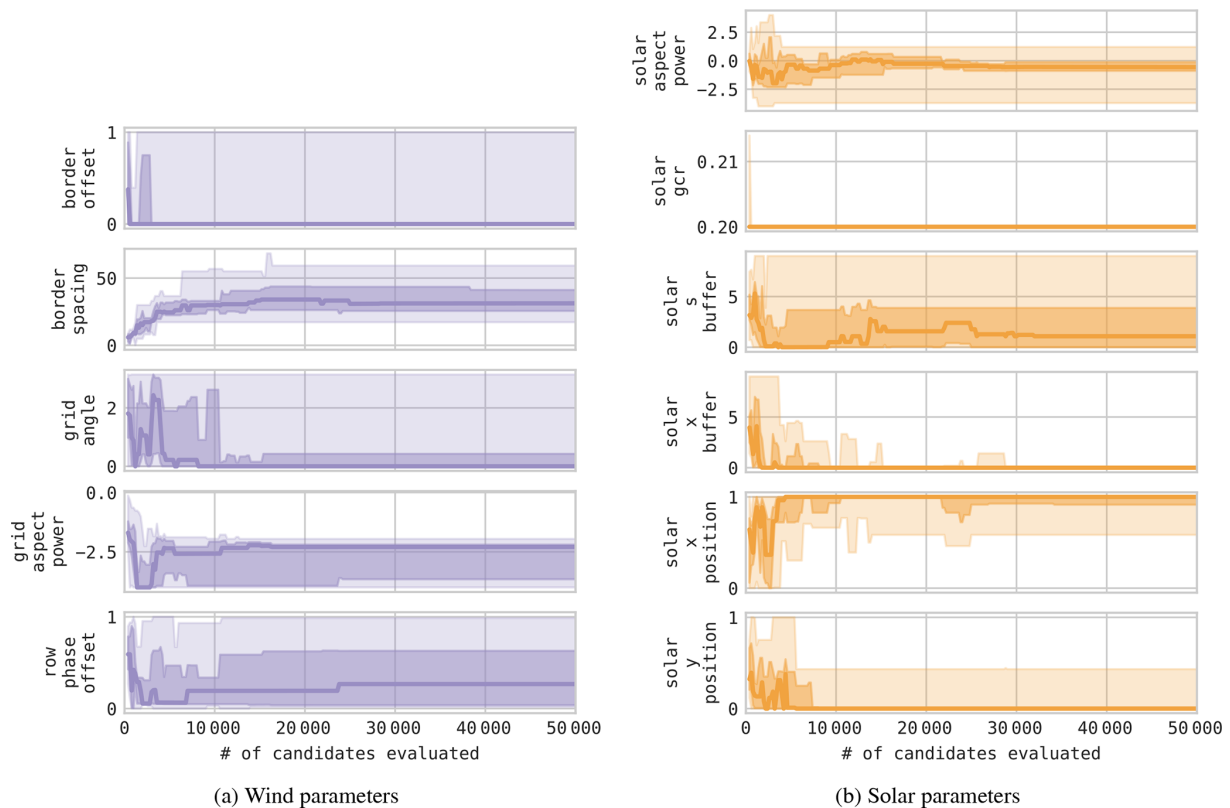
Figure A2. Solution trajectories using CMA-ES on the high-correlation site and irregular boundary. Dark lines indicate median values. The dark fill around the median spans the 25th–75th percentile range, and the lighter fill spans the minimum to maximum range observed over 10 optimization runs.

Author contributions. JK supervised and managed the project in which this research was accomplished. DG 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). AB's work to map hybrid power production potential was used to identify sites of interest for our ex-

perimental runs. CT 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. JK, DG, CT, and AB wrote, edited, and contributed to this paper.

References

Andrew, G. and Gao, J.: Scalable training of L 1-regularized log-linear models, in: Proceedings of the 24th international conference on Machine learning, 33–40, https://doi.org/10.1145/1273496.1273501, 2007.

Bendib, B., Belmili, H., and Krim, F.: A survey of the most used MPPT methods: Conventional and advanced algorithms applied for photovoltaic systems, Renew. Sustain. Energ. Rev., 45, 637–648, https://doi.org/10.1016/j.rser.2015.02.009, 2015.

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), US Department of Energy [code], https://doi.org/10.2172/1440404, 2018.

Broyden, C. G.: A class of methods for solving nonlinear simultaneous equations, Math. Comput., 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 EU PVSEC), 3626–3629, https://doi.org/10.1109/PVSC.2018.8548107, 2018.

Chen, L. and MacDonald, E.: A system-level cost-of-energy wind farm layout optimization with landowner modeling, Energ. Conv. Manage., 77, 484–494, 2014.

Clifton, A., Smith, A., and Fields, M.: Wind Plant Preconstruction Energy Estimates, Current Practice and Opportunities, https://doi.org/10.2172/1248798, 2016.

Croonenbroeck, C. and Hennecke, D.: A comparison of optimizers in a unified standard for optimization on wind farm layout optimization, Energy, 216, 119244, https://doi.org/10.1016/j.energy.2020.119244, 2021.

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, https://doi.org/10.2172/1395453, 2017.

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

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

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, NREL [data set], https://doi.org/10.2172/1214985, 2015.

Dykes, K., King, J., and Diorio, N.: Research Opportunities in the Physical Design Optimization of Hybrid Power Plants, in: 4th International Hybrid Power Systems Workshop, https://www.nrel.gov/docs/fy20osti/74115.pdf (last access: 12 September 2021), 2019.

Fischetti, M. and Pisinger, D.: Optimizing wind farm cable routing considering power losses, Eur. J. Operat. Res., 270, 917–930, https://doi.org/10.1016/j.ejor.2017.07.061, 2018.

Fletcher, R.: Practical methods of optimization, John Wiley & Sons, ISBN 9781118723180, 2000.

Freeman, J., Jorgenson, J., Gilman, P., and Ferguson, T.: Reference Manual for the System Advisor Model's Wind Power Performance Model, https://doi.org/10.2172/1150800, 2014.

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

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, Electricity J., 33, 106739, https://doi.org/10.1016/j.tej.2020.106739, 2020.

Haghi, E., Farshidian, B., and Saboohi, Y.: Developing a model for optimal sizing of a small hydropower/PV hybrid system for electrification, in: 2017 5th IEEE International Conference on Smart Energy Grid Engineering, SEGE 2017, 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.

Hansen, N. and Kern, S.: Evaluating the CMA evolution strategy on multimodal test functions, in: International Conference on Parallel Problem Solving from Nature, Springer, 282–291, https://doi.org/10.1007/978-3-540-30217-9_29, 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, 312–317, https://doi.org/10.1109/ICEC.1996.542381, 1996.

Hansen, N. and Ostermeier, A.: Convergence properties of evolution strategies with the derandomized covariance matrix adaptation: The $(\mu/\mu_l, \lambda)$-CMA-ES, Eufit, 97, 650-654., 1997.

Hansen, N. and Ostermeier, A.: Completely derandomized self-adaptation in evolution strategies, Evolut. Computat., 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), Evolut. Comput., 11, 1–18, 2003.

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 J. Photovolt., 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.

Iwanowski, D.: Correlation Between Wind and Solar Renewable Energy Generation and Electricity Consumption in Poland, Derived to Estimate the Energy Storage Needs, in: IEEE 2018 14th Selected Issues of Electrical Engineering and Electronics (WZEE), 1–4, https://doi.org/10.1109/WZEE.2018.8749030, 2018.

Jurasz, J., Canales, F., Kies, A., Guezgouz, M., and Beluco, A.: A review on the complementarity of renewable energy sources: Concept, metrics, application and future research directions, Solar Energy, 195, 703–724, https://doi.org/10.1016/j.solener.2019.11.087, 2020.

MacAlpine, S. M., Erickson, R. W., and Brandemuehl, M. J.: Characterization of Power Optimizer Potential to Increase Energy Capture in Photovoltaic Systems Operating Under Nonuniform Conditions, IEEE T. Power Electron., 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), https://doi.org/10.3115/1118853.1118871, 2002.

Mamia, I. and Appelbaum, J.: Shadow analysis of wind turbines for dual use of land for combined wind and solar photovoltaic power generation, Renew. Sustain. Energ. Rev., 55, 713–718, https://doi.org/10.1016/j.rser.2015.11.009, 2016.

Mikofski, M., Meyers, B., and Chaudhari, C.: PVMismatch Project, GitHub, https://github.com/SunPower/PVMismatch (last access: 8 December 2021), 2018.

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

Nagpal, S. V., Liu, M. V., and Anderson, C. L.: A comparison of deterministic refinement techniques for wind farm layout optimization, Renew. Energy, 168, 581–592, https://doi.org/10.1016/j.renene.2020.12.043, 2021.

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

Padrón, A. S., Thomas, J., Stanley, A. P. J., Alonso, J. J., and Ning, A.: Polynomial chaos to efficiently compute the annual energy production in wind farm layout optimization, Wind Energ. Sci., 4, 211–231, https://doi.org/10.5194/wes-4-211-2019, 2019.

Pearson, K. and Henrici, O. M. F. E.: VII. Mathematical contributions to the theory of evolution – 2014; III. Regression, heredity, and panmixia, Philos. T. Roy. Soc. Lond. A, 187, 253–318, https://doi.org/10.1098/rsta.1896.0007, 1896.

Rubinstein, R. Y.: Optimization of computer simulation models with rare events, Eur. J. Operat. Res., 99, 89–112, https://doi.org/10.1016/S0377-2217(96)00385-2, 1997.

Sengupta, M., Xie, Y., Lopez, A., Habte, A., Maclaurin, G., and Shelby, J.: The National Solar Radiation Data Base (NSRDB), Renew. Sustain. Energ. Rev., 89, 51–60, https://doi.org/10.1016/j.rser.2018.03.003, 2018.

Spearman, C.: The Proof and Measurement of Association between Two Things, Am. J. Psychol., 15, 72–101, 1904.

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 Energ. Sci., 4, 663–676, https://doi.org/10.5194/wes-4-663-2019, 2019.

Thomas, J. J. and Ning, A.: A method for reducing multi-modality in the wind farm layout optimization problem, J. Phys.: Conf. Ser., 1037, 042012, https://doi.org/10.1088/1742-6596/1037/4/042012, 2018.

Tripp, C., Guittet, D., King, J., Barker, A., and Hamilton, B.: Hybrid Optimization and Performance Platform (HOPP), DOECODE [code], 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, Renew. Sustain. Energ. Rev., 38, 47–63, https://doi.org/10.1016/j.rser.2014.05.057, 2014.

Zhang, J., Hodge, B.-M., and Florita, A.: Investigating the correlation between wind and solar power forecast errors in the western interconnection, in: Energy Sustainability, vol. 55515, American Society of Mechanical Engineers, p. V001T16A003, https://doi.org/10.1115/ES2013-18423, 2013.