

Stochastic Gradient Descent for Wind Farm Optimization

Julian Quick¹, Pierre-Elouan Rethore¹, Mads Mølgaard Pedersen¹, Rafael Valotta Rodrigues¹, and Mikkel Friis-Møller¹

¹Technical University of Denmark, Risø National Laboratory for Sustainable Energy, Frederiksborgvej 399, 4000 Roskilde, Denmark

Correspondence: Julian Quick (juqu@dtu.dk)

Abstract. It is important to optimize wind turbine positions to mitigate potential wake losses. To perform this optimization, atmospheric conditions, such as the inflow speed and direction, are assigned probability distributions according to measured data, which are propagated through engineering wake models to estimate the annual energy production (AEP). This study presents stochastic gradient descent (SGD) for wind farm optimization, which is an approach that estimates the gradient of the AEP using Monte Carlo simulation, allowing for the consideration of an arbitrarily large number of atmospheric conditions. SGD is demonstrated using wind farms with square and circular boundaries, considering cases with 100, 144, 225, and 325 turbines, and the results are compared to a deterministic optimization approach. It is shown that SGD finds a larger optimal AEP in substantially less time than the deterministic counterpart as the number of wind turbines is increased.

1 Introduction

Wind farms are groups of wind turbines that harness the power in the atmospheric boundary layer to provide renewable energy. When a wind turbine absorbs energy from the air, the air downstream of the wind turbine has reduced power, which often reduces the power production of downstream turbines. This is known as the wake effect (Sanderse, 2009; Hasager et al., 2013). When a new wind power plant is to be constructed, optimal turbine locations are determined using engineering wind farm models (Samorani, 2013; Ning. et al., 2020; Annoni et al., 2018). Turbine positions are optimized to exploit the benefits of the local wind resource while avoiding energy losses from turbine wakes. In the wind turbine placement problem, atmospheric conditions, such as the inflow speed and direction, are assigned probability distributions according to measured data. By propagating these probability distributions through the engineering wake model, the annual energy production (AEP) can be estimated. The AEP is often computed using rectangular quadrature, dividing the relevant speeds and directions into equal-sized bins, then computing the expected AEP as the product of the power and probability of each bin, added together, then multiplied by the number of hours per year. The cost of wind farm optimization generally increases with the number of atmospheric conditions considered during AEP computation, and this expensive become more extreme as more complex wake models (e.g., RANS) are considered. For example, there are some memory limitations when computing AEP gradients using automatic differentiation with very large wind farms. This has given rise to studies seeking convergence of the AEP, proposing methods such as polynomial chaos expansion (Padrón et al., 2019; Murcia et al., 2015) or Bayesian quadrature (King et al., 2020), to avoid discretizing the input distributions into evenly spaced intervals. In this study, we present an approach for wind

farm optimization that estimates the gradient of the AEP using Monte Carlo simulation. This does not require that the input be discretized at all, and allows for the consideration of an arbitrarily large number of atmospheric conditions.

Stochastic gradient descent (SGD) is an optimization algorithm commonly used in machine learning when selecting neural network weights (Ketkar, 2017). The algorithm samples the gradient of a stochastic objective, following the mean gradient by a specified distance, then repeating the process, which amounts to optimizing the expected value of the objective. The SGD algorithm is often enhanced to avoid oscillations caused by large changes in the gradient of the objective (Ruder, 2016). This includes methods to reuse previous gradient information (Qian, 1999), dampen oscillations (Riedmiller and Braun, 1993), or incorporate an estimate of the Hessian matrix (Moritz et al., 2016; Byrd et al., 2016; Liu et al., 2018; Najafabadi et al., 2017). Kingma and Ba (2014) introduced the Adam SGD algorithm, which reuses gradient evaluations and dampens oscillations, and is the basis of the SGD method we propose in this study.

Interestingly, SGD is not often applied to problems with nonlinear constraints, although it can be fruitful to include nonlinear constraints in the context of training a machine learning algorithm. For example, when recognizing three-dimensional pictures of people, it can be useful to impose a constraint that any person’s left arm should be close to the same length as their right arm (Márquez-Neila et al., 2017). Many frameworks have been proposed for constrained SGD, including the log-barrier function (Kervadec et al., 2019), penalty functions (Márquez-Neila et al., 2017), blending barrier and penalty functions (Kervadec et al., 2019), and Riemannian geometry (Roy and Harandi, 2017). In this study, we use a penalty term to transform the constrained problem into an unconstrained optimization.

The wind farm layout optimization problem presents a setting where the objective (annual energy production) can be formulated as being stochastic (e.g., the AEP is derived from a probability density function), while the constraints (e.g., boundaries and minimum turbine spacing) are firmly deterministic. This manuscript explores the potential benefits of formulating the wind farm layout optimization problem in this way. As part of this, the Adam algorithm is extended to optimize a stochastic objective with deterministic constraints. To the best of the authors’ knowledge, this exact algorithm has not been published before.

This study benchmarks the performance of the proposed SGD approach when compared to conventional gradient-based optimization within the TOPFARM framework (DTU Wind Energy Systems, 2023b), considering wind farms with different shapes and sizes. We examine the open-source SLSQP algorithm (Kraft, 1988), which is employed in many engineering frameworks (King et al., 2017; Allen et al., 2020; Wu et al., 2020; Zhang et al., 2022; Zilong and Wei, 2022; Kölle et al., 2022; Clark et al., 2022; Simley et al., 2023) and has been used in previous comparisons of optimization algorithms (Lam et al., 2018; Li and Zhang, 2021; Fleming et al., 2022). The TOPFARM framework has been used with SLSQP in several wind farm optimization studies (Riva et al., 2020; Ciavarra et al., 2022; Criado Risco et al., 2023). In future work, this approach can be extended to co-optimize layout and control strategy—the SGD framework can naturally incorporate uncertainty quantification when modeling the potential control strategies for potential layouts [similar to the work in Gebraad et al. (2017); Quick et al. (2020); Howland et al. (2022)].

While there are some wind plant optimization studies that resemble our approach, we are not aware of any studies that have applied SGD to the wind farm optimization problem [although SGD has been applied to other problems in engineering optimization, e.g., De et al. (2020); Sivanantham and Gopalakrishnan (2022)]. Several wind farm optimization studies make

use of gradient-based optimization techniques (Herbert-Acero et al., 2014; Guirguis et al., 2016; Graf et al., 2016; Gebraad et al., 2017; Baker et al., 2019; Riva et al., 2020; Stanley et al., 2021; Croonenbroeck and Hennecke, 2021). Feng and Shen (2015) present a random search approach, moving the wind turbines one by one using a greedy algorithm. Some studies have employed neural networks to forecast power production (Godinho and Castro, 2021), estimate local atmospheric conditions (Stengel et al., 2020), suggest control strategies (Najd et al., 2020), or optimize engineering wake models (Zhang et al., 2021; Zhang and Zhao, 2022; Hussain et al., 2022), which all use SGD algorithms to the train parameters of the neural networks.

The remainder of the manuscript is the following. Section 2 outlines the SGD and deterministic optimization approaches used in this study. Section 3 details the wind farm optimization application cases examined. Section 4 discusses the results of these optimization comparisons. Section 5 provides conclusions and future research directions.

70 2 Methods

When deciding where to put wind turbines, the typical strategy is to maximize wind farm annual energy production (AEP) while ensuring turbines are within the prospective site and are not spaced too closely together. In this study, we examine square and circular wind farms, where the corresponding optimization problems are posed as

$$\begin{aligned}
 & \underset{\mathbf{x}, \mathbf{y}}{\text{maximize}} && \text{AEP}(\mathbf{x}, \mathbf{y}) \\
 & \text{subject to} && (x_i - x_j)^2 + (y_i - y_j)^2 \geq (N_D D)^2, \quad \forall i \neq j \\
 & && x_l \leq x_i \leq x_u \\
 & && y_l \leq y_i \leq y_u
 \end{aligned} \tag{1}$$

75 and

$$\begin{aligned}
 & \underset{\mathbf{x}, \mathbf{y}}{\text{maximize}} && \text{AEP}(\mathbf{x}, \mathbf{y}) \\
 & \text{subject to} && (x_i - x_j)^2 + (y_i - y_j)^2 \geq (N_D D)^2, \quad \forall i \neq j \\
 & && \sqrt{x_i^2 + y_i^2} \leq R,
 \end{aligned} \tag{2}$$

respectively, where \mathbf{x} and \mathbf{y} are the turbine horizontal and vertical locations, x_l and x_u are the lower and upper horizontal square wind farm boundaries, y_l and y_u are the lower and upper square wind farm vertical boundaries, R is the radius of the circular wind farm, D is the rotor diameter, and N_D is the minimum allowable spacing between turbines measured in rotor diameters.

80 From this point forward, we will use a single variable to represent the x- and y-locations, $\mathbf{s} = \{\mathbf{x}, \mathbf{y}\}$. When optimizing square wind farms, x_l , x_u , y_l , and y_u are constant.

The AEP is defined as

$$\text{AEP}(\mathbf{s}) = 8760 \int_0^{2\pi} \int_0^{\infty} P(\mathbf{s}, u_{\infty}, \theta) \pi(u_{\infty}, \theta) du_{\infty} d\theta, \tag{3}$$

where P is power, π is probability, u_∞ is the freestream velocity, and θ is the freestream direction. The 8,760 factor reflects the number of hours per year, converting from units of power to units of energy.

The AEP is typically estimated through rectangular quadrature, where the freestream velocity and direction are discretized using evenly spaced intervals,

$$AEP(\mathbf{s}) \approx 8760 \sum_{d=1}^D \sum_{u=1}^U P(\mathbf{s}, \mathcal{U}_u, \theta_d) \rho(\mathcal{U}_u, \theta_d) \quad (4)$$

where \mathcal{U} is a vector of evenly spaced wind speeds, θ is a vector of evenly spaced wind directions, and $\rho(\mathcal{U}_u, \theta_d)$ is a probability mass function.

The AEP can also be estimated through Monte Carlo integration,

$$AEP(\mathbf{s}) \approx 8760 \frac{1}{K} \sum_{k=1}^K P(\mathbf{s}, u_\infty^{(k)}, \theta^{(k)}), \quad (5)$$

where $u_\infty^{(k)}$ and $\theta^{(k)}$ represent draw k of the probability distribution $\pi(u_\infty, \theta)$.

The associated AEP gradient can also be approximated through Monte Carlo simulation:

$$\frac{d}{d\mathbf{s}} AEP \approx 8760 \frac{1}{K} \sum_{k=1}^K \frac{d}{d\mathbf{s}} P(\mathbf{s}, u_\infty^{(k)}, \theta^{(k)}) \quad (6)$$

2.1 Stochastic Gradient Descent

SGD is built upon the steepest descent algorithm. Early SGD algorithms added a moving average term (sometimes referred to as “momentum”) to avoid spurious oscillations (Tian et al., 2023). The conventional Adam SGD algorithm uses two moving averages: one of the gradient and one of the squared gradient. The ratio of these moving averages is used to determine the search direction. SGD algorithms are often combined with a learning rate scheduler, where the step size of the gradient descent is gradually decreased, allowing the optimization algorithm to hone in on the best solution. While the Adam algorithm is already designed to dynamically change the step size, including a learning rate can further improve the performance. The conventional Adam algorithm is designed for unconstrained optimization algorithms. In the following, we extend the algorithm to allow for deterministic constraints, which is a case that is common in mechanical engineering and unusual in the context of training neural networks. The basic idea is to aggregate the constraints to a penalty term with units that are consistent with the objective. The penalty term is designed so that, initially, the penalty gradients are of similar magnitude to the AEP gradients, and so that the penalty gradients overwhelm the AEP gradients as the optimization continues.

The SGD algorithm is shown in Algorithm 1, where \mathbf{s}_0 are the initial turbine positions, i is the iteration number, α_i is referred to as the constraint multiplier, $\gamma(\mathbf{s})$ is a penalty function, $P(\mathbf{s}, u_\infty^{(k)}, \theta^{(k)})$ is the wind farm power associated with inflow speed $u_\infty^{(k)}$ and $\theta^{(k)}$, K is the number of samples employed in each SGD iteration, β_1 and β_2 are constants, T is the number of SGD iterations, \mathcal{S} is the learning rate scheduler, and η_i is the learning rate. By default, the early stopping option is false.

Algorithm 1 TOPFARM Stochastic Gradient Descent Implementation

```

m ← 0, v ← 0, s ← s0
for i in [0, 1, 2, ..., T - 1]
do
  if early_stopping and  $\eta_i/\eta_0 \leq$  threshold:
    j =  $\alpha_i \frac{\partial \gamma}{\partial \mathbf{s}}$ 
    if |j ≡ 0:
      break
    else:
      j =  $-\frac{8760}{K} \sum_{k=1}^K \frac{\partial}{\partial \mathbf{s}} P(\mathbf{s}, u_\infty^{(k)}, \theta^{(k)}) + \alpha_i \frac{\partial \gamma}{\partial \mathbf{s}}$ 
      m =  $\beta_1 \mathbf{m} - (1 - \beta_1) \mathbf{j}$ 
      v =  $\beta_2 \mathbf{v} - (1 - \beta_2) \mathbf{j}^2$ 
       $\hat{\mathbf{m}} = \frac{\mathbf{m}}{1 - (\beta_1)^i}$ 
       $\hat{\mathbf{v}} = \frac{\mathbf{v}}{1 - (\beta_2)^i}$ 
      s =  $\mathbf{s} - \eta_i \hat{\mathbf{m}} / \sqrt{\hat{\mathbf{v}}}$ 
       $\eta_i = \mathcal{S}(\eta_0, \delta, i)$ 
       $\alpha_i = \alpha_0 \frac{\eta_0}{\eta_i}$ 

```

The spacing between turbines is enforced using a penalty term,

$$\gamma_s = \sum_{\forall i, j > i} \min [(x_i - x_j)^2 + (y_i - y_j)^2 - (N_D D)^2, 0]. \quad (7)$$

Similarly, the distance outside of boundaries is enforced using a penalty term. When considering square wind farms this
 115 penalty term is defined as

$$\gamma_b = \sum_{i=1}^{N_t} \left[\max(x_i - x_{ub}, 0)^2 + \max(x_{lb} - x_i, 0)^2 + \max(y_i - y_{ub}, 0)^2 + \max(y_{lb} - y_i, 0)^2 \right], \quad (8)$$

and, in the case of circular boundaries, it is defined as

$$\gamma_b = \sum_{i=1}^{N_t} \max \left(\sqrt{x_i^2 + y_i^2} - R, 0 \right)^2, \quad (9)$$

where N_t is the number of wind turbines.

120 The total penalty, γ , is defined as the sum of these two penalty terms,

$$\gamma(\mathbf{s}) = \gamma_s(\mathbf{s}) + \gamma_b(\mathbf{s}). \quad (10)$$

The gradient of the penalty term, γ , is scaled before being added to the negative gradient of the AEP using the scaling factor, α_i .

In Algorithm 1, the learning rate (η_i), constraint multiplier (α_i), number of SGD iterations (T), and the samples per SGD iteration (K) are all free parameters. These parameters can be optimized to perform well for individual wind farm optimization problems. But there is no guarantee that these particular parameters will perform well for other wind farm problems—and this meta-optimization can be expensive. In the machine learning community, these parameters are sometimes optimized using evolutionary, grid search, or Bayesian optimization approaches (Alibrahim and Ludwig, 2021). In addition, it is common to schedule the learning rate to decay as the optimization proceeds (You et al., 2019; Denkowski and Neubig, 2017).

We propose a method for setting free parameters to ensure that all units are consistent. The only free parameters we manually set are the number of optimization iterations and the number of power samples per iteration. The optimization generally becomes more accurate and more expensive as these parameters increase, and users are free to balance this tradeoff as they see fit. Our formulation does not guarantee that all intermediate solutions satisfy the constraints, especially in the beginning of the optimization. The constraint multiplier begins on a comparable scale to the AEP, and is scheduled to increase so that the constraint gradients overwhelm the AEP gradients as the optimization progresses. The number of iterations, T , can be based on a prescribed computational budget.

We initially attempted this approach using the widely used default parameter values in the original Adam algorithm, $\beta_1 = 0.9$ and $\beta_2 = 0.999$. The parameters can be thought of as adding momentum to the moving averages of the gradient and squared gradient, \mathbf{m} and \mathbf{v} . We found that these default values gave too much emphasis to gradients from the penalty function, launching the turbines away from the boundaries in a dramatic fashion. Instead, we suggest the parameters $\beta_1 = 0.1$ and $\beta_2 = 0.2$, which encode a shorter memory of the presence of the penalty. With these new default parameters, and the learning rate defined below, we observed successful convergence for a wide variety of test cases.

The learning rate, η_i , can be interpreted as converting $\hat{\mathbf{m}}/\sqrt{\hat{\mathbf{v}}}$ (with unity units) to distance (units of m). In this study, the learning rate is scheduled to decay according to

$$\begin{aligned} S(t=0) &= \eta_0 \\ S(t=T-1) &= \eta_T \end{aligned} \quad , \quad (11)$$

where T is the number of optimization iterations, η_0 is the initial learning rate, and η_T is the scheduled final learning rate. This final learning rate can be thought of as a solution tolerance for the design variables. In this study, we set $\eta_T = 0.1\text{m}$.

The initial learning rate, η_0 , is based on a length scale parameter, L , which corresponds to a reasonable initial step size for the optimization. By setting the initial learning rate according to

$$\eta_0 = L = D/5, \quad (12)$$

where D is the turbine rotor diameter, we encourage the turbines to move at most L every optimization iteration.

The learning rate is scheduled to decay as

$$S(\eta_0, \delta, t) = \eta_0 \prod_{i=0}^t \frac{1}{1 + i\delta}, \quad (13)$$

where δ is a parameter that controls the learning rate length, such that the final learning rate is η_T . The parameter δ is numerically set as

$$\delta(\eta_0, \eta_T, T) = \underset{\delta}{\operatorname{argmin}} |\eta_T - S(\eta_0, \delta, T)|. \quad (14)$$

The constraint multiplier, α_i , can be interpreted as converting the gradient of constrained square distances (in units of m) to AEP gradients. The initial constraint multiplier, α_0 , is set as the mean absolute AEP gradient divided by the length scale, L , so that the separation constraint has a similar scale to AEP gradients,

$$\alpha_0 = \frac{\operatorname{mean}[|\nabla AEP(\mathbf{s}_0)|]}{L}, \quad (15)$$

where $\operatorname{mean}[|\nabla AEP(\mathbf{s}_0)|]$ is the mean of the absolute AEP gradient of the initial guess with respect to each component of the gradient. During each iteration, the constraint multiplier, α_i , is scheduled to increase based on the inverse of the learning rate,

$$\alpha_i = \alpha_0 \frac{\eta_0}{\eta_i}. \quad (16)$$

The wind rose samples, $(u_\infty^{(i)}, \theta^{(i)}) \sim \pi(u_\infty, \theta)$, are randomly selected based on the direction frequency and direction-specific Weibull shape and scale parameters. Note that the tilde (\sim) denotes a shared probability distribution. After a direction is sampled, the wind speed is sampled as a continuous weibull distributed random variable,

$$u_\infty(\theta) \sim \mathcal{W}[u_\infty, A(\theta), k(\theta)], \quad (17)$$

where the probability density of the Weibull distribution, \mathcal{W} , is given by

$$\mathcal{W}(u_\infty, A, k) = \frac{k}{A} \left(\frac{u_\infty}{A}\right)^{k-1} \exp\left[-\left(\frac{u_\infty}{A}\right)^k\right]. \quad (18)$$

170 2.2 Deterministic Approach

The SLSQP algorithm (Kraft, 1988) is selected to be the deterministic optimization algorithm to act as a benchmark to the SGD approach. SLSQP is a conventional deterministic optimization approach. It is employed in many open-source engineering design codes (Allen et al., 2020; Wu et al., 2020; Ciavarra et al., 2022) and has been used in other previous comparisons of optimization algorithms (Lam et al., 2018; Li and Zhang, 2021; Fleming et al., 2022).

175 The spacing and boundary constraints are passed to the optimizer as individual inequality constraints. The spacing constraints are defined as

$$C_{ij} = (x_i - x_j)^2 + (y_i - y_j)^2 - (N_D D)^2 \quad \forall i, j > i, \quad (19)$$

where \mathbf{C} is an upper triangular matrix of nonlinear inequality constraints.

Square wind farm boundaries are represented using four inequality constraints per turbine,

$$\begin{aligned}
 D_{i1} &= x_i - x_u \\
 D_{i2} &= x_l - x_i \\
 D_{i3} &= y_i - y_u \\
 180 \quad D_{i4} &= y_l - y_i,
 \end{aligned} \tag{20}$$

and the circular wind farm boundaries are represented with one inequality constraint per turbine,

$$D_i = \sqrt{x^2 + y^2} - R, \tag{21}$$

where D is a matrix of boundary constraints that must be less than or equal to zero.

3 Application

185 We apply the optimization approaches discussed above to optimize wind power plants of various sizes using the TOPFARM framework (DTU Wind Energy Systems, 2023b). Each farm consists of turbines with 70-meter hub heights, 80-meter rotor diameters, and 2-megawatt rated powers. Power is computed using PyWake (Pedersen et al., 2019), which is an open-source wake modeling tool that has been used in several related studies (Riva et al., 2020; Rodrigues et al., 2022; Ciavarra et al., 2022; van der Laan et al., 2022; Fischereit et al., 2022). Power gradients are computed directly from PyWake using automatic
 190 differentiation. The power of each turbine is estimated by a combination of velocity deficits predicted by the Bastankhah Gaussian wake model (Bastankhah and Porté-Agel, 2014) using the default parameters in the PyWake tool and the squared sum superposition (Pedersen et al., 2019; DTU Wind Energy Systems, 2023a). We require each turbine to be spaced at minimum two rotor diameters apart ($N_D = 2$). This is imposed as an optimization constraint. We considered wind farms with square and circular boundaries. The square wind farm boundaries are determined as

$$\begin{aligned}
 x_l &= 0 \\
 y_l &= 0 \\
 195 \quad x_u &= D(\sqrt{N_t} - 1)\Delta \\
 y_u &= D(\sqrt{N_t} - 1)\Delta,
 \end{aligned} \tag{22}$$

and, in cases with circular wind farm boundaries, the radius is determined as

$$R = D(\sqrt{N_t} - 1)\Delta, \tag{23}$$

where the Δ parameter control the average spacing of the turbines. In this study, $\Delta = 5$.

We use the pyoptsparse driver (Wu et al., 2020) SLSQP (Kraft, 1988) implementation (Virtanen et al., 2020) in TOPFARM.
 200 The optimizer was set to run for 300 maximum iterations with a tolerance of 10^{-1} . The TOPFARM “expected_cost” parameter

is set to 10. The turbine coordinates are normalized from 0 to 1. In each optimization iteration, the AEP, and the corresponding gradient, is computed using rectangular quadrature as described in Equation 4, using 360 wind direction bins and 23 wind speed bins, resulting in 8280 power evaluations.

The wind rose, visualized in Figure 1, is based on PyWake’s Lillgrund example site. A probability mass function is assigned to different direction bins. Each direction bin is associated with Weibull scale and shape parameters describing the distributions of wind speeds within the sector. This probability mass is derived from seven months of measured data used in a previous study (Göçmen and Giebel, 2016). Each direction bin is 30 degree wide. The reasoning behind this is similar that behind the IEC 614400 power curve standard (International Electrotechnical Commission, 2005)—it is crucial that the reference data consider a statistically significant amount of data in each bin. This coarse direction discretization results in a faster convergence of the estimated probability mass function than a finer discretization would. The continuous probability density function $\pi(u_\infty, \theta)$ is approximated as $\rho(\theta)\pi(u_\infty|\theta)$, where ρ is the previously mentioned probability mass function, linearly interpolated across one-degree bins, and $\pi(u_\infty|\theta)$ is parameterized by direction-specific Weibull shape and scale parameters that are also linearly interpolated from the provided data. With this formulation, the likelihood of different wind directions is provided as a probability mass function, $\rho(\theta)$. This probability mass is used as weights passed to the Numpy “choice” function (Harris et al., 2020), allowing the wind direction to be sampled as a discrete random variable. We note that this formulation could be extended to a fully continuous formulation by drawing the direction samples from the inverse of an empirical cumulative direction density function.




Figure 1. Lillgrund wind speed and direction probability mass function with 360 direction bins and five wind speed bins, where the probability mass function is a linear interpolation of coarser measurements.

In all wind farm optimization problems considered, constraint gradients, and the associated penalty function gradients, are computed analytically. The AEP gradient is computed via automatic differentiation. The directions are discretized from 0 to 360 degrees, with one-degree increments. In the deterministic formulation, the discretized wind speed ranges from 3-25 m/s and is divided using increments of 1 m/s.

While each Monte Carlo estimate of AEP has significant error, the average error will be close to zero throughout the course of the SGD optimization. We compare the accuracy of the Monte Carlo approach (Equation 5) and the quadrature approach (Equation 4) to estimate the AEP and the L-2 norm of the AEP gradient. The true values are estimated with a very fine discretization of speed and direction, 0.2 m/s and 0.2 degrees. These are used as reference values to assess the accuracy of the Monte Carlo and quadrature approaches by comparing the errors associated with both approaches as functions of the number of samples and the discretization level, respectively, when analyzing a 100-turbine farm with square boundaries. This convergence analysis is shown in Figure 2. While some realizations of the Monte Carlo approach yield more accurate results than the quadrature approach, the quadrature approach is generally more accurate than Monte Carlo sampling. The Monte Carlo approach requires on average around ten times as many power evaluations to obtain the same accuracy as the deterministic approach.




Figure 2. Convergence of AEP (left) and L-2 norm of the AEP gradient (right) with respect to the number of samples used in the quadrature and Monte Carlo techniques. The grey cloud shows the 90 percent confidence interval associated with the Monte Carlo approach, using 50 samples. The dashed black line shows the average error associated with the Monte Carlo approach. The solid black line shows the error associated with the deterministic approach.

In this study, we select 50 samples for every SGD iteration. Figure 3 shows the measured computational cost of computing AEP gradients using the circular wind farm described in this study, with different wind farm sizes. The minimum measured time is reported as the minimum of thirty identical runs on the DTU Sophia supercomputer (Technical University of Denmark, 2019). The computational time generally scales logarithmically with the number of turbines. This is to be expected, as there are more interaction terms in the wake model as more turbines are considered. The computational time does not scale logarithmically

with the number of wind rose samples. For small numbers of turbines, evaluating 10 wind rose samples is about as expensive as evaluating 50 samples. This scaling changes as the wind farm grows in size, and it gradually becomes more expensive to sample the wind rose. The evaluation time appears to converge to a logarithmic scaling for large numbers of wind rose samples.

240 These scaling results are likely influenced by memory limitations.




Figure 3. Computational time associated with computing the gradient for various wind farm sizes and wind rose sampling strategies. The left panel compares the cost of computing AEP gradients when using different numbers of Monte Carlo samples with the cost of the full factorial wind rose (8,280 samples) for farms with various numbers of wind turbines. The right panel compares the cost of Monte Carlo estimates of the AEP gradient for different numbers of samples of the atmospheric conditions and turbines in the wind farm.

The optimization algorithms are timed based on the time elapsed between the first and final optimization gradient evaluations. Each optimization case is run on [first generation AMD EPYC 7351](#) processors.

4 Results and Discussion

In the following subsections, the performance of SGD and SLSQP are compared for wind farms with square and circular
245 boundaries, and the sensitivity of the SGD algorithm is assessed.

4.1 Square Wind Farm

The performance of SGD is compared to the deterministic counterpart, considering wind farms with 100, 144, 225, and 324 turbines, with square boundary constraints, using 20 different initial starting conditions to obtain statistically significant results. The AEP, constraint violation (γ), and time elapsed associated with each optimization solution are plotted in Figure 4. The
250 SGD approach consistently yields higher AEPs than the SLSQP approach when the number of scheduled SGD iterations, T , is 2,000. There is a large range of computational times associated with the SLSQP approach, though the computational expense of SLSQP generally grows much larger than SGD as the number of turbines is increased. This is largely due to the nature of the turbine spacing constraint, the size of which grows as the number of turbines squared. SLSQP takes requires about

as much computational time as the SGD approach with 500 scheduled iterations when there are 100 turbines in the square farm. As the number of turbines grows, the average time required by SLSQP becomes more costly than SGD with 2,000 scheduled iterations. The computational cost of SLSQP is a strong function of the initial layout, and the variance of the SLSQP optimization time also increases with the wind farm size. This is due to the complex interaction between the linear boundary constraints and nonlinear spacing constraints. As the proposed SGD formulation does not offer an automatic way to set the number of SGD iterations, T , results are shown for different values of T . When T is increased, the optimizer finds solutions with larger AEPs, with a computational cost that is approximately proportional to T . The SGD solution consistently improves as more optimization iterations are scheduled (larger values of T). Results associated with 1,000 SGD iterations tend to yield similar AEP to the SLSQP approach and results with 2,000 SGD iterations tend to yield higher AEPs than the SLSQP designs.

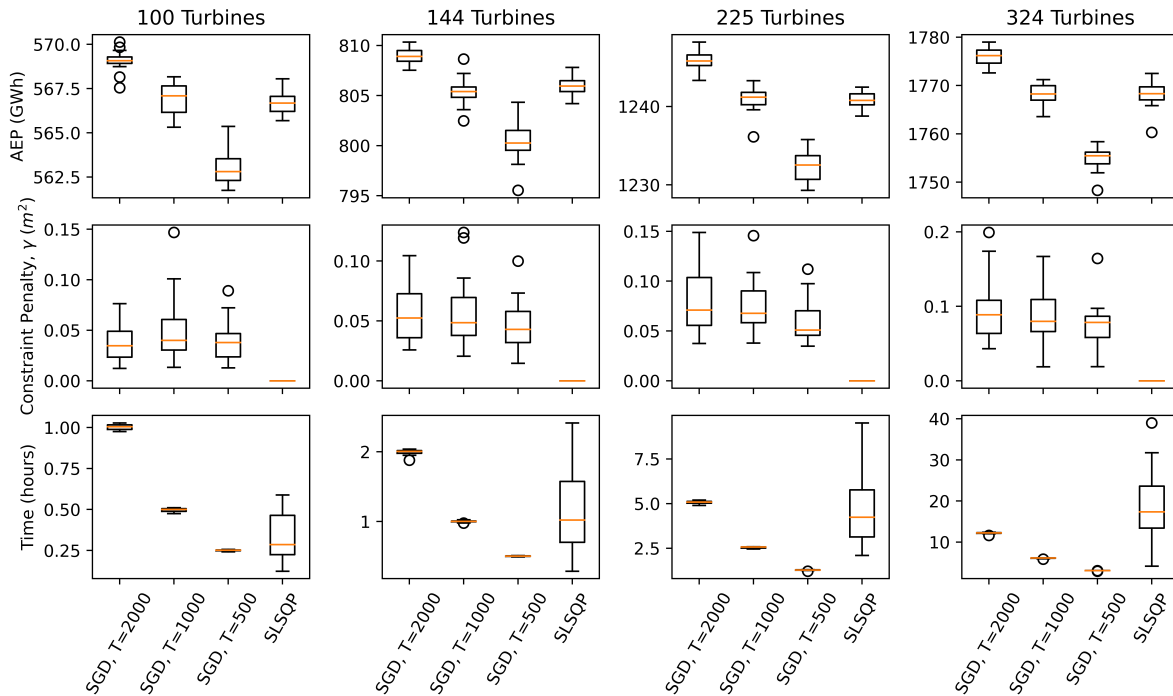


Figure 4. Optimization results associated with SGD and SLSQP for square wind farms with 144, 225, and 324 turbines, using 20 random initial starting conditions. The AEP (top panel), constraint penalty (middle panel), and computational time (bottom panel) are plotted as boxplots. The SGD results are plotted for $T = 500, 1,000,$ and $2,000$ iterations.

The final layouts associated with one of the random initial conditions used in the 324 turbine analysis, when $T = 2,000$ iterations, is shown in Figure 5. The SGD approach generally identifies solutions with the majority of turbines packed into the side boundaries. The deterministic algorithm also packed turbines into the edges of the farm, although not as many turbines were packed into the East and West boundaries as in the SGD results. The layouts found using the SGD approach tend to have interior turbines that generally appear to be more aligned in the North-South direction than in the deterministic solutions.




Figure 5. The final layouts found using SGD (left) and SLSQP (right) using one of the random initial layouts examined in a 324-turbine wind farm with square boundaries for $T = 2,000$ iterations. The spacing constraint is visualized using hollow circles with dashed lines and a radius of one rotor diameter.

The results of the 100-, 144-, 225-, and 324-turbine wind farm optimization cases are summarized in Table 1. The mean time, mean constraint violation, and mean and standard deviation of the AEP are reported with respect to the 20 random initial starting conditions. The constraint violation is reported as $\sqrt{\gamma/N_t}$ to quantify the mean length of the constraint violations of each turbine. The final constraint violations can be reduced by lowering the η_T parameter. In all of these cases, the SLSQP optimization resulted in solutions with zero constraint violations. This is likely because of the linear formulation of the boundary constraints—when a solution satisfies the spacing constraint, any solutions that satisfy the boundary constraints can quickly be found. SGD with 2,000 iterations generally yields solutions with AEP that are 3-5% higher than the solutions found using SLSQP. This is likely because the SGD algorithm is able to better explore the design space by initially relaxing the constraints, allowing for some initial constraint violations.

4.2 Circular Wind Farms

To ensure that the previously presented results are not specific to square wind farms, we performed a similar set of analysis examining circular wind farms. The results yield similar trends to the analysis using square wind farms—SGD becomes significantly less time-consuming than SLSQP as the number of turbines and generally yields solutions with slightly larger AEPs.

Circular wind farms were optimized using 20 random initial layouts, examining different farm sizes, using the SGD and SLSQP optimization algorithms. The results are summarized in Figure 6. The circular wind farm optimization generally took

285 longer than the square wind farm when using the SLSQP optimizer. This is likely due to the more complicated nature of the
 circular boundary when using Cartesian coordinates. These results are similar to the results in Section 4.1—as the number of
 wind turbines and scheduled SGD iterations increases, SGD tends to find solutions with larger AEPs in less computational
 time.

290 The results of the circular wind farm optimization are compared between the SGD and SLSQP optimizers in Table 2, where
 SGD is scheduled to run for 2,000 optimization iterations. SGD generally results in about 0.5% more AEP in significantly
 less time than SLSQP as the number of turbines is increased. The results are also compared in Figure 7. The SGD optimizer
 generally results in more turbines on the boundary edge than the SLSQP optimizer.

Table 1. Results of SGD and deterministic optimizations for various square wind farm sizes. Each optimization case is run using 20 random
 initial starting conditions, and the mean and standard deviation are reported with respect to these 20 initial points. The SGD results are
 associated with $T = 2,000$ iterations.

N_t	Case	Mean Time (hours)	Mean AEP (kWh)	AEP Standard Deviation (kWh)	Mean $\sqrt{\gamma/N_t}$ (m)
100	Deterministic	0.34	5.667e+08	6.223e+05	0.000e+00
	SGD	1.00	5.691e+08	5.347e+05	1.919e-03
144	Deterministic	1.18	8.059e+08	9.188e+05	0.000e+00
	SGD	2.00	8.089e+08	7.727e+05	1.642e-03
225	Deterministic	4.79	1.241e+09	1.004e+06	0.000e+00
	SGD	5.07	1.246e+09	1.110e+06	1.250e-03
324	Deterministic	18.69	1.768e+09	2.556e+06	0.000e+00
	SGD	12.20	1.776e+09	1.810e+06	9.484e-04

Table 2. Results of SGD and deterministic optimizations for various circular wind farm sizes. Each optimization case is run using 20 random
 initial starting conditions, and the mean and standard deviation are reported with respect to these 20 initial points. The SGD results are
 associated with $T = 2,000$ iterations.

N_t	Case	Mean Time (hours)	Mean AEP (kWh)	AEP Standard Deviation (kWh)	Mean $\sqrt{\gamma/N_t}$ (m)
100	Deterministic	1.92	5.383e+08	1.193e+06	2.742e-02
	SGD	0.98	5.407e+08	7.844e+05	2.203e-03
144	Deterministic	7.25	7.638e+08	1.512e+06	8.648e-02
	SGD	1.96	7.676e+08	8.777e+05	1.726e-03
225	Deterministic	38.31	1.174e+09	1.573e+06	2.359e-01
	SGD	4.98	1.180e+09	1.306e+06	1.280e-03




Figure 6. Optimization results associated with SGD and SLSQP for wind farms with 144, 225, and 324 turbines, using 20 random initial starting conditions. The AEP (top panel), constraint penalty (middle panel), and computational time (bottom panel) are plotted as box and whisker plots. The SGD results are plotted for $T = 500, 1,000,$ and $2,000$ iterations.




Figure 7. The final layouts found using SGD (left) and SLSQP (right) using one of the random initial layouts examined in the 225-turbine wind farm with circular boundaries using $T = 2,000$ iterations. The final turbine layouts are shown as filled circles. The spacing constraint is visualized using hollow circles with dashed lines and a radius of one rotor diameter.

4.3 Sensitivity Analysis

295 There are several parameters in the SGD algorithm that were tuned to perform reasonably well. In this section, we investigate the sensitivity of the SGD optimization results with respect to the early stopping option in Algorithm 1, the number of Monte Carlo samples per optimization iteration, the learning rate schedule, as well as the initial and final learning rates.

As the optimization progresses, the constraint multiplier, α_i , becomes large (approaching 10 as the learning rate approaches 0.1), and the gradients of the AEP are overwhelmed by the gradients of the penalty, which take very little time to compute. This situation can be taken advantage of using the early stopping option in Algorithm 1. The solution tends to terminate quickly when the optimizer only follows the deterministic gradient (the optimization engine terminates when the constraint gradients are zero). Figure 8 shows the AEP, constraint violation, and computational time associated five random initial layouts, using threshold parameters of 0.01, 0.05, and 0.1, as well as the SGD algorithm as applied in the previous sections, without the early stopping option activated. The use of each early stopping option results in layouts without constraint violations. As the threshold parameter is increased, the AEP is slightly reduced and the total computational time decreases. A threshold parameter of 0.1 results in approximately 0.3% reduction in AEP and 44% reduction in computation time.

300

305

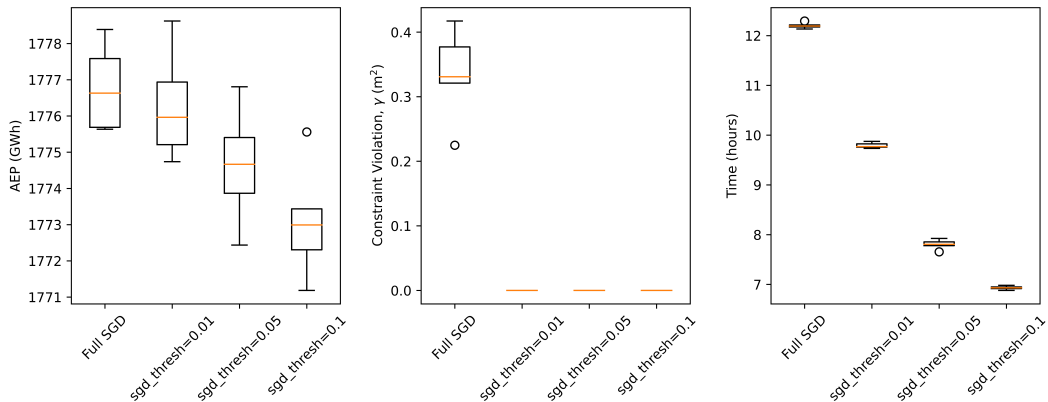


Figure 8. Results using different configurations of the “early stopping” option in TOPFARM, as well as the SGD optimization without early stopping, considering the circular wind farm with 225 turbines, with $T = 2,000$, using five random initial layouts.

The optimization results presented in this study used 50 power samples per iteration ($K = 50$). We found this to produce high-quality results without incurring unacceptable computational expense. Figure 9 shows the behavior of the SGD approach associated with different different values of K , considering 100 turbines with 2,000 scheduled optimization iterations. As K increases, the optimization finds solutions with larger AEPs. There is a small increase in time elapsed and a large increase in the final AEP between the $K = 5$ and $K = 50$ cases, while there is a large increase in time elapsed and a small increase in the final AEP between $K = 50$ and $K = 200$. As K increases, we expect the maximum AEP to reach a plateau and the time

310

and memory required to increase indefinitely. In future work, we plan to explore scheduling K to change as the optimization progresses.




Figure 9. Optimization results associated with SGD for a 100-turbine wind farm using 20 random initial starting conditions and $T = 2,000$. The upper and lower bounds of the results are plotted as a function of the optimization iteration number. The SGD results associated with $K = 5, 50$, and 200 iterations are shown in purple, blue, and yellow, respectively.

315 This study used an exotic learning rate scheduler. We tried several schedulers, and observed this one to be the best at finding sufficiently large AEP solutions that reasonably satisfied the imposed constraints. Figure 10 shows the behavior of the SGD algorithm associated with the presented learning rate scheduler, referred to here as the product scheduler; as well as an exponential and a linear decay scheduler. The exponential scheduler quickly diminishes the learning rate, causing the SGD algorithm to become stuck in local minima. The linear transition from large to lower learning rates prevents the SGD algorithm
 320 from having sufficient time to follow enlarged constraint gradients. It is possible that the algorithm could be improved by using separate schedulers for the learning rate and constraint multiplier. For instance, it might be more effective to use a linear scheduler to decrease the learning rate and an exponential scheduler to increase the constraint multiplier. We leave this question for future work.

The initial and final learning rates, η_0 and η_T , have units of distance and correspond to the initial and final step size of the
 325 optimization algorithm. The final learning rate can be interpreted as the degree to which the constraint are to be satisfied, since this will be the step size the optimization algorithm uses when α_i is large and the constraint gradients overwhelm the AEP gradients. This is illustrated in the left panel of Figure 11, which shows the results of several SGD optimizations using different initial and final learning rates. On average, there is a linear relationship between the constraint violation of the solution and η_T , where the average final constraint violation is approximately two times η_T . The maximum observed constraint violation
 330 is approximately four times η_T . In addition, there is a trade-off between the AEP and constraint violation of the final solution. This tradeoff is influenced by the initial and final learning rates. It is important to tune the initial learning rate. An initial learning rate that is too low will result in very little exploration. Initial learning rates that are too high will result in a rapid

influx of penalty violations that overwhelm AEP gradients throughout the optimization. From our experiments, we found a step size of one fifth of the rotor diameter to produce satisfactory results, as shown in the right panel of Figure 11.

335 5 Conclusions

SGD is a promising optimization tool for wind farm design. Instead of evaluating all anticipated atmospheric conditions during every optimization iteration, SGD randomly samples the defined distributions of atmospheric conditions, resulting in substantially reduced computational time required for each optimization iteration. The total optimization time can be scheduled according to a prescribed computational budget. The presented formulation allows for continuous resolution of uncertain




Figure 10. Influence of the learning rate scheduler on the SGD optimization, considering a 100-turbine wind farm with a square boundary. The product scheduler is shown in yellow. The exponential scheduler is shown in blue. The linear scheduler is shown in purple. The AEP (upper left panel), constraint penalty function (bottom left panel), and the learning rate decay (right panel) are plotted as a function of the number of optimization iterations. The optimization iteration is denoted as t in the legend of the right panel.




Figure 11. (Left) The final constraint penalty plotted against the final learning rate. (Right) The final AEP plotted against the final constraint penalty. The different colors represent different initial learning rates. The results of 20 initial starting positions are plotted as points. The average results of the 20 initial starting conditions are connected as lines. These data are associated with 100-turbine wind farms with square boundaries and $T = 2,000$ iterations.

340 variables, eliminating the need to choose a discretization resolution of atmospheric conditions, such as the wind speed and direction. This technique does not become exponentially more expensive as a greater number of uncertain parameters is included, allowing for consideration of other atmospheric conditions, such as turbulence intensity, air density, veer, and shear (Saint-Drenan et al., 2020; Duc et al., 2019).

The presented SGD approach was shown to become more effective than a deterministic counterpart as the number of wind turbines increased. SGD yielded slightly higher AEPs than the deterministic approach in substantially reduced computational time. The time required to optimize wind farm layouts can be a major bottleneck in corporate workflows, and the time savings associated with the SGD approach allows engineers to access optimization results sooner than a conventional approach. If the inflow conditions were discretized using extremely small bins, or if several atmospheric conditions were to be considered, we expect that the SGD approach would perform the optimization even faster and more effectively than the deterministic approach.

350 The SGD approach is a simple framework that is well suited large-scale stochastic wind power plant design optimization challenges. [This framework is available in the open-source TOPFARM package.](#) Future work includes: exploring separate schedulers for the constraint multiplier and learning rate and scheduling the number of Monte Carlo samples, K , to change as the optimization proceeds.

Code availability. The code used in this study is available from DTU Wind Energy Systems' PyWake and TOPFARM repositories (DTU Wind Energy Systems, 2023a, b)

Author contributions. JQ and PER designed the experiments. JQ, PER, MMP, and RVP developed the problem formulation. JQ developed the SGD formulation. JQ, RVP, MFM, and MMP performed the simulations. JQ, MMP, and RVP prepared the manuscript with contributions from all co-authors.

Acknowledgements. The authors gratefully acknowledge the computational and data resources provided on the Sophia HPC Cluster at the
360 Technical University of Denmark, DOI: 10.57940/FAFC-6M81

References

- H. Alibrahim and S. A. Ludwig. Hyperparameter optimization: Comparing genetic algorithm against grid search and bayesian optimization. In *2021 IEEE Congress on Evolutionary Computation (CEC)*, pages 1551–1559, 2021. <https://doi.org/10.1109/CEC45853.2021.9504761>.
- J. Allen, R. King, and G. Barter. Wind farm simulation and layout optimization in complex terrain. In *Journal of Physics: Conference Series*, volume 1452, page 012066. IOP Publishing, 2020.
- J. Annoni, P. Fleming, A. Scholbrock, J. Roadman, S. Dana, C. Adcock, F. Porte-Agel, S. Raach, F. Haizmann, and D. Schlipf. Analysis of control-oriented wake modeling tools using lidar field results. *Wind Energy Science*, 3(2):819–831, 2018.
- N. F. Baker, A. P. Stanley, J. J. Thomas, A. Ning, and K. Dykes. Best practices for wake model and optimization algorithm selection in wind farm layout optimization. In *AIAA Scitech 2019 forum*, page 0540, 2019.
- 370 M. Bastankhah and F. Porté-Agel. A new analytical model for wind-turbine wakes. *Renewable energy*, 70:116–123, 2014.
- R. H. Byrd, S. L. Hansen, J. Nocedal, and Y. Singer. A stochastic quasi-newton method for large-scale optimization. *SIAM Journal on Optimization*, 26(2):1008–1031, 2016.
- A. W. Ciavarra, R. V. Rodrigues, K. Dykes, and P.-E. Réthoré. Wind farm optimization with multiple hub heights using gradient-based methods. In *Journal of Physics: Conference Series*, volume 2265, page 022012. IOP Publishing, 2022.
- 375 C. E. Clark, G. Barter, K. Shaler, and B. DuPont. Reliability-based layout optimization in offshore wind energy systems. *Wind Energy*, 25(1):125–148, 2022.
- J. Criado Risco, R. Valotta Rodrigues, M. Friis-Møller, J. Quick, M. Mølgaard Pedersen, and P.-E. Réthoré. Gradient-based wind farm layout optimization with inclusion and exclusion zones. *Wind Energy Science Discussions*, 2023:1–24, 2023. <https://doi.org/10.5194/wes-2023-5>. URL <https://wes.copernicus.org/preprints/wes-2023-5/>.
- 380 C. Croonenbroeck and D. Hennecke. A comparison of optimizers in a unified standard for optimization on wind farm layout optimization. *Energy*, 216:119244, 2021.
- S. De, J. Hampton, K. Maute, and A. Doostan. Topology optimization under uncertainty using a stochastic gradient-based approach. *Structural and Multidisciplinary Optimization*, 62(5):2255–2278, 2020.
- M. Denkowski and G. Neubig. Stronger baselines for trustable results in neural machine translation. *arXiv preprint arXiv:1706.09733*, 2017.
- 385 DTU Wind Energy Systems. Pywake.
<https://gitlab.windenergy.dtu.dk/TOPFARM/PyWake>, commit 3a61c8aec34a2505d0b460da2ad831aeaf46802b, 2023a.
- DTU Wind Energy Systems. Topfarm.
<https://gitlab.windenergy.dtu.dk/TOPFARM/Topfarm2>, commit 7d9deb62dd62b7db2fe38dc0a629b137ac06f9c7, 2023b.
- T. Duc, O. Coupiac, N. Girard, G. Giebel, and T. Göçmen. Local turbulence parameterization improves the jensen wake model and its implementation for power optimization of an operating wind farm. *Wind Energy Science*, 4(2):287–302, 2019.
- 390 J. Feng and W. Z. Shen. Solving the wind farm layout optimization problem using random search algorithm. *Renewable Energy*, 78:182–192, 2015.
- J. Fischereit, K. Schaldemose Hansen, X. G. Larsén, M. P. van der Laan, P.-E. Réthoré, and J. P. Murcia Leon. Comparing and validating intra-farm and farm-to-farm wakes across different mesoscale and high-resolution wake models. *Wind Energy Science*, 7(3):1069–1091, 2022.
- 395 P. A. Fleming, A. P. Stanley, C. J. Bay, J. King, E. Simley, B. M. Doekemeijer, and R. Mudafort. Serial-refine method for fast wake-steering yaw optimization. In *Journal of Physics: Conference Series*, volume 2265, page 032109. IOP Publishing, 2022.

- P. Gebraad, J. J. Thomas, A. Ning, P. Fleming, and K. Dykes. Maximization of the annual energy production of wind power plants by optimization of layout and yaw-based wake control. *Wind Energy*, 20(1):97–107, 2017.
- 400 T. Göçmen and G. Giebel. Estimation of turbulence intensity using rotor effective wind speed in lillgrund and horns rev-i offshore wind farms. *Renewable energy*, 99:524–532, 2016.
- M. Godinho and R. Castro. Comparative performance of ai methods for wind power forecast in portugal. *Wind Energy*, 24(1):39–53, 2021.
- P. Graf, K. Dykes, G. Scott, J. Fields, M. Lunacek, J. Quick, and P-E. Rethore. Wind farm turbine type and placement optimization. In *Journal of Physics: Conference Series*, volume 753, page 062004. IOP Publishing, 2016.
- 405 D. Guirguis, D. A. Romero, and C. H. Amon. Toward efficient optimization of wind farm layouts: Utilizing exact gradient information. *Applied energy*, 179:110–123, 2016.
- C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del R'io, M. Wiebe, P. Peterson, P. G'erard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, Sept. 410 2020. <https://doi.org/10.1038/s41586-020-2649-2>. URL <https://doi.org/10.1038/s41586-020-2649-2>.
- C. B. Hasager, L. Rasmussen, A. Peña, L. E. Jensen, and P-E. Réthoré. Wind farm wake: The horns rev photo case. *Energies*, 6(2):696–716, 2013.
- J. F. Herbert-Acero, O. Probst, P-E. Réthoré, G. C. Larsen, and K. K. Castillo-Villar. A review of methodological approaches for the design and optimization of wind farms. *Energies*, 7(11):6930–7016, 2014.
- 415 M. F. Howland, A. S. Ghate, J. B. Quesada, J. J. Pena Martínez, W. Zhong, F. P. Larrañaga, S. K. Lele, and J. O. Dabiri. Optimal closed-loop wake steering—part 2: Diurnal cycle atmospheric boundary layer conditions. *Wind Energy Science*, 7(1):345–365, 2022.
- M. N. Hussain, N. Shaukat, A. Ahmad, M. Abid, A. Hashmi, Z. Rajabi, and M. A. U. R. Tariq. Micro-siting of wind turbines in an optimal wind farm area using teaching–learning-based optimization technique. *Sustainability*, 14(14):8846, 2022.
- International Electrotechnical Commission. IEC 61400-12-1 Wind Turbines-Part 12-1: Power Performance Measurements of Electricity 420 Producing Wind Turbines. *International Electrotechnical Commission (IEC), Geneva, Switzerland*, 1, 2005.
- H. Kervadec, J. Dolz, J. Yuan, C. Desrosiers, E. Granger, and I. B. Ayed. Constrained deep networks: Lagrangian optimization via log-barrier extensions. *arXiv preprint arXiv:1904.04205*, 2019.
- N. Ketkar. Stochastic gradient descent. In *Deep learning with Python*, pages 113–132. Springer, 2017.
- R. King, A. Glaws, G. Geraci, and M. S. Eldred. A probabilistic approach to estimating wind farm annual energy production with bayesian 425 quadrature. In *AIAA Scitech 2020 Forum*, page 1951, 2020.
- R. N. King, K. Dykes, P. Graf, and P. E. Hamlington. Optimization of wind plant layouts using an adjoint approach. *Wind Energy Science*, 2 (1):115–131, 2017.
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- K. Kölle, T. Göçmen, I. Eguinoa, L. A. Alcayaga Román, M. Aparicio-Sanchez, J. Feng, J. Meyers, V. Pettas, and I. Sood. Farmconners 430 market showcase results: wind farm flow control considering electricity prices. *Wind Energy Science*, 7(6):2181–2200, 2022.
- D. Kraft. A software package for sequential quadratic programming. *Forschungsbericht- Deutsche Forschungs- und Versuchsanstalt fur Luft- und Raumfahrt*, 1988.
- R. Lam, M. Poloczek, P. Frazier, and K. E. Willcox. Advances in bayesian optimization with applications in aerospace engineering. In *2018 AIAA Non-Deterministic Approaches Conference*, page 1656, 2018.
- 435 J. Li and M. Zhang. Data-based approach for wing shape design optimization. *Aerospace Science and Technology*, 112:106639, 2021.

- J. Liu, Y. Rong, M. Takác, and J. Huang. On the acceleration of l-bfgs with second-order information and stochastic batches. *arXiv preprint arXiv:1807.05328*, 2018.
- P. Márquez-Neila, M. Salzmann, and P. Fua. Imposing hard constraints on deep networks: Promises and limitations. *arXiv preprint arXiv:1706.02025*, 2017.
- 440 P. Moritz, R. Nishihara, and M. Jordan. A linearly-convergent stochastic l-bfgs algorithm. In *Artificial Intelligence and Statistics*, pages 249–258. PMLR, 2016.
- J. Murcia, P.-E. Réthoré, A. Natarajan, and J. D. Sørensen. How many model evaluations are required to predict the aep of a wind power plant? In *Journal of Physics: Conference Series*, volume 625, page 012030. IOP Publishing, 2015.
- M. M. Najafabadi, T. M. Khoshgoftaar, F. Villanustre, and J. Holt. Large-scale distributed l-bfgs. *Journal of Big Data*, 4(1):1–17, 2017.
- 445 A. H. Najd, G. Goksu, and H. F. Hammood. Pitch angle control using neural network in wind turbines. In *IOP Conference Series: Materials Science and Engineering*, volume 928, page 022118. IOP Publishing, 2020.
- A. Ning., K. Dykes., and J. Quick. *Systems engineering and optimization of wind turbines and power plants*, pages 235–292. Wind Energy Modeling and Simulation - Volume 2: Turbine and System. Institution of Engineering and Technology, 2020.
- A. S. Padrón, J. Thomas, A. P. Stanley, J. J. Alonso, and A. Ning. Polynomial chaos to efficiently compute the annual energy production in
450 wind farm layout optimization. *Wind Energy Science*, 4(2):211–231, 2019.
- M. M. Pedersen, P. van der Laan, M. Friis-Møller, J. Rinker, and P.-E. Réthoré. Dtuwindenergy/pywake: Pywake. Feb 2019. <https://doi.org/10.5281/zenodo.2562662>.
- N. Qian. On the momentum term in gradient descent learning algorithms. *Neural networks*, 12(1):145–151, 1999.
- J. Quick, J. King, R. N. King, P. E. Hamlington, and K. Dykes. Wake steering optimization under uncertainty. *Wind Energy Science*, 5(1):
455 413–426, 2020.
- M. Riedmiller and H. Braun. A direct adaptive method for faster backpropagation learning: The rprop algorithm. In *IEEE international conference on neural networks*, pages 586–591. IEEE, 1993.
- R. Riva, J. Liew, M. Friis-Møller, N. Dimitrov, E. Barlas, P.-E. Réthoré, and A. Beržonskis. Wind farm layout optimization with load constraints using surrogate modelling. In *Journal of Physics: Conference Series*, volume 1618, page 042035. IOP Publishing, 2020.
- 460 R. V. Rodrigues, M. Friis-Møller, K. Dykes, N. Pollini, and M. Jensen. A surrogate model of offshore wind farm annual energy production to support financial valuation. In *Journal of Physics: Conference Series*, volume 2265, page 022003. IOP Publishing, 2022.
- S. K. Roy and M. Harandi. Constrained stochastic gradient descent: The good practice. In *2017 International Conference on Digital Image Computing: Techniques and Applications (DICTA)*, pages 1–8, 2017. <https://doi.org/10.1109/DICTA.2017.8227420>.
- S. Ruder. An overview of gradient descent optimization algorithms, 2016. URL <https://arxiv.org/abs/1609.04747>.
- 465 Y.-M. Saint-Drenan, R. Besseau, M. Jansen, I. Staffell, A. Troccoli, L. Dubus, J. Schmidt, K. Gruber, S. G. Simões, and S. Heier. A parametric model for wind turbine power curves incorporating environmental conditions. *Renewable Energy*, 157:754–768, 2020.
- M. Samorani. The wind farm layout optimization problem. *Handbook of wind power systems*, pages 21–38, 2013.
- B. Sanderse. Aerodynamics of wind turbine wakes: Literature review; ecn-e-09-016. *Energy Research Center of Netherland (ECN): Petten, The Netherland*, 2009.
- 470 E. Simley, D. Millstein, S. Jeong, and P. Fleming. The value of wake steering wind farm control in us energy markets. *Wind Energy Science Discussions*, pages 1–26, 2023.
- G. Sivanantham and S. Gopalakrishnan. Stochastic gradient descent optimization model for demand response in a connected microgrid. *KSII Transactions on Internet and Information Systems (TIIS)*, 16(1):97–115, 2022.

- A. P. Stanley, O. Roberts, J. King, and C. J. Bay. Objective and algorithm considerations when optimizing the number and placement of
475 turbines in a wind power plant. *Wind Energy Science*, 6(5):1143–1167, 2021.
- K. Stengel, A. Glaws, D. Hettinger, and R. N. King. Adversarial super-resolution of climatological wind and solar data. *Proceedings of the
National Academy of Sciences*, 117(29):16805–16815, 2020.
- Technical University of Denmark. Sophia HPC Cluster, 2019. URL <https://dtu-sophia.github.io/docs/>.
- Y. Tian, Y. Zhang, and H. Zhang. Recent advances in stochastic gradient descent in deep learning. *Mathematics*, 11(3):682, 2023.
- 480 M. P. van der Laan, O. García-Santiago, M. Kelly, A. Meyer Forsting, C. Dubreuil-Boisclair, K. S. Seim, M. Imberger, A. Peña, N. N.
Sørensen, and P.-E. Réthoré. A new rans-based wind farm parametrization and inflow model for wind farm cluster modeling. *Wind
Energy Science Discussions*, pages 1–38, 2022.
- P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J.
van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng,
485 E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H.
Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in
Python. *Nature Methods*, 17:261–272, 2020. <https://doi.org/10.1038/s41592-019-0686-2>.
- N. Wu, G. Kenway, C. A. Mader, J. Jasa, and J. R. Martins. pyoptparse: A python framework for large-scale constrained nonlinear opti-
mization of sparse systems. *Journal of Open Source Software*, 5(54):2564, 2020.
- 490 K. You, M. Long, J. Wang, and M. I. Jordan. How does learning rate decay help modern neural networks? *arXiv preprint arXiv:1908.01878*,
2019.
- C. Zhang, S. C. Kramer, A. Angeloudis, J. Zhang, X. Lin, and M. D. Piggott. Improving tidal turbine array performance through the
optimisation of layout and yaw angles. *International Marine Energy Journal*, 5(3):273–280, 2022.
- J. Zhang and X. Zhao. Wind farm wake modeling based on deep convolutional conditional generative adversarial network. *Energy*, 238:
495 121747, 2022.
- Z. Zhang, C. Santoni, T. Herges, F. Sotiropoulos, and A. Khosronejad. Time-averaged wind turbine wake flow field prediction using autoen-
coder convolutional neural networks. *Energies*, 15(1):41, 2021.
- T. Zilong and D. X. Wei. Layout optimization of offshore wind farm considering spatially inhomogeneous wave loads. *Applied Energy*, 306:
117947, 2022.