



A comparison of efficient AEP models for Wind Farm Layout Optimization

Javier Andueza-Bueno¹, Leonardo Alcayaga¹, Mikkel Friis-Møller¹, and Pierre-Elouan Réthoré¹

¹Technical University of Denmark, Department of Wind and Energy Systems, Frederiksborgvej 399, 4000 Roskilde, Denmark

Correspondence: Javier Andueza-Bueno (janbu@dtu.dk)

Abstract. As wind farms scale up to benefit from cost reductions, wake effects create significant challenges, increasing the importance of the Wind Farm Layout Optimization (WFLO) problem. Current methods to maximize Annual Energy Production (AEP) by minimizing wake losses are computationally intensive, requiring hundreds of iterations and multiple optimizations per project. In this work we present the results of testing several AEP models with a variety of simplifications, assumptions and approaches as the objective function in gradient based optimizations, comparing them against a reference. These include the conventional rectangular quadrature in different forms and the novel approaches FLOWERS and Bayesian Quadrature, alongside the Stochastic Gradient Descent (SGD) optimization algorithm. Four different wind farms encompassing different boundaries, wind roses, sizes and types of constraints were used. Three different wake models are considered here, introducing the FLOWERS derivation for Nygaard TurbOPark. The results showed that FLOWERS achieved the greatest balance between efficiency (optimization time) and efficacy (quality of the layout found), with time reductions of over 95 % and layouts comparable to the reference, although it encountered difficulties with strongly unimodal wind roses. Similar layouts were found using rectangular quadrature with a coarse wind direction binning; however, this method was not as efficient as FLOWERS, and worked well only in combination with the Gaussian Bastankhah wake model, with optimization time reductions between 40 and 70 %. Nevertheless, it proved to be more versatile, even for strongly unimodal wind roses. The work also highlights the struggles for probabilistic methods to properly capture the physics and turbine interactions in the problem, not properly reflecting the positive or negative effects of modifying the layout despite their ability to estimate AEP accurately. The similarity in results between using the local wind speed or the free flow wind speed to compute the thrust coefficient in a WFLO framework is pointed out. Finally, SGD was found to obtain consistent layouts, albeit with the highest computational cost out of all methods.

1 Introduction

With the expansion of wind energy as an energy source, wind turbines are placed in groups known as wind farms. Within these groups, turbines shadow each other, thereby reducing the amount of energy available for extraction, on account of the wake that is generated by the turbines upstream. The physics, development and consequences of wakes is a widely researched field of study (Sanderse et al., 2011; Vermeer et al., 2003; Wang et al., 2024; Porté-Agel et al., 2020). This phenomenon is described as a volume of flow characterized by a reduced streamwise velocity — perpendicular to the turbine rotor — and increased levels



of turbulence compared to the free stream airflow (Porté-Agel et al., 2020). The power generation of a turbine operating in such conditions is considerably decreased, by up to 40% in the near wake region (Herbert-Acero et al., 2014). The solution could seem simple, spacing the turbines further apart while not placing them directly downstream, however, space is often limited by the wind farm development area restrictions. Furthermore, wind speed and direction are not unique, but instead coming from an often complex distribution, known as the wind rose, which makes it difficult to determine where turbines will experience the least wake losses, given that downstream turbine production is conditional on the wind direction considered. This wake avoidance problem, known more commonly as the Wind Farm Layout Optimization (WFLO) problem, involves the optimization of turbine position so as to optimize the Annual Energy Production (AEP) or Levelized Cost of Energy (LCOE), among other objective functions (Herbert-Acero et al., 2014). AEP is defined as the total amount of energy generated by a wind farm over the course of one year, while LCOE includes economical factors regarding the performance of the wind farm such as its variable or fixed costs (Herbert-Acero et al., 2014). In this study, we focus on AEP as the objective function, which is to be maximized.

The AEP of a given wind farm is a function of the atmospheric conditions at the site, the positions of the turbines, and their power and thrust curves. While the atmospheric conditions – most commonly wind speed and direction, with their corresponding probability densities – are known from measurement campaigns or atmospheric models, the specific wind speed experienced by each wind turbine is more difficult to obtain, and is usually estimated using engineering wake and superposition models (Porté-Agel et al., 2020). AEP is obtained with a continuous integration of power production over time, wind speed and direction. The current approach is through a numerical integration approach using what is commonly referred to as a wind rose. Based on time series data, wind measurements can be grouped in a determined number of wind direction bins of known probability. The wind speeds falling into each of these bins are then grouped in wind speed bins. The probability of each of these bins is obtained by fitting a Weibull distribution, resulting in known Weibull parameters for each wind direction bin considered (Gonzalez-Rodriguez et al., 2025). The most widely used method to estimate AEP is rectangular quadrature or midpoint rule (Criado Risco et al., 2024; Feng and Shen, 2017; Valotta Rodrigues et al., 2024), which solves a wind farm flow simulation for each combination of wind speed and direction in the chosen discretization, taking into account all interactions between wind turbines in the farm, to then add them together taking into consideration their respective weight based on the wind rose.

Optimization algorithms carry out hundreds to thousands of function and gradients evaluations. Combined to the number of wind farm flow simulations that need to be solved per evaluation, this results in optimization times which may be in the order of hours, depending on several factors (Valotta Rodrigues et al., 2024; LoCascio et al., 2025). For this reason, computationally efficient AEP methods are researched. Methods proposed in the literature to reduced the cost of AEP estimation include Polynomial Chaos (Padrón et al., 2019; Murcia et al., 2015), Bayesian Quadrature (King et al., 2020), Monte Carlo integration (Quick et al., 2023) and FLOW Estimation and Rose Superposition (FLOWERS) (LoCascio et al., 2025, 2022, 2024; Whittaker et al., 2025).

Rectangular quadrature will simulate all atmospheric combinations, using equally as much computation effort for highly likely and for very unlikely cases. Bayesian Quadrature addresses this inefficiency by performing flow modeling on optimally selected atmospheric conditions. Rather than evaluating the flow model at every condition, a Gaussian Process (GP) surrogate



is trained on a set of wind direction and speed combinations to estimate the wind farm's power. Crucially, the trained GP's posterior variance can be estimated a priori, enabling the computation of an optimal wind conditions sampling strategy in which the variance of the AEP estimate is minimized. Using Bayesian Quadrature to integrate the GP over the optimal sampling, an accurate AEP estimate can be obtained with substantially fewer wind speed and direction combinations than traditional quadrature. When compared to Rectangular Quadrature with 4433 simulations, Bayesian Quadrature achieves an error smaller than 1 % with only 200 simulations, reducing computational costs considerably (King et al., 2020). However, it has never been tested as the objective function in WFLO problems, to our knowledge. On the other hand, the use of GP trained on layouts and their respective AEPs to then use in Bayesian Optimization for WFLO has been widely studied (Pujari and Mitra, 2023; Chugh and Ymeraj, 2022).

A more straightforward approach is Monte Carlo, which is characterized by a simple method and a relatively uncomplicated configuration. This method has been employed in a broad spectrum of mathematical problems, encompassing various applications, including optimization, numerical integration, and uncertainty assessments. The method is based on the random sampling of uncertain input variables from their respective probability density functions, followed by the evaluation of the desired function. In the context of AEP, the variables under consideration are wind direction and speed. These variables are drawn from the wind rose (Quick et al., 2023; Padrón et al., 2019). Despite being substantially simpler than other AEP methods such as Bayesian Quadrature, it is one of the most computationally expensive methods. Around 200,000 samples are required to ensure that the 99 % confidence interval for the true AEP is smaller than 1 % of the computed AEP (Padrón et al., 2019). Nonetheless, the Monte Carlo AEP method can be used for layout optimization methods where a very noisy AEP gradient is needed. This is the case of the Stochastic Gradient Descent algorithm (Quick et al., 2023), where a coarse AEP gradient is needed and small number of samples is required. This method is an adaptation of the Adam algorithm (Kingma and Ba, 2017) for WFLO, and it has been shown to yield higher AEP in substantially less time than conventional methods as the number of wind turbines increase (Quick et al., 2023).

All of the AEP methods discussed so far rely on the discretization of the integral to make it tractable, whether it is to run the full discretized wind rose or a certain number of scenarios. The novel FLOW Estimation and Rose Superposition (FLOWERS) aims to go a step further and solve the integral in an analytical manner, rather than numerical. The method, first described in LoCascio et al. (2022) and then improved into a more generalized case in LoCascio et al. (2024) makes use of several simplifications and the Fourier Transform to end up with an integral that is tractable. Initially, in LoCascio et al. (2022) and LoCascio et al. (2024), the method had been developed to use the N.O. Jensen (NOJ) wake model (Jensen, 1983), here referred to as NOJ-FLOWERS. The resulting AEP expression's performance was compared for different wind roses, layouts and wind farm sizes, using the AEP obtained by means of Rectangular Quadrature using NOJ wake model as reference. NOJ-FLOWERS proved itself to compute AEP at least one order of magnitude faster, but it was found to underestimate AEP by up to 30 % in sites with high wind speeds. NOJ wake model is known to underestimate wake velocity at hub height, while overestimating that at the wake edge (Wang et al., 2024). To better represent wind turbine wakes, the Gaussian Bastankhah (BPA) wake model (Bastankhah and Porté-Agel, 2014) was implemented into the FLOWERS approach in Whittaker et al. (2025), here referred to as BPA-FLOWERS. This method provides similar results to NOJ-FLOWERS in terms of computational cost, with



the addition of an increased accuracy with respect to NOJ-FLOWERS when validated against the Cumulative Curl model from FLORIS. NOJ-FLOWERS was shown to underestimate AEP, whereas BPA-FLOWERS yields a slight overestimation (Whittaker et al., 2025). For the case of NOJ-FLOWERS, it has been tested for a gradient based layout optimization problem for small wind farms in LoCascio et al. (2024), and for larger wind farms in LoCascio et al. (2025), achieving promising results. It has a decreased computational time by at least an order of magnitude with respect to conventional methods. Despite the fact that it underestimates AEP and its accuracy is dependent on the wind rose, it still captures the positive or negative impacts of modifying the layout. In addition, it provides a much smoother design space, which will prove advantageous in an optimization scenario (Martins and Ning, 2021). When it comes to BPA-FLOWERS, its performance has not yet been analyzed for optimization purposes.

Regarding WFLO algorithms, they fall into two major groups in the literature: gradient based and gradient free. Gradient based algorithms, sometimes referred to as calculus based algorithms, rely on the derivatives or gradients of the objective function to know the direction in which to explore the design space. These methods require the objective function to be differentiable and continuous (Herbert-Acero et al., 2014; Martins and Ning, 2021). Despite the multimodality of the layout optimization problem (Thomas and Ning, 2018; Thomas et al., 2022), gradient based algorithms have been proven to attain strong results (Guirguis et al., 2016; Fleming et al., 2022). On the other hand gradient free algorithms do not make use of the derivatives or gradients to determine the direction. These algorithms are commonly divided into heuristics and metaheuristics, the latter being more efficient and complex, often achieving near-optimal solutions (Herbert-Acero et al., 2014). However, these algorithms are known for not scaling well with the number of variables, meaning that, unlike gradient based algorithms, an increase in the number of design variables corresponds to a similar increase in the computational cost (Herbert-Acero et al., 2014). Ultimately the choice of algorithm depends greatly on the nature of the problem, the desired quality-to-computational cost ratio, and the gradients availability.

As outlined above, a variety of models have been used to reduce the computational cost of WFLO. While some methods have been used in an optimization framework, others – such as Bayesian Quadrature to integrate a wind farm power GP surrogate model to estimate AEP – have not yet been tested. However, simpler methods of computing AEP that do not require such complex mathematical derivations may be used to efficiently evaluate wind farm layouts, such as using a reduced number of wind direction bins or an average wind speed bin per wind direction. The objective of this paper is to provide a fair comparison of different conventional and new AEP models that can be used for layout optimization, focusing particularly on computational efficiency and the quality of the results achieved. The Stochastic Gradient Descent algorithm, despite not being an AEP model, will be assessed alongside the other methods given its efficient sampling strategy to obtain the gradients. A series of boundary conditions will be imposed on four different real offshore wind farms of various sizes and presenting different atmospheric conditions, with the aim of exploring the WFLO performance of the AEP methods under a variety of conditions. The optimization will be carried out using a gradient based algorithm, thus reflecting how well the model's gradients guide the optimizer to reach optimal layouts. Each of the different AEP methods will be assessed using three different wake models in order to explore its effect on the efficiency and how each method handles the different complexity of the models. The models assessed are N.O. Jensen (Jensen, 1983), Gaussian Bastankhah (Bastankhah and Porté-Agel, 2014) and Nygaard TurbOPark



(NTP) (Nygaard et al., 2020). For the latter, its derivation for FLOWERS is here introduced using a similar implementation of FLOWERS with NOJ wake model (LoCascio et al., 2024, 2025, 2022). The NTP is a top-hat model, which unlike NOJ, assumes the wake expansion rate to be proportional to the local turbulence intensity (TI) in the wake.

The paper is organized as follows. Section 2 introduces the different efficient AEP models used, as well as the baseline used and the optimization problem setup is presented. Section 3 presents and analyses the initial AEP and computational time benchmarking in comparison to the baseline. This is followed by an examination of the optimization results, which includes a study of the different efficacies, the extent to which the models were optimized, and efficiency, the speed at which they optimized. Finally the main results and takeaways are summarized in Section 4.

2 AEP models and optimization setup

The AEP can be estimated using a rectangular quadrature as,

$$\text{AEP} \approx 8760 \sum_{d=1}^{N_\theta} \sum_{u=1}^{N_u} \sum_{n=1}^{N_{WT}} P_n(u_{\text{eff},n}(\theta_d, u_u)) f(\theta_d, u_u), \quad (1)$$

where $f(\theta, u)$ is the joint probability of the wind direction and speed combination, θ and u , with $u_{\text{eff},n}$ indicating the effective wind speed experienced by wind turbine n , taking into account the wake effects from the upstream turbines, thus being a function of u and θ . P_n is the power output from turbine n and N_{WT} is the number of turbines in the farm. The computational cost of evaluating AEP in this manner is directly proportional to the number of bins used in the discretization, with higher amounts resulting in higher computation times, but also higher accuracy (Gonzalez-Rodriguez et al., 2025). Studies have shown that the use of 360 wind direction bins results in the most consistent and reliable results in layout optimization, although with significant computational costs (Feng and Shen, 2015).

The efficient AEP models used can be divided into four different categories, referring to the approach used to reduce the computational cost of evaluating the AEP: coarser binning, FLOWERS, probabilistic and stochastic. In addition, a baseline model is introduced and used as the ground truth, which is the model to be surpassed. All method in this category are computed using PyWake (Pedersen et al., 2023), an open source wake modeling tool used in several studies (Valotta Rodrigues et al., 2024; Quick et al., 2023; Criado Risco et al., 2024). A summary of the AEP models used in this study can be found in Table 1.

2.1 Baseline

A conventional model is used as the reference. The target layout will be the optimized output of this model, and the time taken to achieve it will be reduced where possible. This model consists of the rectangular quadrature in Eq. 1 using fine binning, consisting of 1° wind direction bins, and 1 m s^{-1} wind speed bins. Considering common turbine operating conditions of $3\text{--}25 \text{ m s}^{-1}$, the total number of flow simulations to carry out amounts to 8280. This aligns with the guidelines provided by Feng and Shen (2015) to obtain consistent and reliable results. This method is also used to evaluate the final optimized layouts to provide a fair comparison under each specific wake model.



2.2 Coarser binning

The following methods aim to reduce the computational cost by decreasing the total number of farm flow simulations required for each AEP evaluation with respect to the baseline, with 8280 simulations in total. The AEP methods in this category are Coarse WD, Average WS and Uniform C_T .

165 Reducing the number of wind direction bins, or equivalently, increasing the wind direction bin width from 1 to n degrees reduces the number of scenarios to be considered by a factor of n . Since the number of wind turbine interactions to be computed per scenario is the same regardless of the number of wind direction bins, the computation is theoretically also sped up by a factor of n . For the first approximation, referred to as Coarse WD, a value of 5° is selected as wind direction width, that is, 72 wind directions. The total number of wind farm simulations to compute is therefore 1656.

170 Similarly to how a high number of wind direction bins causes a high number of scenarios to be computed, a high number of wind speed bins has a similar effect. For this third approach, referred to as Average WS, a unique value for wind speed will be used for each wind direction bin considered, the average wind speed. The total number of scenarios to compute amounts to 360, one per wind direction in a fine binning. Therefore Eq. 1 is transformed into

$$\text{AEP} \approx 8760 \sum_{d=1}^{N_\theta} \sum_{n=1}^{N_{WT}} P_n(u_{\text{eff},n}(\theta_d, \overline{u_d})) f(\theta_d). \quad (2)$$

175 Lastly, software used for AEP evaluations use engineering models to quickly determine the wind speed downstream. One of the main inputs for these models is the thrust coefficient, C_T , to indicate the thrust, and varies non linearly with wind speed (Herbert-Acero et al., 2014). Being a function of wind speed, it is necessary to know the effective wind speed of the turbines upstream, and therefore the incident wind speeds are solved in sequence, creating a coupled system of interactions, limiting the computation speed. Using a uniform C_T makes this system redundant, thus speeding up calculations. The C_T used in this approach is therefore a function of the freestream wind speed, as well as using a similar procedure as the method above, using a single wind speed bin per direction bin. The thrust coefficient then becomes $C_T(\overline{u_\infty})$ rather than $C_T(u_{\text{eff}})$, therefore this method is referred to as Uniform C_T .

2.3 FLOW Estimation and Rose Superposition (FLOWERS)

185 All AEP methods introduced so far depend on discretizing the integral to make it computationally feasible, whether through a full discretized wind rose or a set of scenarios. FLOWERS seeks to advance this by solving the integral analytically instead of numerically. The method, first described in LoCascio et al. (2022) and later generalized in LoCascio et al. (2024), employs simplifications and the Fourier Transform to derive a tractable integral

The main simplification corresponds to the use of an average wind speed per wind direction bin, reducing the joint probability distribution of the atmospheric states, $f(\theta, u)$ into a single probability, that of the wind direction, $f(\theta)$. In addition, similarly to the Uniform C_T model explained above, the coupled systems of interactions between turbines to determine the local wind speed, used to obtain the power and thrust coefficients C_T and C_P , is incompatible with FLOWERS's analytical integration



approach. To solve this, the freestream wind speed is used to obtain C_T and C_P for all wind turbines, that is $C_T(u_\infty)$ and $C_P(u_\infty)$. Other simplifications include rotor treated as a single point, atmospheric homogeneity, and all turbines are assumed to be identical, as well as some minor mathematical simplifications and approximations, found in LoCascio et al. (2024) and Whittaker et al. (2025).

In the literature, FLOWERS can be found for the top-hat wake model described in Jensen (1983), with a thorough analysis in LoCascio et al. (2022, 2024, 2025), including the AEP model and the analytical gradients needed for gradient based optimization. In addition, it has been derived for BPA wake model (Bastankhah and Porté-Agel, 2014) with an initial analysis in Whittaker et al. (2025), introducing only the AEP model, but not the analytical gradients.

A new FLOWERS implementation is here introduced, using the top-hat wake model NTP (Nygaard et al., 2020) which, in addition to NOJ's approach, assumes the wake expansion rate to be proportional to the local turbulence intensity on the wake, referred to as NTP-FLOWERS. This approach improves the prediction of wake recovery over long distances, as the expansion of the wake slows down as the turbulence generated by the turbine decays. This makes the approach particularly relevant for large wind farms.

2.3.1 NTP-FLOWERS AEP

NTP wake model is a top-hat model, therefore, similarly to NOJ, the velocity deficit is uniformly distributed across the surface of the wake perpendicular to the axis of the wind turbine. Unlike NOJ, NTP's wake expansion is proportional to the local TI rather than a constant expansion rate. This local TI is the local and wake induced turbulence added in quadrature. The resulting wake deficit for a downstream distance of x is

$$\frac{\Delta u}{u_0}(x) = \left(1 - \sqrt{1 - C_T}\right) \left(\frac{D}{D_w(x)}\right)^2 \quad (3)$$

where D is the wind turbine rotor diameter and $D_w(x)$ is the wake diameter at a distance x . Whereas for NOJ $D_w(x)$ can be defined as $D + 2kx$, where k is a constant expansion coefficient, for the NTP model, after including the local turbulence, $D_w(x)$ becomes

$$D_w(x) = D + \frac{AI_0D}{\beta} \left(\sqrt{(\alpha + \beta x/D)^2 + 1} - \sqrt{1 + \alpha^2} - \ln \left[\frac{(\sqrt{(\alpha + \beta x/D)^2 + 1} + 1)\alpha}{(\sqrt{1 + \alpha^2} + 1)(\alpha + \beta x/D)} \right] \right), \quad (4)$$

where I_0 is the background atmospheric turbulence and α and β are auxiliary positive variables defined as $\alpha = c_1 I_0$ and $\beta = c_2 I_0 / \sqrt{C_T(u_0)}$ with $c_1 = 1.5$ and $c_2 = 0.8$ as given in IEC (2010). Finally A is a calibrated expansion coefficient, which recommended value for offshore wind farms is 0.6 (Nygaard et al., 2020). Such a wake model implementation into FLOWERS would not be tractable given the complexity and nature of the equation. Nevertheless, the only difference with respect to the NOJ model used in LoCascio et al. (2024) is the wake radius, no longer a function of the constant wake coefficient k . With the goal of following a similar implementation, Eq. 3 is transformed into




Figure 1. Visual representation of effective expansion wake coefficient k_{eff} , representing the wake shape obtained using Eq. 4 and that using k_{eff} for an arbitrary downstream distance x .

$$\frac{\Delta u}{u_0}(x) = \frac{1 - \sqrt{1 - C_T}}{2k_{\text{eff}}(x)/D + 1} \quad (5)$$

where $k_{\text{eff}}(x)$ represents the effective wake expansion coefficient at a downstream distance x which yields an equal wake diameter to that obtained using Eq. 4 in the NTP wake model. This $k_{\text{eff}}(x)$ is obtained for every pair of wind turbines i and j in the wind farm, and is defined as

$$k_{\text{eff}}(x) = \frac{1}{2} \frac{\Delta D_W}{\Delta x} = \frac{D_w(x) - D}{2x} \quad (6)$$

where x is the downstream position from the wake generating turbine. A visual representation of the effective wake expansion coefficient is displayed in Fig. 1. This adaptation matches the formulation in NOJ-FLOWERS, enabling the exact same derivation for the AEP model. The presence of C_T in the β factor in Eq. 4 makes it impossible to include it among the discrete inputs as a function of wind directions that are transformed into continuous form via the discrete Fourier transform. For this reason, a universal thrust coefficient, or $\overline{C_T}$, is used, similar to the one introduced for Gaussian-FLOWERS in Whittaker et al. (2025), based on weight averaging the thrust coefficient by expected power production, defined as

$$\overline{C_T} = \sum_{i=1}^{N_\theta} f_i C_T(U_{0,i}) \frac{C_P(U_{0,i}) U_{0,i}^3}{\sum_I C_P(U_{0,i}) U_{0,i}^3}. \quad (7)$$

Using this effective expansion wake coefficient, the FLOWERS AEP implementation is close to that found in LoCascio et al. (2024), with the substitution of k with k_{eff} .

2.3.2 NTP-FLOWERS AEP gradients

The derivation of AEP gradients follows LoCascio et al. (2024) with the additional gradient terms from the introduction of k_{eff} as derived here. Knowing that k_{eff} in Eq. 6 is only a function of the downstream distance x , or in the relative polar coordinate



system being used, a function of r_{ij} , only terms sensitive to changes in r_{ij} need to be computed, that is $\overline{\Delta p}_{ij}(r_{ij}, \gamma_{ij})$, where γ_{ij} is the relative polar angle. The derivative of $\overline{\Delta p}_{ij}(r_{ij}, \gamma_{ij})$ with respect to r_{ij} then becomes

$$240 \quad \frac{\partial \overline{\Delta p}_{ij}}{\partial r_{ij}} = \left. \frac{\partial \overline{\Delta p}_{ij}}{\partial r_{ij}} \right|_{\text{NOJ}} + \frac{\partial \overline{\Delta p}_{ij}}{\partial k_{\text{eff}}} \frac{\partial k_{\text{eff}}}{\partial r_{ij}} + \frac{\partial \overline{\Delta p}_{ij}}{\partial \theta_{ij}^{(c)}} \frac{\partial \theta_{ij}^{(c)}}{\partial k_{\text{eff}}} \frac{\partial k_{\text{eff}}}{\partial r_{ij}} \quad (8)$$

where $\left. \frac{\partial \overline{\Delta p}_{ij}}{\partial r_{ij}} \right|_{\text{NOJ}}$ refers to those gradients already derived in LoCascio et al. (2024). $\theta_{ij}^{(c)}$ is the critical angle representing the wake boundary, also introduced in LoCascio et al. (2024), and is a function of r_{ij} and $k_{\text{eff}}(r_{ij})$. For simplicity, the derivation of the different terms in Eq. 8 is omitted. The sensitivity of $\overline{\Delta p}_{ij}$ to changes in k_{eff} is

$$\begin{aligned} \frac{\partial \overline{\Delta p}_{ij}}{\partial k_{\text{eff}}} = & \frac{a_0 \theta_{ij}^{(c)}}{(2k_{\text{eff}} r_{ij} + 1)^2} \left[\frac{8\pi^2 r_{ij} (\theta_{ij}^{(c)})^2}{3(2k_{\text{eff}} r_{ij} + 1)^2} \right. \\ & \left. - \frac{4r_{ij}}{2k_{\text{eff}} r_{ij} + 1} \left(1 + \frac{8\pi^2 k_{\text{eff}} r_{ij} (\theta_{ij}^{(c)})^2}{3(2k_{\text{eff}} r_{ij} + 1)} \right) \right] \\ & + \sum_{m=1}^{M-1} \left[\frac{a_m \cos(2\pi m \theta_{ij}) + b_m \sin(2\pi m \theta_{ij})}{\pi m (2k_{\text{eff}} r_{ij} + 1)^2} \left(\frac{2r_{ij} T_m}{m^2 (2k_{\text{eff}} r_{ij} + 1)^2} \right. \right. \\ & \left. \left. - \frac{4r_{ij}}{2k_{\text{eff}} r_{ij} + 1} \left(\sin(2\pi m \theta_{ij}^{(c)}) + \frac{2k_{\text{eff}} r_{ij} T_m}{m^2 (2k_{\text{eff}} r_{ij} + 1)} \right) \right) \right], \end{aligned} \quad (9)$$

245 where

$$T_m = \left[(2\pi m \theta_{ij}^{(c)})^2 - 2 \right] \sin(2\pi m \theta_{ij}^{(c)}) + 4\pi m \theta_{ij}^{(c)} \cos(2\pi m \theta_{ij}^{(c)}). \quad (10)$$

a_m and b_m correspond to the Fourier terms obtained from the discrete Fourier transform of the discrete wind rose components found in LoCascio et al. (2024). Moreover, the derivative with respect to r_{ij} of the k_{eff} is

$$\frac{\partial k_{\text{eff}}}{\partial r_{ij}} = \frac{AI_0}{2r_{ij}} \left(\frac{\alpha + \beta r_{ij}}{\sqrt{(\alpha + \beta r_{ij})^2 + 1} + 1} + \frac{1}{\alpha + \beta r_{ij}} \right) - \frac{k_{\text{eff}}}{r_{ij}}, \quad (11)$$

250 the derivative of $\overline{\Delta p}_{ij}$ with respect to $\theta_{ij}^{(c)}$ is

$$\begin{aligned} \frac{\partial \overline{\Delta p}_{ij}}{\partial \theta_{ij}^{(c)}} = & \frac{a_0}{(2k_{\text{eff}} r_{ij} + 1)^2} \left[1 + \frac{8\pi^2 k_{\text{eff}} r_{ij} (\theta_{ij}^{(c)})^2}{2k_{\text{eff}} r_{ij} + 1} \right] \\ & + \sum_{m=1}^{M-1} \frac{2(a \cos(2\pi m \theta_{ij}) + b \sin(2\pi m \theta_{ij})) \cos(2\pi m \theta_{ij}^{(c)})}{(2k_{\text{eff}} r_{ij} + 1)^2} \left(1 + \frac{2k_{\text{eff}} r_{ij} (2\pi m \theta_{ij}^{(c)})^2}{m^2 (2k_{\text{eff}} r_{ij} + 1)} \right) \end{aligned} \quad (12)$$

and finally the sensitivity of $\theta_{ij}^{(c)}$ to changes in k_{eff} is defined as



$$\frac{\partial \theta_{ij}^{(c)}}{\partial k_{\text{eff}}} = \frac{1}{2\pi \sqrt{1 + k_{\text{eff}}^2 - \frac{1}{4r_{ij}}}} \frac{\sqrt{1 + k_{\text{eff}}^2 - \frac{1}{4r_{ij}} - \frac{k_{\text{eff}}}{2r_{ij}}}}{1 + k_{\text{eff}}} \quad (13)$$

In addition to the equations introduced in LoCascio et al. (2024), Eq. 9, Eq. 10, Eq. 11, Eq. 12, Eq. 13 represent the gradients of NTP-FLOWERS with respect to the turbine coordinates.

2.4 Probabilistic approach

The models presented so far use the whole wind rose indiscriminately, regardless of their weight and relevance in the AEP estimation. This category introduces probabilistic AEP approaches, meaning that the wind conditions used to estimate AEP are used based on the probabilistic weight from the wind distribution. The two methods falling into this category are Bayesian Quadrature (BQ) and Optimal-design Weighted Quadrature (OWQ).

The BQ method for AEP estimation, introduced in (King et al., 2020), uses a trained Gaussian Process estimate to compute AEP with a reduced number of wind farm flow simulations using Bayesian Quadrature. The resulting AEP equation is

$$\text{AEP} \approx 8760 \mathbf{w}^T \mathbf{K}^{-1} \mathbf{P} \quad (14)$$

where $\mathbf{w}$ is the kernel-weighted probability integrals, converting the AEP estimation into a simple point weighted sum over the discrete samples, in this case, the wind conditions sampled. $\mathbf{K}$ is the covariance matrix formed by evaluating the Matérn kernel at all sample points, with hyperparameters fitted via the Gaussian Process regression on the sampled flow model outputs. $\mathbf{P}$ is the power output coming from the wind farm flow modeling simulations at the sampled points, computed using PyWake.

The use of a Gaussian Process opens the possibility of obtaining the uncertainty of the AEP estimation, meaning that the set of wind conditions to use in the Bayesian Quadrature can be optimized to minimize the uncertainty. The unconstrained optimization is

$$\max_{s_i} \mathbf{w}^T \mathbf{K} \mathbf{w}, \quad (15)$$

where s_i is the wind direction and speed combinations, i being the number of samples that are desired to be optimized. The recommended number of samples by King et al. (2020) is 200, but 360 was here used, so that it can be compared with other methods using 360 atmospheric combinations (Average WS, Uniform C_T or FLOWERS). In order to avoid being trapped in local minima, a total number of 10 random starts are optimized. An analysis on the choice of number of random starts is provided in Appendix A. Given that none of the components involved in the optimization are a function of the wind farm's layout, these optimal points can be obtained beforehand, and be used for any layout, as long as the atmospheric site conditions remain constant.

Regarding the AEP gradients with respect to the turbine positions, the only component of Eq. 14 which depends on the specific layout is $\mathbf{P}$. Therefore, the gradients become



$$\frac{\partial AEP}{\partial x} = 8760 \mathbf{w}^T \mathbf{K}^{-1} \frac{\partial \mathbf{P}}{\partial x}, \quad (16a)$$

$$\frac{\partial AEP}{\partial y} = 8760 \mathbf{w}^T \mathbf{K}^{-1} \frac{\partial \mathbf{P}}{\partial y}, \quad (16b)$$

Here we propose the novel method Optimal-design Weighted Quadrature (OWQ), which uses the same optimized sampling strategy obtained for Bayesian Quadrature, to compute AEP by summing the power contribution from each optimally-selected wind speed and direction combination, weighted by its normalized occurrence probability. The used number of samples in this case is also 360. AEP can then be obtained using

$$AEP \approx 8760 \sum_{i=1}^{N_i} \sum_{n=1}^{N_{WT}} P_n(u_{\text{eff},n}(\theta_i, u_i) \tilde{f}(\theta_i, u_i), \quad (17)$$

where N_i is the number of samples, and $\tilde{f}(\theta_i, u_i)$ is the normalized probability of occurrence of the wind speed and direction combination i , defined as

$$\tilde{f}(\theta_i, u_i) = \frac{f(\theta_i, u_i)}{\sum_{i=1}^{N_i} f(\theta_i, u_i)} \quad (18)$$

2.5 Stochastic sampling of wind rose

The models presented so far are deterministic, using the entire wind rose data, or base their sampling strategy on it, to estimate AEP. This category introduces a stochastic AEP approach, meaning that the wind conditions used to estimate AEP are randomly drawn from the wind distribution, resulting in different AEP estimation values per evaluation, or in the method here introduced, gradients. Unlike the approaches seen so far, Stochastic Gradient Descent (SGD) is not an objective function, but an optimization algorithm, commonly used in machine learning when selecting neural network weights. This last method uses its adaptation to the WFLO problem developed by Quick et al. (2023). Using the Monte Carlo estimation of the AEP to compute its associated gradient by using the equation found in Eq. 19, and an adaptation of the Adam algorithm (Kingma and Ba, 2017), this coarse gradient is enough to guide SGD towards the nearest local optima. Monte Carlo's approximation of the gradients can be defined as

$$\frac{d}{ds} AEP \approx 8760 \frac{1}{K_{MC}} \sum_{k=1}^{K_{MC}} \sum_{n=1}^{N_{WT}} \frac{d}{ds} P_n(u_{\text{eff},n}(\theta_k, u_k)) \quad (19)$$

where $K_{MC} = 50$ is the number of samples drawn from the probability distributions. S is the set of coordinates of the wind turbines. The original Adam algorithm is intended for unconstrained optimization, therefore in Quick et al. (2023), a constraint aggregation to a penalty term with units that are consistent with the objective is introduced. The purpose of this penalty term is to ensure that both the penalty and the AEP gradients initially have a similar magnitude, enabling the algorithm to explore the



Category	AEP method	N_{WD}	N_{WS}	Total flow simulations	Notes
Baseline	Baseline	360	23	8280	Combinations
Coarser binning	Coarse WD	72	23	1656	Combinations
	Average WS	360	1	360	Combinations
	Uniform C_T	360	1	360	Combinations
FLOWERS	FLOWERS	360	1	–	Fourier
Probabilistic	BQ	360	360	360	Pairs of samples
	OWQ	360	360	360	Pairs of samples
Stochastic	SGD	50	50	50	Pairs of samples

Table 1. Summary of AEP methods. N_{WD} and N_{WS} indicate the number of wind direction and speed bins, respectively. *Combinations* refers to those models where the different atmospheric conditions are combined, *Pairs of samples* to those where the conditions are paired, and *Fourier* for the case where the wind rose is transformed using Fourier Transform.

design space. However, as the optimization process continues, the penalty gradients will become dominant over the AEP ones, ensuring a viable solution. The constraint multiplier is responsible for this increase in the dominance of the penalty gradients, growing as the learning rate decreases.

This learning rate is the main parameter in SGD. It controls the step size at each iteration, also determining how much the turbine positions are adjusted in response to the estimated gradient. A well-chosen learning rate balances convergence speed and solution quality, avoiding overshooting or slow progress. It is scheduled to decay as the optimization progresses, controlled by the learning rate scheduler, dependent on the initial and final learning rates, $D/5$ and 0.1 respectively in this study, corresponding to those recommended in Quick et al. (2023). The last parameter that needs to be considered is the number of iterations. When it is increased, larger final AEP values are obtained, with a computational cost that grows proportional to the optimization iterations (Quick et al., 2023). A total of 2000 iterations are considered here, which has been proven to achieve consistent results (Quick et al., 2023). An initial analysis of this method showed that lower iteration counts resulted in proportionally faster optimization times, but with the layout found presenting poorer AEPs. For the Adam algorithm momentum constants, the values used are those found in Quick et al. (2023), corresponding to $\beta_1 = 0.1$ and $\beta_2 = 0.2$.

2.6 WFLO setup

The performance of the different AEP methods studied here is tested as objective functions in a series of WFLO problems, encompassing a variety of farm sizes, wind roses and constraints. In order to test different conditions, two convex (C) and two non-convex (NC) boundaries are utilized. A non-convex boundary is characterized by inward-curving edges, in contrast to a convex boundary where straight lines between interior points can exit the region (i.e. concave polygon). While gradient based algorithms can efficiently and easily optimize convex boundaries, they may struggle and require more iterations to optimize non-convex regions (Thomas et al., 2023). Moreover, distancing constraints (DC), limiting the minimum distance between turbines within the wind farm, further reduce the efficiency of the optimization, scaling very poorly with the number



of turbines, given the increase in turbine to turbine distances to consider. For this reason, such constraints will be included into two cases, with the other two being unconstrained in this sense. This is a reasonable decision, given that turbines tend to separate in the initial stages of optimization to minimize wake losses while still meeting the minimum safety standards (Valotta Rodrigues et al., 2024). To verify this claim, a minimum safety distance of $2D$ is verified for each optimized layout. This simplification is also for the sake of testing the AEP models's ability to handle large wind farms without making the optimization time intractable, given that the optimizer's handling of the distancing constraints is not the main goal of this study. Real wind farms are used here, where the ones using convex boundaries are a convex-hull of the actual wind farm region, therefore not fully resembling the actual setup. The wind farms used and their configurations are presented in Table 2, with their boundaries shown in Fig. 2, and the wind roses in Fig. 3, where Sofia and Hornsea 2 do not correspond to the original ones, but to other locations, so as to have a larger variety of wind conditions. The wind roses are obtained from the New European Wind Atlas (NEWA), a publicly available dataset produced with the Weather Research and Forecasting (WRF) mesoscale model that provides 30 years of wind data (1989–2018) (Hahmann et al., 2020; Dörenkämper et al., 2020). The optimization problem is therefore

$$\text{340} \quad \underset{x,y}{\text{Maximize}} \quad AEP(x,y)$$

Subject to:

$$\begin{aligned} & \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2} \geq N_D D, \quad \forall i, j \in \{1, 2, \dots, N_{WT}\}, i \neq j, \quad (\text{when applicable}) \\ & (x_i, y_i) \in \mathcal{P}, \quad \forall i \in \{1, 2, \dots, N_{WT}\}, \end{aligned}$$

where $N_D D$ indicates the minimum distance between turbines in turbine diameters, for those cases in which such a constraint is applied. $\mathcal{P}$ represents the specific wind farm boundaries. Gradient based algorithms, which usually minimize, are usually known as local search algorithms since despite searching for the global minima, they very often get stuck, converging to suboptimal solutions (Herbert-Acero et al., 2014; Guirguis et al., 2016; Thomas et al., 2023), showing a great sensitivity to the initial conditions. In consequence, WFLO studies using gradient based algorithms rely on multi-start approaches (Quick et al., 2023; LoCascio et al., 2025; Criado Risco et al., 2024), in which a set of optimization problems are solved with random initial conditions. In this case, a total of 30 multi-starts are used. Lastly, all turbine coordinates are normalized with respect to the maximum respective coordinate possible given the different boundaries, so that all coordinates are between 0 and 1. The optimization problem is carried out using the TOPFARM framework (Pedersen et al., 2025), used in a variety of wind farm optimization studies (Wolf Ciavarrá et al., 2022; Riva et al., 2020; Criado Risco et al., 2024). The gradients are obtained using PyWake's automatic differentiation for all of the AEP models, with the exception of FLOWERS, in which analytically derived gradients are used. The reported times are all obtained using node clusters from DTU Sophia supercomputer (Technical University of Denmark, 2019), in single node operation, with each node being composed of 2×AMD EPYC 7351 CPUs, at 2.9 GHz, with 16 cores and 128 GB of RAM.

For the different wake models use, the respective wake expansions coefficients, for those where it is assumed constant, are $k_{NOJ}=0.05$ and $k_{BPA} = 0.03$, corresponding to typical off-shore values (Bastankhah and Porté-Agel, 2014). The atmospheric



Name	Turbines	Boundary	Distancing	WT Power
Hornsea 3	231	Convex	No	14 MW
Sofia	100	Convex	Yes	14 MW
Thor	72	Non-convex	No	14 MW
Hornsea 2 – Phase 1	55	Non-convex	Yes	8 MW

Table 2. Wind farm configuration characteristics

3




Figure 2. Boundaries for the different wind farms being studied, not scaled (C=Convex, NC=Non-convex, DC=Distancing constraints).

360 TI , I_0 , used for the NTP wake model is assumed to be equal to 0.1. To provide a comparison of the optimized layouts, for each different wake model, these are evaluated using the baseline AEP model regardless of the model that has been used for the optimization, considered the highest fidelity model in this study.

When computing the AEP, the contributions from the different scenarios – the combinations of wind speed and direction – are summed. As these are independent of each other, parallel computation is possible with the use of more than one CPU, where
 365 several scenarios are evaluated simultaneously, considerably speeding up the overall computation. This is only possible when wind distributions are discretized and flow simulations are independent of each other, that is all methods except FLOWERS. The maximum number of CPUs is limited due to the overhead times, where a point is reached beyond which increasing the number of CPUs does not result in a useful reduction of computational time (Valotta Rodrigues et al., 2024). In order to determine the optimal number of CPUs that provide a balance between computational speed and resources used an analysis is
 370 presented in Appendix B, justifying the use of 8 CPUs for all methods except SGD and FLOWERS which use a single CPU.

3 Results

As a means to study the different methods capability to evaluate AEP and its corresponding gradients, as well as its computational efficiency, before testing them in an optimization framework an initial comparison is carried out. With these results in mind, the WFLO optimization outcomes for the series of case studies previously introduced are subsequently analyzed.




Figure 3. Wind roses for the different wind farms studied (C=Convex, NC=Non-convex, DC=Distancing constraints).

3.1 AEP method benchmarking – AEP

The difference with respect to the baseline is displayed in Fig. 4 for AEP estimation and computational time, for all three wake models considered. These results are obtained for regular grid wind farms of increasing number of wind turbines, using the wind rose corresponding to Hornsea 3, found in Fig. 3. It is important to note that the time results presented are comparing CPU hours. The SGD algorithm does not perform AEP evaluations in the implementation used here, so it is excluded from this analysis. A model may be very suitable to evaluate a wind farm's AEP, but might not prove useful for layout optimization. Besides efficiently computing AEP, WFLO requires an objective function that properly reflects the design space, in this case, the wind farm physics and turbine interactions. In other words, the accuracy of a model's determination of the AEP with respect to the baseline does not by itself determine whether it is suitable for layout optimization. Likewise, a model presenting a large error does not render it useless for layout optimization, as it may still capture the physics in a way that can guide the optimizer towards the optimal solution.

Regarding the AEP estimations, some models are more accurate than others, with Coarse WD, BQ and OWQ presenting the closest values to the reference, with differences smaller than 5 % regardless of the wind farm size and wake model. The fact that Coarse WD barely differs from the baseline reflects that coarser wind direction bins barely affect the accuracy of the calculations. How this affects how well wind farm interactions are captured will be seen later. BQ and OWQ, both probabilistic methods, are capable of achieving values that are very close to the reference values with as little as 360 wind farm flow simulations, aligning with the results obtained by King et al. (2020). On the other hand, methods using an average wind speed per wind direction bin, Average WS, Uniform C_T and FLOWERS, fail to accurately estimate AEP, with the two former ones resulting in overestimations for all wake models. For the case of FLOWERS, the error increases with the number of turbines, and while NOJ-FLOWERS results in underestimations, BPA-FLOWERS shows overestimations, matching the conclusions reached in LoCascio et al. (2025) and Whittaker et al. (2025) respectively. NTP-FLOWERS, being a very similar wake model to that of NOJ-FLOWERS, in addition to the parallel derivation, presents an underestimation of the AEP.




Figure 4. AEP and its computational time (CPU hours) difference of the different AEP methods with respect to the baseline using Hornsea 3 wind rose from Fig. 3.

All AEP models considered here are faster than the baseline. The fastest of the methods is FLOWERS, achieving a time reduction of more than 99 % for all wind farm sizes and wake models, due to its continuous approach. The remaining methods all achieved time reductions of around 60 to 80 %, and the use of a single wind speed per direction bin in Average WS and Uniform C_T is noticeable: they are slightly faster than simply using coarser wind direction bins, Coarse WD. However, the effect of implementing a Uniform C_T is negligible in larger wind farms. Finally, both probabilistic methods, BQ and OWQ achieve a considerable reduction of the computational time, of the order of 70-90 % depending on the wake model, which in addition to the small errors, makes them a very promising AEP model, at least when it comes to evaluating AEP. Notably, for smaller wind farms BQ is slower than the baseline, reflecting that it does not take as much advantage of parallelization as other methods. However, it is demonstrated that all AEP methods implemented here fulfill the main objective of this study, achieve significant computational time reductions in AEP computation.

The accuracy of FLOWERS, and especially NOJ-FLOWERS, has been found to be dependent on the wind rose used, with strongly unimodal wind roses resulting in larger errors (LoCascio et al., 2024). This increased error is attributed to the use of an average wind speed and uniform thrust coefficient. For this reason, an AEP benchmarking is carried out to see how the AEP predictions are affected when using strongly unimodal wind roses such as the one used for Sofia. The results obtained are displayed in Fig. 5, where the increased error of FLOWERS can be observed, increasing to 40 % difference for NOJ- and NTP-FLOWERS, a significant increase with respect to the 20 % from the previous analysis. On the other hand, BPA-FLOWERS




Figure 5. AEP difference of the different AEP methods with respect to the baseline using Sofia wind rose from Fig. 3.

shows a substantial decrease, demonstrating its superiority over the other wake model counterparts. A similar behavior is observed for Average WS and Uniform C_T , where they displayed a slightly increased accuracy. The most accurate models in the previous analysis, BQ and OWQ, with consistent differences of less than 5 %, have now surged to errors of 40 and 60 % respectively, clearly demonstrating their difficulty to adapt to this sort of wind roses.

3.2 AEP method benchmarking – Gradients

Besides requiring hundreds or thousands of AEP evaluations, gradient based optimizers also need gradients to guide them, evaluated a similar number of times (Valotta Rodrigues et al., 2024). For this reason, the time needed to evaluate such gradients plays a key role in the overall optimization time. Figure 6 shows the difference in computational time with respect to the reference method, for an increasing number of turbines. Results are very similar to those found in the AEP benchmarking, where all gradient evaluations are faster than the baseline with the exception of BQ, which takes little advantage of parallelization, as was already noted in Appendix B, making OWQ faster than its probabilistic counterpart. Once again, even when compared to parallelized methods, FLOWERS computes the gradients in a time that can be considered negligible compared to the other methods. Similar to what was seen for the AEP evaluations, Uniform C_T seems to outperform Average WS in all wake models, owing to most importantly taking better advantage of parallelization. In this analysis, SGD emerges as a model with extremely rapid gradient evaluation times thanks to only solving 50 flow models, approaching reductions similar to those of FLOWERS, despite using a single CPU.

3.3 Optimization results

The results displayed in Fig. 7 show the difference, in percentage, with respect to the baseline of the median final optimized AEP and total computational time for each wind farm, wake model and AEP method used as objective function. The optimized layouts are evaluated using the baseline method for comparison purposes. The median is used to avoid the effect of outliers




Figure 6. AEP gradients computational time difference of the different AEP methods with respect to the baseline using CPU hours.

skewing the results. Red colors indicate poorer performance in the respective category compared to the baseline, while green colors reflect improved performance. First an analysis of the final AEP will be provided (efficacy comparison), to then move
 435 onto the optimization time (efficiency comparison) and a combined analysis. Table C1 contains the median and standard deviation of optimizer iterations across wake models, by wind farm and AEP model.

3.3.1 Efficacy comparison - Optimized AEP

The final optimized AEP – when evaluated using the same wake model and AEP method, in this case the corresponding wake model and the baseline AEP method – provides a measurement of how well the method optimizes a wind farm layout. In the
 440 results shown in Fig. 7, it is immediately apparent that in general, all simplified AEP methods reach worse layouts than the baseline, with SGD, FLOWERS – with the exception of Sofia – and BPA-Coarse WD, as those attaining the closest results to the baseline. FLOWERS, despite its numerous simplifications obtains some of the best results, due to its continuous approach results in a very smooth design space, sometimes even surpassing the baseline, regardless of the considered wake model. Nevertheless, as mentioned in LoCascio et al. (2024), the assumptions may lead to inaccuracies in strongly unimodal wind
 445 roses such as Sofia’s, due to the method not properly capturing the physics within the farm, and therefore presenting the worst results of all when it comes to such wind farm, with AEP values $\sim 2\%$ lower. A similar number of iterations with respect to other wind farms confirms that the issue is related to how the model captures the physics, providing gradients that guide the optimizer towards a failed optimized layout, rather than simply becoming stuck in a suboptimal layout. In addition, the low standard deviations in the iteration count suggest a very smooth design space in which the optimizer will not have issues
 450 with multi-modality, agreeing with those results found in LoCascio et al. (2025). For the case of SGD, these consistent results with virtually no difference with the baseline align with those found in Quick et al. (2023), proving once again the strength of the algorithm in WFLO, showing no difficulty under any of the cases considered here. The last top performer, Coarse WD, despite its coarse direction binning managed to obtain the same results as the baseline, but only for the BPA wake model.




Figure 7. Optimized AEP and optimization time for the different wake models, wind farms and AEP methods. Results are displayed as the difference in percentage, with respect to the baseline, of the median value out of the 30 random starts. (C=Convex, NC=Non-convex, DC=Distancing constraints)

This can be easily explained because of the gaussian shaped wake which avoids the over sensitiveness to lateral displacement present in top-hat models (Wang et al., 2024), resulting in a very similar design space to that of using 360 direction bins in the baseline. On the other hand, when using this coarser direction binning in top-hat models, given the overestimation at the wake edges and uniform distribution of the velocity deficit (Wang et al., 2024), the resulting design space is too coarse, resulting in substantially suboptimal final layouts, with layouts yielding around 0.2 and 0.4 % lower AEPs. This conclusion can be validated with the iteration count, where for this method, Gaussian results yield an average of 98 iteration across wind farms – barely any difference with respect to the baseline with 96 iterations. On the other hand, top-hat wake models NOJ and NTP



result in an average of 134 and 140 respectively, around 30 % more than the finer binning from the baseline, meaning that the optimizer had issues navigating the problem for the latter and therefore resulting in poorer results, although this can be a consequence of other optimization parameters.

Both probabilistic methods, BQ and OWQ, despite computing AEP accurately as shown in Section 3.1, where they presented the smallest errors out of all methods considere, show rather poor results when used in layout optimization, with negative differences between 0.2 and 1 %. This behavior could be related to the use of only 360 samples, but seeing that both Average WS and Uniform C_T , also using 360 samples, do not present such poor outcomes, sample count alone does not appear to be a sufficient explanation. A more likely cause lies in how the BQ sample locations are chosen. The optimal set of points is obtained by minimizing the integrated variance of the AEP estimate with respect to the wind speed and direction joint probability density functions (PDF) and the kernel covariance function alone, independent of the power response surface itself. This is by design, since it allows the same sample set to be reused across layouts without recomputation, but it also means the points concentrate where inflow is probable rather than where AEP is most sensitive to turbine position, such as narrow sector partial wake conditions. These conditions contribute little to the AEP integral, so the estimate remains accurate, but they are precisely the conditions a layout optimizer needs resolved to obtain a useful gradient. Average WS and Uniform C_T , in contrast, retain full unimodal coverage even at 360 samples, so sensitivity information is preserved across the wind rose regardless of probability weighting. Nothing in the BQ sampling strategy ensures that the selected points capture enough of this turbine interaction physics to effectively guide the optimizer, even though they are more than sufficient to guide the AEP estimate itself. Nonetheless, the fact that OWQ presents better optimized layouts, with values ranging from -0.21 to -0.57 % indicates that its representation of the design space is more accurate than that of BQ, with values ranging from -0.18 to -0.95 %. However, with only slightly lower iteration counts on average, but more than double the average standard deviation in the iteration count, it may be suggested that the design space provided by OWQ is more chaotic and more challenging for the optimizer.

Finally, the two methods using an average wind speed per wind direction bin, Average WS and Uniform C_T , present very similar results with virtually the same iteration counts, both median and standard deviation, with the exception of the non-convex boundaries where they yield significantly different and sparser results. This demonstrates that the use of a unique value of thrust coefficient for all wind turbines barely affects the overall outcome. However, both of these methods present a sensitivity to the wind rose being used, where the use of more unimodal wind roses, like those found in Sofia and Hornsea 2, result in poorer optimization performances with notably bad layouts in the former, a direct consequence of the non linearity of the power and thrust curves. From these, it can be seen that all methods studied here, with the exception of SGD and Coarse WD, suffer from the strongly unimodal wind rose found in Sofia.

3.3.2 Efficiency comparison - Optimization time

The purpose of this study was to compare efficient AEP methods for layout optimization, where the computational efficiency is assessed by means of the optimization time, defined as the time taken to run the optimization. The results found in Fig. 7, showing the difference in percentage with the baseline of the median of all multi-starts, demonstrate that all AEP methods



495 took shorter than the baseline to run the optimization, with the exception of SGD and two cases for BQ. The high optimization times for SGD can be explained by the high iteration count needed to obtain consistent results, a total of 2000 (Quick et al., 2023), around 20 times more than the number of iterations needed by the baseline on average. Consequently, it takes between 2 and 15 times longer than the baseline depending on the wake model being used, NTP-SGD being the fastest one on average compared to the baseline, without any clear pattern on the use of convex, non-convex boundaries nor distancing constraints. It is important to note that a different hyperparameter selection for for this algorithm may potentially results in lower optimization times, maintaining a similar quality of the obtained layouts. On the opposite side of the spectrum, we find FLOWERS, with reductions between 90 and 99 % in the optimization time, with the exception of Sofia for the BPA-FLOWERS. This decrease, resulting in the fastest optimizations observed in this study, thanks to its fast evaluations, as seen in Section 3.1, and the smooth design space, proven to be advantageous in gradient based optimization (LoCascio et al., 2024, 2025). Together with the high quality layouts found, FLOWERS is therefore the top performer of this study, albeit with clear weaknesses when using strongly unimodal wind roses, where the performance is substantially poorer than the baseline.

For the probabilistic methods, BQ and OWQ, the time decrease is not as large as for the other methods, with BQ even presenting increased optimization times in the Sofia wind farm, again signaling the difficulty of the model to adapt to unimodal roses. A longer optimization time is commonly a consequence of more iterations required by the optimizer – reflecting a poor representation of the design space – or slower AEP and gradients evaluations. Another likely reason for this is the fact that OWQ takes a greater advantage of parallelization, resulting in faster evaluation times. Nonetheless, it might also be a result of a better reflection of the problem’s physics, due to the samples being summed based on probability of occurrence rather than weights, with 150 iterations on average, compared to the 175 needed by BQ.

It was previously concluded that the two methods Average WS and Uniform C_T demonstrated that the use of a single thrust coefficient for all turbines barely affected the optimization. From results obtained in Fig. 7, the computational time differs significantly for the more complex cases where non-convex boundaries are needed. The iteration counts are similar, without any difference that could justify these computational time differences. On the other hand, the computational time using parallelization, as reported in Section 3.1, reveals that Uniform C_T takes slightly less time to evaluate AEP and gradients, and therefore contributing to the optimization time differences perceived.

520 Finally, having obtained remarkable results when applied with the BPA wake model, Coarse WD achieves reductions between 40 and 70 %, therefore showing that this combination of wake model and AEP method is a top performer in WFLO, owing to the smoothening effect of using a gaussian shaped model. Nonetheless, this needs further exploration to determine how coarse the bins can be without the reduction in number of bins becoming detrimental for the model.

4 Conclusions

525 This research aimed to provide an investigation of different computationally efficient AEP methods, comparing their performance and potential uses in WFLO. The AEP models consisted of reducing the number of wind direction and speed bins; using a Uniform C_T ; the novel FLOWERS approach, which uses assumptions, simplifications and Fourier transform to analytically



solve the integral; and the probabilistic BQ method which uses a trained GP to compute AEP with a reduced number of wind farm flow simulations, with the addition of a similar probabilistic derivation, OWQ. In addition, the SGD algorithm was included in the study, given its simplified wind rose sampling methodology to obtain the gradients. All simplified models were studied for NOJ, BPA and NTP wake models. For FLOWERS, the derivation for the latter wake model was introduced. They were tested on four different WFLO scenarios, encompassing convex and non-convex boundaries, different number of wind turbines, distancing constraints and a variety of wind roses to ensure that a broad and representative spectrum of cases were covered. A multi-start approach was employed to expand the exploration of the optimization space, given the multimodality of the problem being considered.

First, an analysis of the AEP estimation compared to the baseline or reference point, consisting of 1° and 1 m s^{-1} direction and speed bins in the rectangular quadrature method, was carried out. It showed how some models, especially Coarse WD, BQ and OWQ, with differences of less than 5 % managed to properly estimate AEP, while the rest resulted in significant errors, particularly FLOWERS as result of its numerous assumptions and simplifications. All methods, with the exception of Coarse WD, were shown to be highly sensitive to a strongly unimodal wind rose, even in initially most accurate models, BQ and OWQ, whose accuracy decreased the most out of all the AEP simplifications. This demonstrated that the optimized sampling strategy implemented in these two methods failed to adapt to this sort of wind conditions. It is important to mention that all methods studied here were developed with the goal of reducing the optimization time in a gradient based WFLO problem, in which the objective function's ability to reflect wind farm interactions to provide useful gradients is more important than accurately computing AEP. Regarding the computational time all methods were more efficient than the baseline, with FLOWERS obtaining a speed up of more than 95 % for all wind farm sizes. Other methods achieved a time reduction between 60 and 90 %, also achieving an efficiency increase, as intended. Similar results were found for the gradients evaluation, showing a superior FLOWERS with negligible computational times with compared to the other methods, and a BQ that barely took advantage of parallelization, thus distancing it from other computationally efficient methods. Finally, it was found that Uniform C_T computed the gradients faster than Average WS, especially when parallelized, proving the advantage of using a single C_T value.

Regarding the efficacy of the AEP methods in a WFLO, the median optimized AEP shows that with a few exceptions for FLOWERS, all methods studied here yielded poorer layouts compared to those obtained using the baseline as objective function. The strongly unimodal wind rose used for the Sofia wind farm posed a challenge for most models, with the exception of Coarse WD and SGD, which did not see their performance affected. Surprisingly FLOWERS performed poorly in this scenario – with optimized AEPs between 1.5 and 2 % lower compared to the baseline – especially given how easily it improved upon the baseline in the other cases. It was stated in LoCascio et al. (2024) and LoCascio et al. (2025) that this type of wind roses reduced the accuracy of FLOWERS, therefore confirming here that this is the model's main weakness. The other two methods using similar assumptions to FLOWERS are Average WS and Uniform C_T , also with noticeable difficulties under these conditions, although with AEPs between 0.4 and 0.8 % lower. The complications suffered by BQ and OWQ where the unimodal wind rose also resulted in a lack of capability to capture its physics, resulting in the worst result for this method out of all wind farms. Nevertheless, their results in other wind farms demonstrated a similar issue, not of the same magnitude, showing



that using only the wind conditions with the highest probability is not enough to properly optimize a layout. In addition, simply summing over the optimal sample selection as in OWQ yields better results than using the BQ approach, with values of 0.37 and 0.51 % lower than the reference, respectively. On the other hand SGD proved to be highly versatile, achieving virtually the same results as the baseline agreeing with the results found in Quick et al. (2023), while Coarse WD bins did so only for the BPA wake model, thanks to the smoothening effect coming from the gaussian shaped speed deficit. Finally, the similarity between the results obtained by Average WS and Uniform C_T , demonstrated that the use of a unique thrust coefficient for all turbines in each wind flow considered barely alters the final result.

The primary objective of this study was to ascertain which AEP models were computationally efficient. In order to evaluate this, the median optimization time derived from the 30 random scenarios was analyzed. While all methods, except SGD, managed to reduce the optimization times significantly, the reductions achieved by FLOWERS – of more than 95 % – were remarkable thanks to its fast AEP and gradients evaluation and smooth design space, proven to help the optimizer’s convergence (Martins and Ning, 2021), and confirming the results found in LoCascio et al. (2025). Combined with the layouts close to the baseline found by FLOWERS, with the exception of Sofia, FLOWERS proved itself to be one of the most solid methods in the study, achieving balance between efficiency and efficacy. On the other hand, SGD, as a consequence of its fixed budget of 2000 iterations, resulted in the slowest optimizations of the entire study, with values between 2 and 15 times slower, although these results can be potentially improved with a more suitable hyperparameter choice. Moreover, BQ and OWQ did not achieve such substantial time reductions, sometimes even taking longer than the reference, once again reflecting the model’s difficulties to capture the physics of the problem. It was concluded that the better optimization results and optimization times from OWQ made it the superior model in the probabilistic AEP category. In the case of the coarser binning category, Coarse WD was found to achieve consistent time reductions, which when combined with a gaussian wake model, managed to find a balance between efficiency and efficacy. While Average WS and Uniform C_T displayed similar results in terms of final layouts, the optimization time differed, especially in the non-convex cases (Thor and Hornsea 2) where the optimizer requires more function and gradients evaluations, probably caused by the faster gradient evaluations in Uniform C_T , also taking more advantage of parallelizing the flow evaluations.

In general, despite all methods managing to optimize more efficiently than the selected baseline, the quality of the optimized layouts differed significantly across AEP models, resulting in a variety of efficacy and efficiency balances. FLOWERS was shown to display remarkable optimization times and layouts, with a strong weakness when dealing with strongly unimodal wind roses. Coarse WD in combination with the BPA wake model on the other hand found layouts close to those found by the baseline while decreasing the optimization time, although not nearly as substantially as FLOWERS, but with the added advantage of not struggling with unimodal wind roses, demonstrating a strong versatility. These results, in addition to the simple and straightforward setup in these two methods, make them the clear choice when dealing with layout optimization in a project, providing similar results to those from the baseline but with time reductions around 95 % for FLOWERS and 40 % for BPA-Coarse WD. Moreover, highly intricate configurations, such as those belonging to the probabilistic category, BQ and OWQ, considerably diminished their appeal for layout optimization. Finally, the new FLOWERS implementation introduced here, NTP-FLOWERS, demonstrated similar behavior to the other two FLOWERS implementations, with time reductions



around 98 % and similar or better layouts than those obtained by the baseline, albeit with the flaw of having trouble with the strongly unimodal wind rose in Sofia that characterized FLOWERS in this study.

600 In future work, the smoothening effect provided by using the BPA wake model in combination with a coarser wind direction binning should be studied more extensively with different bin sizes and wind farms. Further analysis is required to establish the maximum size of the wind direction bins before the model becomes impractical for WFLO. The BQ and OWQ models were found to compute AEP very efficiently and accurately, but not being very suitable for WFLO. In this study, engineering wake models were used to train the GP in which the model is based, but higher fidelity models may be used. This would potentially
605 result in a high fidelity surrogate model for AEP evaluations. Finally, FLOWERS was observed to demonstrate flaws when used with strongly unimodal wind roses, however this may be specific to the case used in this study. This should be investigated using other study cases.

Code and data availability.

The code and wind roses used to obtain the results presented here can be found in <https://doi.org/10.5281/zenodo.23211846>
610 Andueza-Bueno (2026).

Appendix A: Analysis of the number of multi-starts to determine optimal sampling strategy in BQ and OWQ

Bayesian Quadrature (BQ) and Optimal-design Weighted Quadrature (OWQ) carry out an optimization problem to minimize the uncertainty of the AEP estimate, resulting in an optimal sampling strategy of the wind conditions used in the computation. To avoid becoming trapped in local optima, a total number of 10 random starts were optimized to obtain the results that are
615 presented in this study. In order to validate this as a proper number of multi-starts, and analysis is provided here to demonstrate how the number of multi-starts performed affects the AEP estimation. For an 80 turbine wind farm with a fixed layout, the sampling strategy optimization problem was carried out for an increasing number of multi-starts, for a total of 20 times. The mean and standard deviation of the AEP, compared to the AEP obtained using the baseline, are shown in Fig. A1.

The estimated AEP values converge at around 5 multi-starts, reflecting that the choice of 10 was acceptable. However, an
620 examination of the standard deviation of the 20 runs indicates that convergence occurs at 20 multi-starts. This suggests that the number used may result in a higher variability in the AEPs obtained than could be achieved with higher amounts, meaning that for some runs, the uncertainty of the AEP estimate was not properly minimized. This was only noticed after the study had been completed, but the difference in standard deviation between 10 and 20 multi-starts is low enough as to assume that results would not differ significantly from those obtained here.




Figure A1. Mean AEP and standard deviation out of 20 runs using BQ for an increasing number of multi-starts in the optimal sampling, compared to the that using the baseline for an 80 turbine wind farm.

625 Appendix B: Justification of number of CPUs used for each AEP method

The maximum number of CPUs to be used for parallelization in AEP modeling is very often determined by the overhead time. The study done in Valotta Rodrigues et al. (2024) shows that the speedup achieved with parallelization remains relatively constant with the number of turbines for large wind farms. The results obtained are displayed in Fig. B1 and Fig. B2, where for each AEP method and wake model the computation time is reported using an increasing number of CPUs. The value at which the performance stops improving is considered as the optimal number of CPUs. The analysis are carried out for 100 turbines, an amount for which BQ decreases its computational time when parallelizing AEP calculations, but not gradients. In this case, the number of CPUs is the one where the performance stops improving for the AEP, given that it does not affect the gradients evaluation time negatively.

In general, the time needed to obtain AEP and gradients decreases significantly until around 8 CPUs, although NTP based methods further decrease until 16 CPUs. Given that the speed up obtained using such number of CPUs is very reduced compared to 8 CPUs, in order to achieve an efficient allocation of resources, all parallelization methods are allocated 8 CPUs. Regarding SGD, given that it only performs 50 flow simulations, parallelization does not provide any advantages.

Appendix C: Median number of iterations for each optimization set up, AEP model and wake model



Wind Farm	AEP Model	NOJ		BPA		NTP	
		Median	Std	Median	Std	Median	Std
Hornsea 3	Baseline	78.0	6.34	39.0	6.56	65.5	9.38
	Coarse WD	147.5	19.72	41.5	6.04	150.0	17.14
	Average WS	71.5	7.01	50.5	8.77	75.0	5.93
	FLOWERS	74.0	8.66	89.5	11.21	71.0	9.43
	Uniform C_T	78.5	7.32	48.0	7.72	74.0	5.93
	BQ	156.0	34.89	155.5	58.60	154.5	26.30
	OWQ	126.0	16.64	123.0	31.38	148.5	11.94
Sofia	Baseline	73.5	12.44	38.0	5.40	81.0	13.58
	Coarse WD	116.5	24.68	37.0	6.13	127.0	21.91
	Average WS	50.5	8.19	31.0	5.16	53.5	9.11
	FLOWERS	70.0	9.43	287.0	8.38	72.0	10.67
	Uniform C_T	48.5	8.33	29.0	4.04	53.5	9.11
	BQ	254.0	65.53	178.5	34.35	221.0	56.86
	OWQ	85.0	21.33	62.5	15.78	95.0	20.50
Thor	Baseline	131.5	46.92	131.0	53.43	105.0	45.60
	Coarse WD	124.0	34.81	132.0	53.33	131.0	34.22
	Average WS	257.5	188.09	851.0	198.79	189.0	114.92
	FLOWERS	166.0	26.94	143.0	40.13	165.5	15.67
	Uniform C_T	223.5	133.98	767.0	248.83	189.0	114.92
	BQ	129.5	46.40	140.5	30.26	120.0	48.47
	OWG	152.0	254.28	199.0	156.63	171.0	166.33
Hornsea 2	Baseline	140.0	38.96	175.5	26.20	148.5	32.43
	Coarse WD	146.0	24.76	179.5	35.75	150.5	25.55
	Average WS	467.0	299.56	944.5	308.98	226.5	268.02
	FLOWERS	160.5	11.38	184.5	19.15	160.0	15.46
	Uniform C_T	191.0	293.28	921.5	321.00	226.5	268.02
	BQ	223.0	59.05	173.0	29.46	198.5	42.83
	OWG	167.5	252.64	283.5	111.14	193.0	190.84

Table C1. Median and standard deviation of optimizer iterations across wake models, by wind farm and AEP model.




Figure B1. Time taken to compute AEP for a 100 turbine wind farm for the different methods and wake models, for an increasing number of CPUs. FLOWERS cannot be parallelized so it is excluded from the analysis.




Figure B2. Time taken to compute gradients for a 100 turbine wind farm for the different methods and wake models, for an increasing number of CPUs. FLOWERS cannot be parallelized, so it is excluded from the analysis.



Author contributions.

640 JAB wrote the code for FLOWERS TurbOPark, BQ and OWQ, assembled the optimization setup, obtained and analyzed the results and wrote the initial draft of the paper. LA, MFM and PER conceptualized the project, developed the methodology and supervised the writing (review and editing).

Competing interests.

The authors declare that they have no conflict of interests.

645 *Acknowledgements.* The authors gratefully acknowledge the computational and data resources provided by the Sophia HPC cluster at the Technical University of Denmark, DOI: <https://doi.org/10.57940/FAFC-6M81> (Technical University of Denmark, 2019). AI tools were used during the writing of the report. Claude Opus models were used to assist during the development of the code base iteratively, under close supervision of JAB, particularly to increase code performance and during the construction of specific functions. JAB thoroughly reviewed and edited the whole code base. JAB wrote the initial draft, and used Claude models to, under his supervision, enhance the clarity of the explanations and review and edit grammar and wording. The authors iteratively reviewed all changes and modifications.

650



References

- Wind turbines – Part 1: Design requirements, 2010.
- Andueza-Bueno, J.: javandu45/fastAEP: Fast AEP, <https://doi.org/10.5281/zenodo.23211846>, 2026.
- Bastankhah, M. and Porté-Agel, F.: A new analytical model for wind-turbine wakes, *Renewable Energy*, 70, 116–123, <https://doi.org/10.1016/j.renene.2014.01.002>, 2014.
- Chugh, T. and Ymeraj, E.: Wind Farm Layout Optimisation using Set Based Multi-objective Bayesian Optimisation, <https://doi.org/10.48550/arXiv.2203.17065>, arXiv:2203.17065 [stat.ML], 2022.
- Criado Risco, J., Valotta Rodrigues, R., Friis-Møller, M., Quick, J., Mølgaard Pedersen, M., and Réthoré, P.-E.: Gradient-based wind farm layout optimization with inclusion and exclusion zones, *Wind Energy Science*, 9, 585–600, <https://doi.org/10.5194/wes-9-585-2024>, 2024.
- Dörenkämper, M., Olsen, B. T., Witha, B., Hahmann, A. N., Davis, N. N., Barcons, J., Ezber, Y., García-Bustamante, E., González-Rouco, J. F., Navarro, J., Sastre-Marugán, M., Sile, T., Trei, W., Žagar, M., Badger, J., Gottschall, J., Sanz Rodrigo, J., and Mann, J.: The Making of the New European Wind Atlas – Part 2: Production and evaluation, *Geoscientific Model Development*, 13, 5079–5102, <https://doi.org/10.5194/gmd-13-5079-2020>, 2020.
- Feng, J. and Shen, W. Z.: Solving the wind farm layout optimization problem using random search algorithm, *Renewable Energy*, 78, 182–192, <https://doi.org/10.1016/j.renene.2015.01.005>, 2015.
- Feng, J. and Shen, W. Z.: Design optimization of offshore wind farms with multiple types of wind turbines, *Applied Energy*, 205, 1283–1297, <https://doi.org/10.1016/j.apenergy.2017.08.107>, 2017.
- Fleming, P. A., Stanley, A. P. J., Bay, C. J., King, J., Simley, E., Doekemeijer, B. M., and Mudafort, R.: Serial-Refine Method for Fast Wake-Steering Yaw Optimization, *Journal of Physics: Conference Series*, 2265, 032 109, <https://doi.org/10.1088/1742-6596/2265/3/032109>, 2022.
- Gonzalez-Rodriguez, A. G., Riega-Medina, J. M., Ruano-Ruano, I., and Muñoz-Diez, J. V.: Practices and Considerations in Wind Data Processing for Accurate and Efficient Wind Farm Energy Calculation, *Energies*, 18, 3402, <https://doi.org/10.3390/en18133402>, 2025.
- Guirguis, D., Romero, D. A., and Amon, C. H.: Toward efficient optimization of wind farm layouts: Utilizing exact gradient information, *Applied Energy*, 179, 110–123, <https://doi.org/10.1016/j.apenergy.2016.06.101>, 2016.
- Hahmann, A. N., Sile, T., Witha, B., Davis, N. N., Dörenkämper, M., Ezber, Y., García-Bustamante, E., González-Rouco, J. F., Navarro, J., Olsen, B. T., and Söderberg, S.: The making of the New European Wind Atlas – Part 1: Model sensitivity, *Geoscientific Model Development*, 13, 5053–5078, <https://doi.org/10.5194/gmd-13-5053-2020>, 2020.
- Herbert-Acero, J., Probst, O., Réthoré, P.-E., Larsen, G., and Castillo-Villar, K.: A Review of Methodological Approaches for the Design and Optimization of Wind Farms, *Energies*, 7, 6930–7016, <https://doi.org/10.3390/en7116930>, 2014.
- Jensen, N.: A note on wind generator interaction, no. 2411 in Risø-M, Risø National Laboratory, ISBN 87-550-0971-9, 1983.
- King, R., Glaws, A., Geraci, G., and Eldred, M. S.: A Probabilistic Approach to Estimating Wind Farm Annual Energy Production with Bayesian Quadrature, in: AIAA Scitech 2020 Forum, American Institute of Aeronautics and Astronautics, Orlando, FL, ISBN 978-1-62410-595-1, <https://doi.org/10.2514/6.2020-1951>, 2020.
- Kingma, D. P. and Ba, J.: Adam: A Method for Stochastic Optimization, <https://arxiv.org/abs/1412.6980>, 2017.
- LoCascio, M. J., Bay, C. J., Bastankhah, M., Barter, G. E., Fleming, P. A., and Martínez-Tossas, L. A.: FLOW Estimation and Rose Superposition (FLOWERS): an integral approach to engineering wake models, *Wind Energy Science*, 7, 1137–1151, <https://doi.org/10.5194/wes-7-1137-2022>, 2022.



- LoCascio, M. J., Bay, C. J., Martínez-Tossas, L. A., Bastankhah, M., and Gorié, C.: FLOWERS AEP: An Analytical Model for Wind Farm Layout Optimization, *Wind Energy*, 27, 1563–1580, <https://doi.org/10.1002/we.2954>, 2024.
- 690 LoCascio, M. J., Bay, C. J., Martínez-Tossas, L. A., Thomas, J. J., and Gorié, C.: Efficient wind farm layout optimization with the FLOWERS AEP model and analytic gradients, *Journal of Renewable and Sustainable Energy*, 17, 013 302, <https://doi.org/10.1063/5.0237778>, 2025.
- Martins, J. R. R. A. and Ning, A.: *Engineering Design Optimization*, Cambridge University Press, 2021.
- Murcia, J. P., Réthoré, P. E., Natarajan, A., and Sørensen, J. D.: How Many Model Evaluations Are Required To Predict The AEP Of A Wind Power Plant?, *Journal of Physics: Conference Series*, 625, 012 030, <https://doi.org/10.1088/1742-6596/625/1/012030>, 2015.
- 695 Nygaard, N. G., Steen, S. T., Poulsen, L., and Pedersen, J. G.: Modelling cluster wakes and wind farm blockage, *Journal of Physics: Conference Series*, 1618, 062 072, <https://doi.org/10.1088/1742-6596/1618/6/062072>, 2020.
- 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 Energy Science*, 4, 211–231, <https://doi.org/10.5194/wes-4-211-2019>, 2019.
- Pedersen, M. M., van der Laan, P., Friis-Møller, M., Forsting, A. M., Riva, R., Román, L. A. A., Risco, J. C., Quick, J., Christiansen, J. P. S., Olsen, B. T., Rodrigues, R. V., and Réthoré, P.-E.: DTUWindEnergy/PyWake: PyWake, <https://doi.org/10.5281/zenodo.6806136>, 2023.
- 700 Pedersen, M. M., Friis-Møller, M., Réthoré, P.-E., Simutis, E., Riva, R., Quick, J., Dimitrov, N. K., Rinker, J., and Dykes, K.: DTUWindEnergy/TopFarm2: Release of v2.6.1, <https://doi.org/10.5281/zenodo.17540961>, 2025.
- Porté-Agel, F., Bastankhah, M., and Shamsoddin, S.: Wind-Turbine and Wind-Farm Flows: A Review, *Boundary-Layer Meteorology*, 174, 1–59, <https://doi.org/10.1007/s10546-019-00473-0>, 2020.
- 705 Pujari, K. N. and Mitra, K.: Wind Farm Layout Optimization Under Uncertainty using Bayesian Approach, in: 2023 Ninth Indian Control Conference (ICC), pp. 90–95, IEEE, Visakhapatnam, India, ISBN 979-8-3503-9525-9, <https://doi.org/10.1109/ICC61519.2023.10442007>, 2023.
- Quick, J., Rethore, P.-E., Mølgaard Pedersen, M., Rodrigues, R. V., and Friis-Møller, M.: Stochastic gradient descent for wind farm optimization, *Wind Energy Science*, 8, 1235–1250, <https://doi.org/10.5194/wes-8-1235-2023>, 2023.
- 710 Riva, R., Liew, J., Friis-Møller, M., Dimitrov, N., Barlas, E., Réthoré, P.-E., and Beržonskis, A.: Wind farm layout optimization with load constraints using surrogate modelling, *Journal of Physics: Conference Series*, 1618, 042 035, <https://doi.org/10.1088/1742-6596/1618/4/042035>, 2020.
- Sanderse, B., Van Der Pijl, S. P., and Koren, B.: Review of computational fluid dynamics for wind turbine wake aerodynamics, *Wind Energy*, 14, 799–819, <https://doi.org/10.1002/we.458>, 2011.
- 715 Technical University of Denmark: Sophia HPC Cluster, <https://doi.org/10.57940/FAFC-6M81>, 2019.
- Thomas, J. J. and Ning, A.: A method for reducing multi-modality in the wind farm layout optimization problem, *Journal of Physics: Conference Series*, 1037, 042 012, <https://doi.org/10.1088/1742-6596/1037/4/042012>, 2018.
- Thomas, J. J., McOmber, S., and Ning, A.: Wake expansion continuation: Multi-modality reduction in the wind farm layout optimization problem, *Wind Energy*, 25, 678–699, <https://doi.org/https://doi.org/10.1002/we.2692>, 2022.
- 720 Thomas, J. J., Baker, N. F., Malisani, P., Quaeghebeur, E., Sanchez Perez-Moreno, S., Jasa, J., Bay, C., Tilli, F., Bieniek, D., Robinson, N., Stanley, A. P. J., Holt, W., and Ning, A.: A comparison of eight optimization methods applied to a wind farm layout optimization problem, *Wind Energy Science*, 8, 865–891, <https://doi.org/10.5194/wes-8-865-2023>, 2023.
- Valotta Rodrigues, R., Pedersen, M. M., Schøler, J. P., Quick, J., and Réthoré, P.-E.: Speeding up large-wind-farm layout optimization using gradients, parallelization, and a heuristic algorithm for the initial layout, *Wind Energy Science*, 9, 321–341, <https://doi.org/10.5194/wes-9-321-2024>, 2024.
- 725



Vermeer, L., Sørensen, J. N., and Crespo, A.: Wind Turbines Wake Aerodynamics, 2003.

Wang, L., Dong, M., Yang, J., Wang, L., Chen, S., Duić, N., Joo, Y. H., and Song, D.: Wind turbine wakes modeling and applications: Past, present, and future, *Ocean Engineering*, 309, 118 508, <https://doi.org/10.1016/j.oceaneng.2024.118508>, 2024.

730 Whittaker, C., LoCascio, M. J., Martínez-Tossas, L. A., Bay, C. J., and Bastankhah, M.: Gaussian FLOWERS: Wind-rose-based analytical integration of Gaussian wake model for extremely fast AEP estimation, *Journal of Renewable and Sustainable Energy*, 17, 013 306, <https://doi.org/10.1063/5.0245886>, 2025.

Wolf Ciavarra, A., Valotta Rodrigues, R., Dykes, K., and Réthoré, P.-E.: Wind farm optimization with multiple hub heights using gradient-based methods, *Journal of Physics: Conference Series*, 2265, 022 012, <https://doi.org/10.1088/1742-6596/2265/2/022012>, 2022.