

Thank you for your suggestions. We have incorporated them in the final manuscript has detailed below. The specific changes can be seen in the difference pdf file.

In the updated version the authors have addressed most of my suggestions / comments from my previous report. Based on the changes, and after reading the manuscript again, I have the following final suggestions with respect to this manuscript. When these are incorporated the

End section 6 and discussion conclusion section

** Here, the authors discuss the results obtained using the different starting layouts. For example, it is found that the optimal layout obtained using the Amalia layout as starting configuration is not as good as using a rectangular grid. It should be added that the Amalia layout has been obtained using different optimization constraints than used in the present study (or simply by using a slightly different wake mode). Given the local maxima in the optimizations procedure this causes the optimization algorithm to get stuck in this local optimum that is created.*

We have noted that the Amalia layout has been obtained by solving a different optimization problem as shown in the diff pdf file.

* In the abstract the authors should also add a sentence on the sensitivity of the starting configuration, which clearly is also an important consideration in this optimization.

We have incorporated this suggestion in the abstract.

* "With PC based on regression, we have reduced by a factor of five the number of simulations required to accurately compute the AEP, thus enabling the use of more expensive, higher-fidelity models or larger wind farm optimizations."

==> Please specify reduced by a factor of five compared to "???"

==> Later on in the abstract a reduction factor of 3 is mentioned. So make sure what the factor 3 and 5 exactly refer to.

We have rewritten the second half of the abstract to incorporate both of the above suggestions, i.e., we have added a sentence on the sensitivity of the starting configuration and mentioned exactly what the factor of 3 and 5 refer to.

* in the abstract and conclusion specify that the numbers have been obtained for this one wind farm test case.

We have explicitly specified in both the abstract and the conclusion that the number obtained have been for the cases considered as can be seen in the diff pdf file.

Minor:

*Around equation 36 there is a sentence: "36 wind directions the chosen intervals are {[0, 360]; [10, 370]; [20, 380]; ... ; [340, 700]; [350, 710]"}
==> The mentioned wind angle ranges are somewhat strange

We've change this sentence to make it less confusing and easier to understand.

* Section 6 first sentence

AEP has already been introduced before.

Thanks, we have deleted this sentence.

*Section 6.2 In what follows we will compare the PC-R (the best performing method) with the rectangle rule (the method currently used in practice)
==> please formulate more generally; the rectangle rule is not the only method used

We have changed it to the more general statement of "to the method most commonly used in practice." In section 3, when describing the methods, we mention other methods besides the rectangle rule.

Polynomial chaos to efficiently compute the annual energy production in wind farm layout optimization

Andrés Santiago Padrón¹, Jared Thomas², Andrew P. J. Stanley², Juan J. Alonso¹, and Andrew Ning²

¹Department of Aeronautics & Astronautics, Stanford University, Stanford, CA, 94305, USA

²Department of Mechanical Engineering, Brigham Young University, Provo, UT, 84602, USA

Correspondence to: Andrew Ning (aning@byu.edu)

Abstract.

In this paper, we develop computationally-efficient techniques to calculate statistics used in wind farm optimization with the goal of enabling the use of higher-fidelity models and larger wind farm optimization problems. We apply these techniques to maximize the Annual Energy Production (AEP) of a wind farm by optimizing the position of the individual wind turbines. The AEP (a statistic) is the expected power produced by the wind farm over a period of one year subject to uncertainties in the wind conditions (wind direction and wind speed) that are described with empirically-determined probability distributions. To compute the AEP of the wind farm, we use a wake model to simulate the power at different input conditions composed of wind direction and wind speed pairs. We use polynomial chaos (PC), an uncertainty quantification method, to construct a polynomial approximation of the power over the entire stochastic space and to efficiently (using as few simulations as possible) compute the expected power (AEP). We explore both regression and quadrature approaches to compute the PC coefficients. PC based on regression is significantly more efficient than the rectangle rule (the method most commonly used to compute the expected power). With PC based on regression, we have reduced on average by a factor of five the number of simulations required to accurately compute the AEP ~~,-thus enabling the when compared to the rectangle rule for the different wind farm layouts considered. In the wind farm layout optimization problem, each optimization step requires an AEP computation. Thus, the ability to compute the AEP accurately with fewer simulations is beneficial as it reduces the cost to perform an optimization, which enables the use of more expensive,-computationally expensive higher-fidelity models or larger wind farm optimizations~~the consideration of larger or multiple wind farm optimization problems. We perform a large suite of gradient-based optimizations to compare the optimal layouts obtained when computing the AEP with polynomial chaos based on regression and the rectangle rule. We consider three different starting layouts (Grid, Amalia, Random) and find that the optimization has many local optima and is sensitive to the starting layout of the turbines. We observe that starting from a good layout (Grid, Amalia) will, in general, find better optima than starting from a bad layout (Random) independent of the method used to compute the AEP.~~We show how the PC method can be used to efficiently compute the gradients of the AEP and we perform a large suite of gradient-based optimizations with different initial turbine locations and with different numbers of samples~~ For both PC based on regression and the rectangle rule, we consider both a *coarse* (~ 225) and a *fine* (~ 625) number of simulations to compute the AEP. We find that ~~for roughly one-third of the computational cost, the optimizations with the optimizations with coarse~~ PC based on regression result in optimized layouts that produce comparable AEP as the optimized

layouts found with the ~~rectangle rule but using only one-third of the samples (one-third of the computational cost)~~fine rectangle rule. Furthermore, for the same computational cost, for the different cases considered, polynomial chaos finds optimal layouts with 0.4 % higher AEP on average than those found with the rectangle rule.

1 Introduction

5 In 2015, wind energy growth accounted for almost half of the global electricity supply growth. In the United States, it accounted for 41 % of new power capacity, raising the wind energy supply to 4.7 % of the total electricity generated in 2015 and on target to reach 10 % by 2020 (U.S. Department of Energy, 2015; AWEA, 2016; GWEC, 2016). Most of the current and upcoming wind energy comes from large turbines (greater than 1 MW) situated in clusters—wind farms. A problem with putting turbines together in confined spaces is that they operate in the wakes of other turbines, i.e., in regions of reduced speed and increased
10 turbulence. This leads to an underproduction of power and decreased (10–20 %) energy output for the farm (Barthelmie et al., 2007, 2009; Briggs, 2013) when compared to ideal conditions. This loss in energy capture results in millions of dollars of loss for operators and investors and increased economic uncertainty for new installations. Many current wind farms have grid-like layouts, where wind turbines are aligned in rows, which further exacerbate the wake losses. By optimizing the layout of the wind farm, the wake losses can be minimized, with a corresponding increase in energy production and revenue.

15 Wind farm optimization is a complex, multi-disciplinary and high-dimensional problem. The wind farm may contain dozens or even hundreds of wind turbines, where each turbine may be parameterically described using several design variables. Furthermore, the wind conditions (wind direction, wind speed, wind turbulence, etc.) are stochastic (uncertain), and thus we need a statistic to evaluate the performance of the wind farm. A common statistic is the expected power or the Annual Energy Production (AEP). Many model simulations are needed to estimate the statistic (Padrón et al., 2016; Murcia et al., 2015). The
20 statistic is usually the objective function of the optimization (Herbert-Acero et al., 2014); thus, the wind farm optimization is an optimization under uncertainty problem (Fig. 1). Optimization under uncertainty (OUU) differs from deterministic optimization in that it contains a nested uncertainty quantification loop to compute statistics. In the OUU problem, for every optimization step many model evaluations are needed to compute the relevant statistics. Thus, even with a very small number of design variables per turbine, the total number of variables and simulations required by the wind farm optimization can grow
25 very rapidly (Gebraad et al., 2017) and quickly make the problem infeasible, especially when using a high-fidelity model for the wind farm simulation.

We see three approaches to improving wind farm optimization capabilities. Each approach focuses on the different blocks of the optimization under uncertainty problem (Fig. 1c). The first approach is to improve the modeling quality of entire wind farms, i.e., improve the models in all the disciplines (aerodynamics, structures, controls, electrical, acoustics, atmospheric
30 physics, policy, economics, etc.) that are relevant to building and operating a wind farm, as well as the interaction between the different turbines. The second approach is to improve the optimization problem formulation and the algorithms to solve the optimization. And the third approach is to improve the treatment of the stochastic nature of the problem, i.e., develop better




Figure 1. Examples of applications that require many model evaluations. In deterministic optimization (a), we evaluate the model at different values of the design variables while searching for an optimum response. In uncertainty quantification (b), we query the model multiple times at different instances of the uncertain variables to generate an ensemble of responses from which we can compute statistics and probabilities of the model response. In optimization under uncertainty (OUU), we optimize a statistic. OUU is computationally expensive as it requires many model evaluations because of the nested loop in the problem formulation (c).

uncertainty quantification methods to efficiently compute the relevant wind farm statistics (and their gradients with respect to the design variables of the problem) used in the optimization under uncertainty problem.

The first approach increases the fidelity of the model whereas the second (optimization) and third (uncertainty quantification) approaches seek to reduce the number of model evaluations, as this enables the study of larger and more realistic problems.

5 Here, we focus on the uncertainty quantification approach (the third approach), as it has not been considered in detail before. The most recent and thorough review of the wind farm optimization literature (Herbert-Acero et al., 2014) does not mention it. It only mentions the first two approaches. In the existing work in the literature, the third approach typically focuses on simple integration methods to compute the statistics, which quantify the effect of the stochastic inputs (Kusiak and Song, 2010; Kwong et al., 2012; Chowdhury et al., 2013; Fleming et al., 2016; Gebraad et al., 2017). Simple integration methods, 10 such as the rectangle rule, are inefficient in the number of samples (simulations of the model) needed to accurately estimate a statistic, such as the AEP. They are especially inefficient if multiple stochastic inputs are considered simultaneously. Normally only the wind direction and/or the wind speed are considered as stochastic input variables. Recent work (Padrón et al., 2016; Murcia et al., 2015) is starting to move beyond these simple integration techniques to compute the AEP and instead using the uncertainty quantification method of polynomial chaos to compute the AEP.

15 In this paper, which is meant as a comprehensive introduction to uncertainty quantification methods applied to wind-farm simulations, we describe in detail the polynomial chaos (PC) method and show that, for the efficient (small number of model simulations) computation of the AEP, the PC method based on regression should be used. An additional benefit of the PC method is that it makes it feasible to consider multiple uncertain variables (e.g., wind direction, wind speed, wind turbulence,

wake model parameter) that impact the computation of the AEP. An example where a wake model parameter is considered as an uncertain variable in addition to the wind speed and wind direction can be found in Padrón (2017). In addition, the PC method can be used to efficiently compute the gradient of statistics, such as the AEP. Eldred (2011) describes how to compute the gradients for PC based on quadrature. Here we show how to compute the gradients for PC based on regression (Sect. 4.3).

5 The use of gradients allows us to efficiently tackle much larger optimization problems (Gebraad et al., 2017). To compute the gradients of the wake model, we use the recently developed Floris wake model with the modifications by Thomas et al. (2017) that provide analytic and continuous gradients of the wake model.

We first discuss the details of computing the power and the AEP of a wind farm in Sect. 2. Then, we discuss uncertainty quantification in Sect. 3 and the polynomial chaos method in Sect. 4. Finally, we discuss the details of the problem formulation
10 in Sect. 5 and the results in Sect. 6.

2 Computing the power and the annual energy production of a wind farm

We first describe the aerodynamic wake model we use (Sect. 2.1). The wake model gives an estimate of the hub-height velocity at each wind turbine, from which we can compute the power produced by the wind farm (Sect. 2.2). Then, to obtain the Annual Energy Production (AEP) we need to integrate the power over all wind conditions that occur in a year (Sect. 2.3) and weigh
15 the results proportionally to the frequency with which such wind conditions manifest themselves.

2.1 Floris

The Floris (FLoW Redirection and Induction in Steady-state)¹ (Gebraad et al., 2016) wake model is an enhancement of the Jensen wake model (Jensen, 1983) and the wake deflection model presented in (Jiménez et al., 2009). The Floris model builds on the Jensen model by defining three separate wake zones with differing expansion and decay rates (controlled by tunable
20 coefficients) to more accurately describe the velocity deficit across the wake region. A simple overlap ratio of the area of the rotor in each zone of each shadowing wake to the full rotor area is used to determine the effective hub velocity of a given turbine. A simple overview of the Floris model, showing the zones and overlap areas, is shown in Fig. 2. We use the Floris wake model with changes to remove discontinuities and add curvature to regions of non-physical zero gradient to make the model more suitable for gradient-based optimization (Thomas et al., 2017). In this work, we use the parameter values recommended
25 in (Gebraad et al., 2016) and (Thomas et al., 2017) and set the yaw-offset angle of each turbine to zero.

2.2 Computing the power of a wind farm

We will consider the power of the wind farm to be a function of three classes of variables: uncertain variables ξ , design variables \mathbf{x} , and parameters θ ,

$$P = P(\xi, \mathbf{x}, \theta). \tag{1}$$

¹We use the name Floris for the model, instead of FLORIS, the name used in (Gebraad et al., 2016).




Figure 2. Schematic of the Floris wake model. The model has three zones with varying diameters, $D_{w,q}$, that depend on tuning parameters k_e and $m_{e,q}$. The effective hub velocity is computed using the overlap ratio, A_q^{OL} , of the part of the rotor-swept area overlapping each wake zone respectively to the total rotor-swept area.

Table 1. The variables used for calculating the power.

Uncertain ξ	wind direction.
	freestream wind speed.
Design \mathbf{x}	x - The x location of each turbine.
	y - The y location of each turbine.
Parameters θ	yaw angles, turbine characteristics, and wake model parameters.

Uncertain variables are variables that follow a probability distribution, design variables are variables that an optimizer can vary, and parameters are important constants that govern the behavior of the system. The classification of the variables is problem dependent. For instance, the rotor yaw could be considered a design variable, a parameter, or an uncertain variable to account for yaw-offset measurement error. A tunable parameter of a wake model, such as the wake expansion coefficient, could be considered as a parameter or as an uncertain variable given by a particular distribution.

For the problems considered in this work, Table 1 lists in which category we place each variable that influences the power computation. The uncertain variables are the wind direction and wind speed with probability distributions described in Sect. 5.1.

The power of the wind farm for a given wind direction and wind speed is equal to the sum of the power produced by each turbine $P = \sum_{i=1}^{n_{turb}} P_i$. The power of each turbine is calculated from

$$P_i = \frac{1}{2} \rho C_P A U_i^3, \quad (2)$$

where ρ is the air density, A is the rotor swept area, C_P is the power coefficient, and U_i is the effective hub velocity for each turbine, which is calculated by the wake model and is a function of the three types of variables described above, $U_i = f(\boldsymbol{\xi}, \mathbf{x}, \boldsymbol{\theta})$. The power coefficient captures both the aerodynamic and electromechanical properties of the wind turbine. It is a complex function of many variables (Herbert-Acero et al., 2014) and is usually reported by wind turbine manufacturers as a function of the tip-speed ratio, which depends on wind speed at hub height U_i . A simple expression for C_P can also be computed using the classical actuator disk theory (Sanderse, 2009).

2.3 Computing the Annual Energy Production (AEP) of a wind farm

The Annual Energy Production (AEP) is an important metric used to describe a wind farm. The AEP is a statistic. Specifically, it is a mean, as it is a function of the expected power multiplied by the number of hours in a year,

$$AEP = 8760 \frac{hr}{yr} E[P(\boldsymbol{\xi})]. \quad (3)$$

The expected power, $E[P(\boldsymbol{\xi})]$, or the mean of the power, μ_P , is defined as

$$\mu_P = E[P(\boldsymbol{\xi})] = \int_{\Omega} P(\boldsymbol{\xi}) \rho(\boldsymbol{\xi}) d\boldsymbol{\xi}, \quad (4)$$

where $\boldsymbol{\xi} = (\xi_1, \xi_2, \dots, \xi_n)$ is a vector of random variables, which we refer to as the uncertain variables, $\rho(\boldsymbol{\xi})$ is the joint probability density function of the uncertain variables, Ω is the domain of the uncertain variables, and P is the power produced by the wind farm (Eq. (1)). Common uncertain variables are the wind direction and the freestream wind speed.

The expected power, and hence the AEP, is normally computed as a weighted average, which amounts to the rectangle rule of integration (Sect. 3.1.1). Other uncertainty quantification methods (Sect. 3) can be used to compute the expected value (AEP). Specifically, we can compute the AEP efficiently by polynomial chaos (Sect. 4).

3 Uncertainty quantification

Uncertainty quantification (UQ) is the process of (1) characterizing input uncertainties, and then (2) propagating these input uncertainties through a computational model with the goal of quantifying their effect on the model's output. There are many sources of uncertainty in the modeling of a problem, and different classifications of the uncertainties have been proposed (Kennedy and O'Hagan, 2001; Oberkampf et al., 2001; Beyers and Sendhoff, 2007). A common classification is to divide the uncertainty into aleatory and epistemic uncertainties (Oberkampf et al., 2001).

In this work, we consider aleatory uncertainties that arise from the variability in the inputs to our model caused by changing environmental conditions. We describe this input variability as random variables with associated probability distributions. Thus, the first step of characterizing the input uncertainties is concerned with finding the probability distributions that describe the model's inputs. This process is known as statistical inference, model calibration and inverse uncertainty quantification (Smith, 2014). Here, we assume that this step has been completed, i.e., we have distributions that characterize the uncertain inputs (Sect. 5.1). We focus on the second step of propagating the input uncertainties to find the statistics that describe the output.

3.1 Uncertainty propagation methods

The goal of uncertainty propagation methods is to compute the statistics that describe the effect of the uncertain inputs on the model output. There are several methods to propagate the uncertainties and compute statistics (Le Maître and Knio, 2010; Smith, 2014), where each method has its advantages/disadvantages depending on the type and size of the problem. The most common methods are sampling or Monte Carlo methods (Caflisch, 1998). Other methods include direct integration methods and stochastic expansion methods. Direct integration methods are the currently used method to compute the AEP of a wind farm; we briefly describe them below (Sect. 3.1.1). We describe the stochastic expansion method of polynomial chaos in detail in Sect. 4, and compare the different methods to propagate the uncertainties to compute the AEP in Sect. 6.2.

3.1.1 Direct numerical integration (Rectangle rule)

As the name implies, this method numerically evaluates the integrals in the definition of the statistics. The integrals to evaluate the mean (or expected value) and the variance are

$$\mu_R = E[R] = \int_{\Omega} R(\boldsymbol{\xi})\rho(\boldsymbol{\xi}) d\boldsymbol{\xi}, \quad (5)$$

$$\sigma_R^2 = \text{Var}[R] = E[(R(\boldsymbol{\xi}) - E[R(\boldsymbol{\xi})])^2] \quad (6)$$

$$= E[R(\boldsymbol{\xi})^2] - (E[R(\boldsymbol{\xi})])^2 \quad (7)$$

$$= \int_{\Omega} R(\boldsymbol{\xi})^2 \rho(\boldsymbol{\xi}) d\boldsymbol{\xi} - \mu_R^2, \quad (8)$$

where $R(\boldsymbol{\xi})$ is the model output and $\boldsymbol{\xi}$ the uncertain input variables. The random vector $\boldsymbol{\xi} = (\xi_1, \xi_2, \dots, \xi_n)$ with joint probability distribution $\rho(\boldsymbol{\xi})$ describes the input variability over the domain Ω . Each random input can follow a particular distribution $\rho_i(\xi_i)$. For the case of independent random variables, the joint distribution is the product of each univariate distribution $\rho(\boldsymbol{\xi}) = \prod_{i=1}^n \rho_i(\xi_i)$.

There are many quadrature methods to evaluate integrals (Ascher and Greif, 2011). We describe the rectangle rule, as this is what is currently used in the wind farm community to compute the AEP.

Rectangle rule. The rectangle rule, or mid-point rule, is the simplest and most straightforward quadrature method. To approximate the mean or expected value,

$$\mu_R = E[R] = \int_{\Omega} R(\boldsymbol{\xi})\rho(\boldsymbol{\xi}) d\boldsymbol{\xi} \quad (5 \text{ revisited})$$

with the rectangle rule, we divide the domain of the uncertain variable² $\Omega = [a, b]$ into m equal subintervals of length $\Delta\xi = (b - a)/m$. Next, we construct rectangles with base $B = \Delta\xi$ and height equal to the product of the response of the model and the density evaluated at the mid-point of the subinterval $H = R(\xi_j)\rho(\xi_j)$. Then, the rectangle rule approximates the expected

²For simplicity, we describe the 1-dimensional case, $\boldsymbol{\xi} = \xi$.

value by adding up the areas of the m rectangles

$$E[R] = \int_a^b R(\xi)\rho(\xi) d\xi \approx \sum_{j=1}^m R(\xi_j)\rho(\xi_j)\Delta\xi. \quad (9)$$

A simple improvement is to integrate the density exactly within each subinterval

$$E[R] = \int_a^b R(\xi)\rho(\xi) d\xi \approx \sum_{j=1}^m R(\xi_j) \int_{\xi_{j-1/2}}^{\xi_{j+1/2}} \rho(\xi) d\xi. \quad (10)$$

- 5 This improvement is easily done as the density is known. This modification is helpful for a small number of evaluations m . We will use this modified rectangle rule and simply refer to it as the rectangle rule.

4 Polynomial chaos

- Polynomial chaos (PC) is the name of an uncertainty quantification (UQ) method that approximates a function with a polynomial expansion made up of orthogonal polynomials. This function has *random variables* as inputs, and we are interested in the effects of the random (uncertain) inputs on the output of this function. Statistics of the output can describe the effects of the inputs. We use the polynomial chaos method to efficiently compute these statistics and the gradients of these statistics.

- We first describe the polynomial chaos method in Sect. 4.1. We then discuss two methods—quadrature and regression—to compute the coefficients used in the PC expansion (Sect. 4.2). In Sect. 4.3, we describe how to compute the gradients of statistics from the polynomial chaos expansion. And, in Sect. 4.4, we discuss how the PC method can be extended to include correlated random inputs.

4.1 Polynomial chaos expansion

Let $R(\xi)$ be a function of interest that depends on the uncertain variable ξ . We can approximate the function by a polynomial expansion

$$R(\xi) \approx \hat{R}(\xi) = \sum_{i=0}^p \alpha_i \phi_i(\xi). \quad (11)$$

- 20 The approximate response $\hat{R}(\xi)$ is a polynomial of order p . Usually, the larger the polynomial order the closer the approximation is to the true response $R(\xi)$.

- The polynomial basis $\{\phi_i(\xi)\}_{i=0}^p$ is determined by the distribution of the uncertain variable—the polynomial basis is orthogonal with a weight function that corresponds up to a constant to the probability density function of the uncertain variable. Common random (uncertain) variables (Normal, Uniform, Exponential, Beta) have corresponding classical orthogonal polynomials (Hermite, Legendre, Laguerre, Jacobi) (Eldred et al., 2008). Empirically-determined distributions, such as those obtained from wind conditions, do not have corresponding classical orthogonal polynomials. For the distributions obtained

from the wind conditions, we need to numerically generate custom orthogonal polynomials in order to preserve the optimal convergence property of the polynomial chaos expansion (Oladyshkin and Nowak, 2012). Details about the numerical generation of orthogonal polynomials can be found in Gautschi (2004) and an example of the generation of orthogonal polynomials for wind distributions in Padrón (2017). In addition to the optimal convergence properties, the use of orthogonal polynomials allows us to analytically compute statistics from the polynomial chaos expansion (Sect. 4.1.1).

In addition to the orthogonal polynomials, the other component of the expansion Eq. (11) are the coefficients α_i . The coefficients can be computed either by quadrature or regression as described in Sect. 4.2.

For the case of multiple uncertain variables $\boldsymbol{\xi} = (\xi_1, \xi_2, \dots, \xi_n)$ and using a multi-index $\mathbf{i} = (i_1, i_2, \dots, i_n)$, we write the multi-dimensional polynomial approximation as

$$10 \quad R(\boldsymbol{\xi}) \approx \hat{R}(\boldsymbol{\xi}) = \sum_{\mathbf{i} \in \mathcal{I}_p} \alpha_{\mathbf{i}} \Phi_{\mathbf{i}}(\boldsymbol{\xi}). \quad (12)$$

The multi-dimensional basis functions $\Phi_{\mathbf{i}}(\boldsymbol{\xi})$ are given by products of the 1-dimensional orthogonal polynomials

$$\Phi_{\mathbf{i}}(\boldsymbol{\xi}) = \prod_{j=1}^n \phi_{i_j}(\xi_j). \quad (13)$$

When the uncertain variables are independent, the multi-dimensional basis functions are also orthogonal (Sect. 4.4). The values of the elements i_j of the multi-index depend on how the expansion is truncated, i.e., on how the index set \mathcal{I}_p is defined. There are two common ways in which to define the index set: *total-order expansion* and *tensor-product expansion*.

In *total-order expansion* a total polynomial order bound p is enforced:

$$\mathcal{I}_p = \{\mathbf{i} : |\mathbf{i}| \leq p\}, \quad |\mathbf{i}| = i_1 + i_2 + \dots + i_n, \quad (14)$$

which for an expansion of total order p with n uncertain variables results in an expansion with $N_{TO} = \frac{(n+p)!}{n!p!}$ terms.

In *tensor-product expansion* a per-dimension polynomial order bound p_j is enforced

$$20 \quad \mathcal{I}_p = \{\mathbf{i} : i_j \leq p_j, j = 1, \dots, n\}, \quad (15)$$

which results in an expansion with $N_{TP} = \prod_{i=1}^n (p_i + 1)$ terms.

An example showing the multi-dimensional basis polynomials Eq. (13) for both the *total-order expansion* and *tensor-product expansion* can be found in Padrón (2017).

Note that for both *total-order expansion* and *tensor-product expansion* the number of terms exhibits an exponential increase with an increase in the number of uncertain dimensions n . This result is known as the *curse of dimensionality*. The *tensor-product expansion* is the preferred approach when the coefficients are computed with quadrature (Sect. 4.2.1) because of increased monomial coverage and accuracy (Eldred and Burkardt, 2009). The *total-order expansion* is the preferred approach when the coefficients are computed with regression (Sect. 4.2.2) because it keeps the sampling requirements lower (Eldred and Burkardt, 2009).

4.1.1 Mean and variance from the polynomial chaos expansion

The mean and variance of the function of interest $R(\boldsymbol{\xi})$ are a function of the coefficients α_i of the polynomial chaos expansion. The statistics are obtained by substituting the polynomial chaos expansion Eq. (12) into the definitions of the mean Eq. (5) and variance Eq. (8), and by integrating the expansion and simplifying using the orthogonality of the polynomials.

5 The mean is the zeroth coefficient

$$\mu_R = \alpha_0 \quad (16)$$

and the variance is the sum of the product of the square of the coefficients—excluding the zeroth coefficient—with the inner product $\langle \Phi_i^2(\boldsymbol{\xi}) \rangle$,

$$\sigma_R^2 = \sum_{i \in \mathcal{I}_p \setminus \{0\}} \alpha_i^2 \langle \Phi_i^2(\boldsymbol{\xi}) \rangle, \quad (17)$$

10 where the inner product is defined as $\langle \Phi_i^2(\boldsymbol{\xi}) \rangle = \int_{\Omega} \Phi_i(\boldsymbol{\xi}) \Phi_i(\boldsymbol{\xi}) \rho(\boldsymbol{\xi}) d\boldsymbol{\xi}$ and $\mathbf{0}$ is the first multi-index—the one with all zero elements.

4.2 Calculating polynomial chaos coefficients

The coefficients of the polynomial chaos expansion Eq. (12) can be calculated via quadrature or by linear regression.

4.2.1 Quadrature

15 To obtain the coefficients of the polynomial chaos expansion

$$R(\boldsymbol{\xi}) = \sum_{i \in \mathcal{I}_p} \alpha_i \Phi_i(\boldsymbol{\xi}), \quad (18)$$

via quadrature, we take the inner product of both sides of Eq. (18) with respect to $\Phi_j(\boldsymbol{\xi})$ to yield

$$\langle R, \Phi_j \rangle = \sum_{i \in \mathcal{I}_p} \alpha_i \langle \Phi_i, \Phi_j \rangle. \quad (19)$$

Making use of the orthogonality of the polynomials and solving for the coefficients in Eq. (19), we obtain

$$20 \alpha_i = \frac{\langle R(\boldsymbol{\xi}), \Phi_i(\boldsymbol{\xi}) \rangle}{\langle \Phi_i^2(\boldsymbol{\xi}) \rangle} = \frac{1}{\langle \Phi_i^2(\boldsymbol{\xi}) \rangle} \int_{\Omega} R(\boldsymbol{\xi}) \Phi_i(\boldsymbol{\xi}) \rho(\boldsymbol{\xi}) d\boldsymbol{\xi}, \quad (20)$$

where the domain Ω is the Cartesian product of 1D domains Ω_j for each dimension, $\Omega = \Omega_1 \times \cdots \times \Omega_n$, and $\rho(\boldsymbol{\xi}) = \prod_{j=1}^n \rho_j(\xi_j)$ is the joint probability density of the stochastic parameters. The inner product $\langle \Phi_i^2(\boldsymbol{\xi}) \rangle$ is known analytically or inexpensively computed. Thus, most of the computational expense in solving for the coefficients resides in evaluating the model $R(\boldsymbol{\xi})$ in the multi-dimensional integral $\int_{\Omega} R(\boldsymbol{\xi}) \Phi_i(\boldsymbol{\xi}) \rho(\boldsymbol{\xi}) d\boldsymbol{\xi}$. This integral is solved with quadrature (numerical integration). Note that the

25 zero coefficient in Eq. (20) reduces to the definition of the mean

$$\mu_R = \alpha_0 = \int_{\Omega} R(\boldsymbol{\xi}) \rho(\boldsymbol{\xi}) d\boldsymbol{\xi}, \quad (21)$$

which the direct numerical integration methods attempt to compute directly (Sect. 3.1.1).

4.2.2 Regression

To obtain the coefficients of the polynomial chaos expansion Eq. (12) via regression, we construct a linear system

$$\Phi \alpha = R \quad (22)$$

- 5 and solve for the coefficients α that best represent a set of responses R . The set of responses is generated by evaluating the model at m realizations of the uncertain vector ξ . The m uncertain vectors are most commonly obtained by sampling the density of the uncertain variables (Hosder et al., 2007).

Each row of the matrix Φ contains the orthogonal polynomials Φ_j evaluated at a sample ξ_i

$$\begin{bmatrix} \Phi_0(\xi_1) & \cdots & \Phi_{n-1}(\xi_1) \\ \vdots & \ddots & \vdots \\ \Phi_0(\xi_m) & \cdots & \Phi_{n-1}(\xi_m) \end{bmatrix} \begin{bmatrix} \alpha_0 \\ \vdots \\ \alpha_{n-1} \end{bmatrix} = \begin{bmatrix} R_1 \\ \vdots \\ R_m \end{bmatrix}. \quad (23)$$

- 10 The size of the $m \times n$ matrix is determined by the number of samples m and by how the polynomial chaos expansion is truncated (Sect. 4.1) which results in n terms. It is common to specify a total order expansion along with a collocation ratio $cr = m/n$ to determine the number of samples m . The collocation ratio determines if the system is overdetermined $cr > 1$ or underdetermined $cr < 1$.

- For overdetermined systems, the most popular method (and the one we use) to estimate the coefficients is *least squares*, in
 15 which we pick coefficients $\alpha = (\alpha_0, \alpha_1, \dots, \alpha_{n-1})$ that minimize the residual sum of squares

$$\alpha = \arg \min \|\Phi \alpha - R\|_2^2. \quad (24)$$

For underdetermined systems, solving a regularized least squares problem is preferred (Doostan and Owhadi, 2011).

For a given number of samples m in the linear system Eq. (23) we can use cross-validation (Hastie et al., 2009) to pick the best polynomial order n to approximate the response.

20 4.3 Gradients of statistics with polynomial chaos

Let $R(\xi, \mathbf{x})$ be a function of interest that depends on uncertain variables ξ and also on design variables \mathbf{x} . We assume independence between the design and uncertain variables³. Now the polynomial chaos expansion—over the uncertain variables—becomes

$$R(\xi, \mathbf{x}) \approx \hat{R}(\xi, \mathbf{x}) = \sum_{\mathbf{i} \in \mathcal{I}_p} \alpha_{\mathbf{i}}(\mathbf{x}) \Phi_{\mathbf{i}}(\xi). \quad (25)$$

³For most applications, the design and uncertain variables are independent. For instance, the design variables are the wind turbines location and the uncertain variables are the wind conditions.

This expansion is only valid for a particular design vector—the coefficients $\alpha_i(\mathbf{x})$ are a function of the design variables. Therefore, the statistics are also a function of the design variables. Specifically the mean and the variance are

$$\mu_R(\mathbf{x}) = \alpha_0(\mathbf{x}), \quad (26)$$

$$\sigma_R^2(\mathbf{x}) = \sum_{i \in \mathcal{I}_p \setminus \mathbf{0}} \alpha_i^2(\mathbf{x}) \langle \Phi_i^2(\boldsymbol{\xi}) \rangle. \quad (27)$$

5 4.3.1 Gradients of the statistics with polynomial chaos

We want to know the gradients of the statistics with respect to the design variables, and we proceed to derive them below. For simplicity, we drop the subscript from the statistics $\mu_R = \mu$, the explicit variable dependence $R(\boldsymbol{\xi}, \mathbf{x}) = R$, the bolded notation, and we use the following notation for the gradient $\frac{df}{dx} \equiv \nabla f$.

The gradient of the mean from Eq. (26) is

$$10 \quad \frac{d\mu}{dx} = \frac{d\alpha_0}{dx}, \quad (28)$$

and the gradient of the variance from Eq. (27) is

$$\frac{d\sigma^2}{dx} = \sum_{i \in \mathcal{I}_p \setminus \mathbf{0}} \langle \Phi_i^2 \rangle \frac{d\alpha_i^2}{dx} = 2 \sum_{i \in \mathcal{I}_p \setminus \mathbf{0}} \langle \Phi_i^2 \rangle \alpha_i \frac{d\alpha_i}{dx}. \quad (29)$$

Both, the mean and the variance gradients, depend on the gradient of the coefficients $\frac{d\alpha_i}{dx}$.

4.3.2 Gradients of the coefficients

- 15 The gradient of the coefficients can be computed with quadrature or regression, similarly to how the coefficients can be calculated with quadrature (Sect. 4.2.1) or regression (Sect. 4.2.2).

Quadrature. We start from the equation for the coefficients Eq. (20) and take the gradient to obtain

$$\frac{d\alpha_i}{dx} = \frac{1}{\langle \Phi_i^2 \rangle} \int_{\Omega} \frac{dR}{dx} \Phi_i \rho d\boldsymbol{\xi} = \frac{\langle \frac{dR}{dx}, \Phi_i \rangle}{\langle \Phi_i^2 \rangle}. \quad (30)$$

Replacing this equation into the gradient of the mean Eq. (28) we obtain

$$20 \quad \frac{d\mu}{dx} = \left\langle \frac{dR}{dx} \right\rangle. \quad (31)$$

And replacing Eq. (30) into the gradient of the variance Eq. (29) we obtain

$$\frac{d\sigma^2}{dx} = 2 \sum_{i \in \mathcal{I}_p \setminus \mathbf{0}} \alpha_i \left\langle \frac{dR}{dx}, \Phi_i \right\rangle. \quad (32)$$

To obtain the gradients of the statistics with respect to each design variable we need to evaluate the multi-dimensional integral containing $\frac{dR}{dx}$. The integral is evaluated with quadrature (Sect. 4.2.1) and requires the computation of the gradient of

the response at each of the quadrature points. Ideally, one would use adjoint methods Giles and Pierce (2000) or algorithmic differentiation Griewank and Walther (2008) to compute the gradients, $\frac{dR}{dx}$, efficiently.

Regression. We start from the linear system Eq. (22) and take the gradient to obtain

$$\frac{d\Phi\alpha}{dx} = \frac{dR}{dx} \quad (33)$$

$$\Phi \frac{d\alpha}{dx} = \frac{dR}{dx} \quad (34)$$

$$\begin{bmatrix} \Phi_0(\xi_1) & \cdots & \Phi_{n-1}(\xi_1) \\ \vdots & \ddots & \vdots \\ \Phi_0(\xi_m) & \cdots & \Phi_{n-1}(\xi_m) \end{bmatrix} \begin{bmatrix} \frac{d\alpha_0}{dx_1} & \cdots & \frac{d\alpha_0}{dx_d} \\ \vdots & \ddots & \vdots \\ \frac{d\alpha_{n-1}}{dx_1} & \cdots & \frac{d\alpha_{n-1}}{dx_d} \end{bmatrix} = \begin{bmatrix} \frac{dR_1}{dx_1} & \cdots & \frac{dR_1}{dx_d} \\ \vdots & \ddots & \vdots \\ \frac{dR_m}{dx_1} & \cdots & \frac{dR_m}{dx_d} \end{bmatrix}. \quad (35)$$

To solve for the gradient of the coefficients, we solve the linear system one column at a time of the $\frac{d\alpha}{dx}$ matrix with the corresponding column of the matrix of the gradients $\frac{dR}{dx}$. The linear system for the multiple right hand sides can be solved with the methods described in Sect. 4.2.2.

Again, the gradient of the mean Eq. (28) and the gradient of the variance Eq. (29) are a function of the gradient of the coefficients. Thus, to obtain the gradient of the mean take the first row of the $\frac{d\alpha}{dx}$ matrix; and for the gradient of the variance use the gradient of the coefficients from all the other rows.

4.3.3 Gradients of the statistics by direct numerical integration

Similarly to computing the mean and variance with direct numerical integration (Sect. 3.1.1), we can also compute the gradients of the mean and variance directly with numerical integration by differentiating the definitions of the mean, Eq. (5), and the variance, Eq. (8).

4.4 Correlated random variables

As described, the polynomial chaos method assumes that the uncertain variables are statistically independent. In wind farm layout optimization, the uncertain variables of wind speed and wind direction are usually correlated. There are different approaches to use polynomial chaos for problems with inputs made up of correlated (dependent) random variables. One approach is to perform a (linear or non-linear) variable transformation (Feinberg and Langtangen, 2015) to uncorrelate the variables before applying polynomial chaos. Another approach is to construct polynomials that are orthogonal to the multivariate distribution instead of orthogonal polynomials for each dimension (Navarro et al., 2014). A third approach, which is the most applicable for the wind farm layout optimization problem, consists of breaking up the domain of the uncertain variables into smaller domains (elements). Then, constructing a polynomial chaos expansion for each of the small domains and combining them in a piece-wise manner to cover the whole domain. This approach of breaking up the domain is known as multi-element polynomial chaos (Wan and Karniadakis, 2006). In the context of the wind farm layout optimization problem, the new distributions for the uncertain variables in the small elements can be considered independent. For example, the wind direction and wind speed at a particular wind farm site are usually correlated, but if the wind direction domain is broken into sectors, it can

be assumed that within each sector the wind direction and wind speed are independent (Herbert-Acero et al., 2014; Murcia et al., 2015).

5 Problem details

Here, we describe the details of computing AEP in the wind farm. We first describe and discuss the inputs, which are the wind direction and wind speed (uncertain variables) (Sect. 5.1) and the wind turbine positions (design variables) (Sect. 5.2). Then, we describe details about the output, the average AEP error (Sect. 5.3), that we use to compare the different methods we consider to compute the AEP (Sect. 5.4).

5.1 Probability distributions of the uncertain wind conditions

We consider the wind direction and the wind speed as uncertain variables with probability distributions shown in Fig. 3. The distributions show the likelihood of a particular wind direction or wind speed occurring during a year. For simplicity, we assume that the wind direction and wind speed distributions are independent (Sect. 4.4). We construct the independent distributions starting from the wind measurements taken near the Princess Amalia wind farm (Sect. 5.2) by the NoordZeeWind meteorological mast during a year (Brand et al., 2012). For a wind rose of the measurements, see Gebraad et al. (2017). The wind rose has a wind direction bin width of 5° and a wind speed bin width of 1 m/s throughout the operational range of the turbine.

We construct the wind direction distribution (Fig. 3a) by linearly interpolating the wind measurements of the likelihood of the wind coming from a particular binned wind direction for all wind speeds. The zero degree direction is set at North and increases clockwise.

For the wind speed, instead of linearly interpolating the data, we fit a Weibull distribution⁴ and then truncate it⁵ (Fig. 3b). We truncate the distribution to correspond to the operational range of the turbine from the cut-in speed of 3 m/s to the cut-out speed of 25 m/s.

The probability distributions are the weight functions of the orthogonal polynomials used in the polynomial chaos method (Sect. 4.1). For both of the distributions (weight functions), we use a histogram of 50 equal width bins to describe each distribution. Using these weight functions, we then numerically generate their corresponding orthogonal polynomials.

5.2 The wind farm layouts

To showcase the results, we will focus on four representative layouts: Grid, Amalia, Optimized and Random (Fig. 4). Each layout has sixty turbines and in the figures each turbine is represented by a dot to scale—the diameter of the dot represents

⁴The fitted Weibull— $\rho(\xi; \alpha, \beta) = \frac{\alpha}{\beta} \left(\frac{\xi}{\beta}\right)^{\alpha-1} e^{-(\xi/\beta)^\alpha}$, $\xi \geq 0$ —has shape parameter $\alpha = 1.8$ and scale parameter $\beta = 12.55$. The Weibull distribution is the preferred distribution to model the wind speed distribution (Belu and Koracin, 2013; Herbert-Acero et al., 2014).

⁵The truncated distribution is no longer a full Weibull distribution and needs to be scaled to ensure it is a valid probability density function (it integrates to 1).




Figure 3. The uncertain variables probability distributions. The vertical lines show the cut-in and cut-out speed of a single wind turbine.




Figure 4. Representative wind farm layouts used in the results. Each dot represents a wind turbine to scale—the dot represents the swept area of the rotor.

the rotor diameter. The Grid layout fits in a box of equivalent area to that of the Amalia layout. The Amalia layout, which is grid-like, is that of the Princess Amalia wind farm located 23 km off the coast of the Netherlands. The Optimized layout is a representative optimal layout obtained by running the optimization problem (Sect. 6.3.1); specifically, it is the layout shown in (Fig. 11a). When we refer to this particular optimized layout, we will capitalize the word optimized. The Random layout was generated by random sampling and keeping the turbines that are contained within the convex hull of the Amalia wind farm while also satisfying a minimum inter-turbine spacing of three diameters. We will refer to the Grid and Amalia as grid-like layouts and the Optimized and Random as non-grid-like layouts.

In reality, the 60 turbines in the Princess Amalia wind farm are the Vestas V80 model. For each of the layouts in our study, we use the NREL 5-MW reference turbine (Jonkman et al., 2009), as this turbine has an open source design.

5.3 Convergence metric — the average AEP error

We use an ensemble of 10 AEP results to compute the average AEP error. The average AEP error allows us to better illustrate the differences between the different methods used to compute the AEP and to avoid drawing conclusions from one-off solutions. We found that averaging over 10 AEP results is enough to illustrate the difference between methods and to smooth out the convergence of the AEP error (Fig. 8). The ensemble of results account for the fact that the zero (starting) position for the wind direction is arbitrary (it could be North, South, East, West, etc.). Also, averaging the AEP errors helps smooth out the AEP convergence curves by reducing the sensitivity of the AEP to the quadrature (sample) points used to compute the AEP. We generate the ensemble of AEP results by selecting 10 different input sets. For example, for the rectangle rule, if we consider 36 wind directions the 10 sets are $\{[0, 10, 20, \dots, 340, 350 \text{ degrees}]; [1, 11, 21, \dots, 341, 351]; \dots; [9, 19, 29, 349, 359]\}$. For the polynomial chaos based on quadrature, the quadrature points are the numerically generated Gaussian quadrature points for the interval. Thus, to create 10 different sets, we pick 10 different intervals, i.e., we pick different starting positions. ~~When considering 36 wind directions the chosen intervals are $\{[0, 360]; [10, 370]; [20, 380]; \dots; [340, 700]; [350, 710]\}$~~ We set the starting positions for the intervals at $i \times 360/10$ degrees for $i = 0..9$. For both the rectangle and polynomial chaos based on quadrature, the wind speed points for each set are the same. For the polynomial chaos based on regression and for Monte Carlo, the wind directions and wind speed pairs are generated by sampling the distribution. Thus, to obtain 10 different sets, 10 different samplings are performed. We use the average AEP error as the convergence metric

$$\text{average AEP error} = \frac{1}{10} \sum_{i=1}^{10} \left| \frac{AEP_i - \text{baseline AEP}}{\text{baseline AEP}} \right| \times 100\%. \quad (36)$$

5.3.1 Baseline AEP

We take as the baseline or true AEP the AEP computed with 200,000 Monte Carlo samples. We picked 200,000 MC samples to ensure the 99 % confidence interval for the true AEP was smaller than 1 % of the computed AEP value for all layouts. We consider an AEP within 1 % of the baseline AEP to be accurate, and we will use it as a reference for the results. AEP predictions of real wind farms usually have an error of 10–20 % (Barthelmie et al., 2007; Briggs, 2013) due to uncertainty in wind conditions and to the errors of wake models, thus resolving the AEP to less than 1 % is unnecessary.

5.4 Methods to compute the AEP

Here, we provide the details of the methods used to compute the AEP, as well as, the abbreviations for the methods:

rect	rectangle rule	(Sect. 3.1.1)
PC-Q	polynomial chaos based on quadrature	(Sect. 4.2.1)
PC-R	polynomial chaos based on regression	(Sect. 4.2.2)
MC	Monte Carlo	(Caffisch, 1998)

For the quadrature-based methods (rect, PC-Q), we use tensor product quadrature to compute multi-dimensional integrals. We use the same number of points for each dimension because we did not see any benefit in favoring a particular dimension. For

PC-Q, we use Gaussian quadrature. For Monte Carlo, we use the traditional Monte Carlo method, i.e., random samples. For PC-R, we use Latin-Hypercube sampling (McKay et al., 1979) to generate the samples needed to construct the linear system. We solve the linear system with least-squares. For a given number of samples, given by the total polynomial order p , we use 10-fold cross-validation to find the least-squares best fit from polynomials of total order 1 up to total order p . The sampling methods and the polynomial chaos methods we use are implemented in the open-source DAKOTA toolkit (Adams et al., 2017).

6 Results

~~The Annual Energy Production (AEP) of a wind farm is obtained by integrating the power output of the farm over the different wind conditions. Thus, we~~ We first characterize the power output of the wake model for different input conditions (Sect. 6.1). Next, we focus on the convergence of the AEP (Sect. 6.2), and then on the wind farm layout optimization problem to maximize the AEP (Sect. 6.3).

6.1 Power response as a function of the uncertain variables

The power production, computed with the Floris wake model, for the four wind farm layouts (Sect. 5.2) as a function of both the wind direction and the wind speed is shown in Fig. 5. The wind direction is measured from North and increases clockwise. The power response as a function of the wind direction is highly oscillatory, and for layouts with structure⁶, it is periodic. The peaks in the contour lines identify wind directions for which there is a poor performance (low power). The grid-like layouts (top row of Fig. 5) have larger peaks due to wind turbines being aligned along particular directions and thus experiencing full-wake conditions. The worst wind direction for the Grid layout is directly from North (0°) or South (180°) when rows of ten turbines are aligned.

The power response as a function of the wind direction has similar shapes for different wind speeds, until the speed is high enough (larger than the rated speed of the layout) that the response gets clipped at 300MW (all sixty turbines are operating at their rated power 5MW). As a function of the wind speed, for each wind direction, the power response starts after the cut-in speed of 3 m/s and is cubic (concave up) until individual turbines start reaching rated power. After that, the power increases more slowly (concave down) until the farm reaches its rated power. The less grid-like layouts show a smoother transition to the maximum power because each turbine reaches its rated speed (11.4 m/s) at different freestream speeds.

The power response as a function of the uncertain variables (wind direction and wind speed) is a complicated function, which needs to be integrated, weighted by the likelihood of the wind conditions, to obtain the AEP of the wind farm. Next, we show that using polynomial chaos, instead of the rectangle rule, results in better estimates of the AEP.

6.2 AEP convergence: polynomial chaos vs. rectangle rule

We consider the AEP as a function of two uncertain variables: the wind speed and the wind direction. The AEP is usually considered a function of these two variables. We compare the convergence of the AEP for the different methods to compute

⁶Wind turbines are aligned in particular directions. The layouts are grid-like and have symmetry planes.




Figure 5. Power contours as a function of wind direction and wind speed for the Grid, Amalia, Optimized and Random layouts. The grid-like layouts (top row) have larger peaks due to wind turbines being aligned along particular directions and thus experiencing full-wake conditions. For all layouts as the speed increases the power increases until the wind farm reaches its rated power 300 MW.

the AEP: polynomial chaos (based on quadrature and regression), the rectangle rule and Monte Carlo (Fig. 6). The polynomial chaos based on regression (PC-R) performs the best for all layouts. It is followed, by the polynomial chaos based on quadrature (PC-Q) and the rectangle rule, which perform similarly. The worst performer is Monte Carlo. The slow convergence of statistics with Monte Carlo is well known. Monte Carlo will start to outperform the other methods when the AEP is a function of a large number (5–10) of uncertain variables, as it does not suffer from the *curse of dimensionality*.

The superior performance of the polynomial chaos based on regression, especially for the grid like-layouts (Grid and Amalia), is due to the following: the polynomial chaos fit based on regression does not chase all the high-frequency oscillations in the power response (Fig. 5), i.e., it smooths out the response. The PC-R fit is usually not higher than an eight total-order polynomial (Sect. 4.1). Whereas, the PC-Q order fit is higher, as it is directly proportional to the number of samples per dimension⁷. A downside of the PC-R being able to predict the mean (AEP) accurately, is that it can underpredict the true variance (standard deviation) of the response (Fig. 7). Usually, the standard deviation of the power response (energy) over

⁷For the two-dimensional problem at 625 samples (25×25 grid) the polynomial order in each dimension is 24.

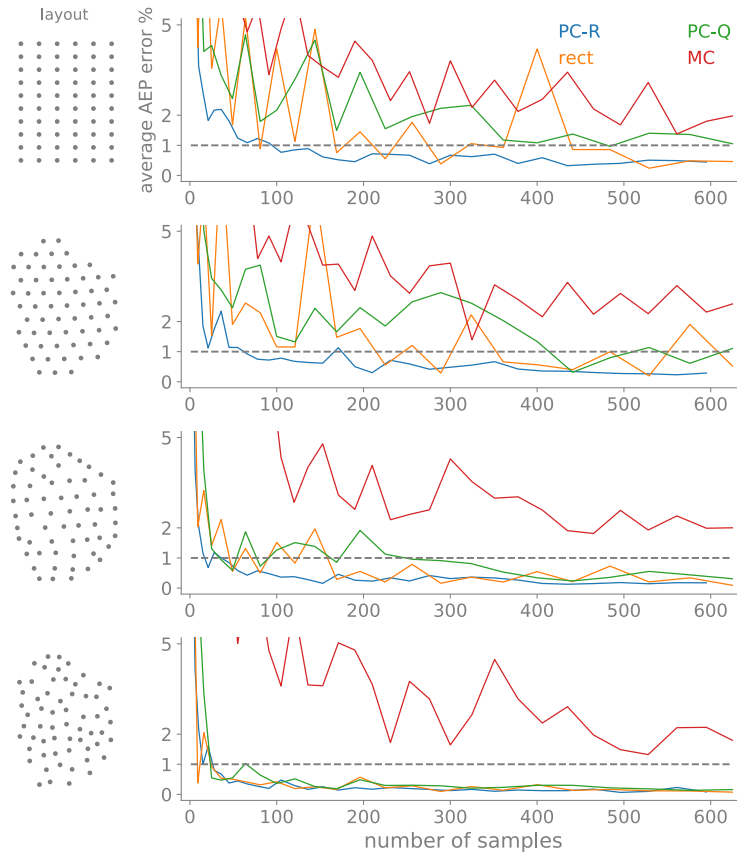


Figure 6. The average AEP error as a function of the number of samples for polynomial chaos (based on regression and quadrature), the rectangle rule and Monte Carlo. The AEP is a function of two uncertain variables the wind direction and wind speed. The polynomial chaos based on regression performs best for all layouts.

a year is not considered as a function to optimize. A common objective in wind farm optimization is to maximize the total amount of energy produced over a year independent of the variability in the power production over the year. For a wind farm, the variability of the energy produced over a year is less important than the variability caused by the changing wind conditions during the day.

- 5 In what follows we will compare the PC-R (the best performing method) with the rectangle rule (the method **currently most commonly** used in practice) to quantify the reduction of samples needed to compute the AEP accurately. Also, we will sometimes use polynomial chaos to refer to polynomial chaos based on regression. Figure 8 only keeps PC-R and the rectangle rule results from Fig. 6, and in addition, the figure shows the average AEP error computed with 10 and 100 sets of samples for each method. For the rectangle rule, there is hardly any difference between the average AEP error computed with 10 or 100 sets. For the PC-R method, the average error with 100 sets shows a smoother convergence. In general, averaging the AEP error

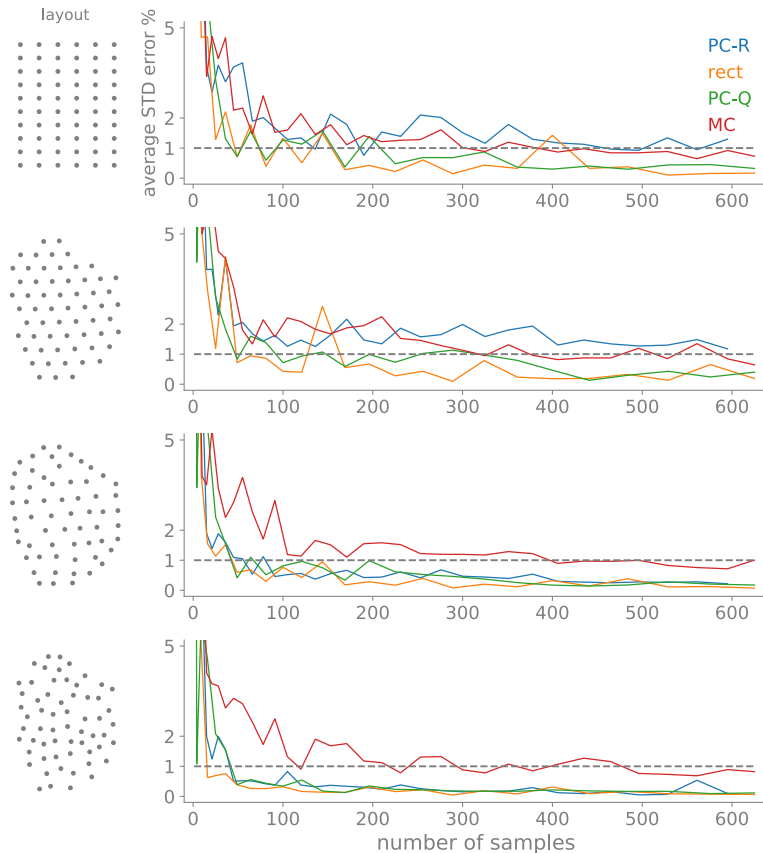


Figure 7. The average standard deviation of the energy (STD) error as a function of the number of samples for polynomial chaos (based on regression and quadrature), the rectangle rule and Monte Carlo. The STD is a function of two uncertain variables the wind direction and wind speed. Note that the PC-R response is biased for the Grid and Amalia layouts. The PC-R underpredicts the true STD at the expense of computing the mean (AEP) more accurately (see Fig. 6).

over 10 sets of samples is enough to minimize the AEP’s sample location sensitivity and to clearly see the differences between the methods used for computing the AEP (Sect. 5.3).

In Fig. 8, we see that the PC-R convergence curve is consistently below the rectangle rule curve, i.e., PC-R has a smaller error for the same number of samples or the same error for a smaller number of samples. Using the 1 % average AEP error as a metric, we see that PC-R achieves this error with fewer samples than the rectangle rule. The reduction in the number of samples is on the order of four times for the Grid layout, eight times for the Amalia, four times for the Optimized and no significant improvement for the Random. These reductions in the number of samples are considerable. In addition to providing faster convergence of the AEP, the polynomial chaos based on regression method converges the AEP more smoothly—less oscillatory, more monotone convergence. Smooth convergence is always a desired property, and it especially useful when




Figure 8. The average AEP error as a function of the number of samples for both the rectangle rule and polynomial chaos based on regression (this is the same as Fig. 6, with only keeping the PC-R and the rectangle rule results). In addition, we show the average AEP error computed with 10 (solid line) and 100 (dashed line) sets of samples (Sect. 5.3). The AEP is a function of two uncertain variables the wind direction and wind speed. The polynomial chaos method computes the AEP more accurately with fewer samples.

performing an optimization. Both methods, PC-R and the rectangle rule, perform better for the less grid-like layouts because there is less variability in the power responses as a function of wind direction (see Fig. 5). This is beneficial as in the first few iterations of an optimization the starting grid-like layouts become non-grid-like and start to resemble the Optimized layout.

The polynomial chaos based on regression is not only better on average, as we have seen in Fig. 8, but it is also better in general, as shown in Fig. 9. The shaded area in Fig. 9 shows the spread between the maximum and minimum AEP for 10 realizations of each of the number of samples (Sect. 5.3). The solid line shows the average of those 10 realizations and the dashed lines the $\pm 1\%$ of the baseline AEP. We see that the spread is significantly smaller for the polynomial chaos based on regression and that by around 300 samples the predictions are almost always within 1% of the true AEP for the grid-like layouts (Grid and Amalia) and around 100 samples for the non-grid-like layouts (Optimized and Random). In contrast, for the

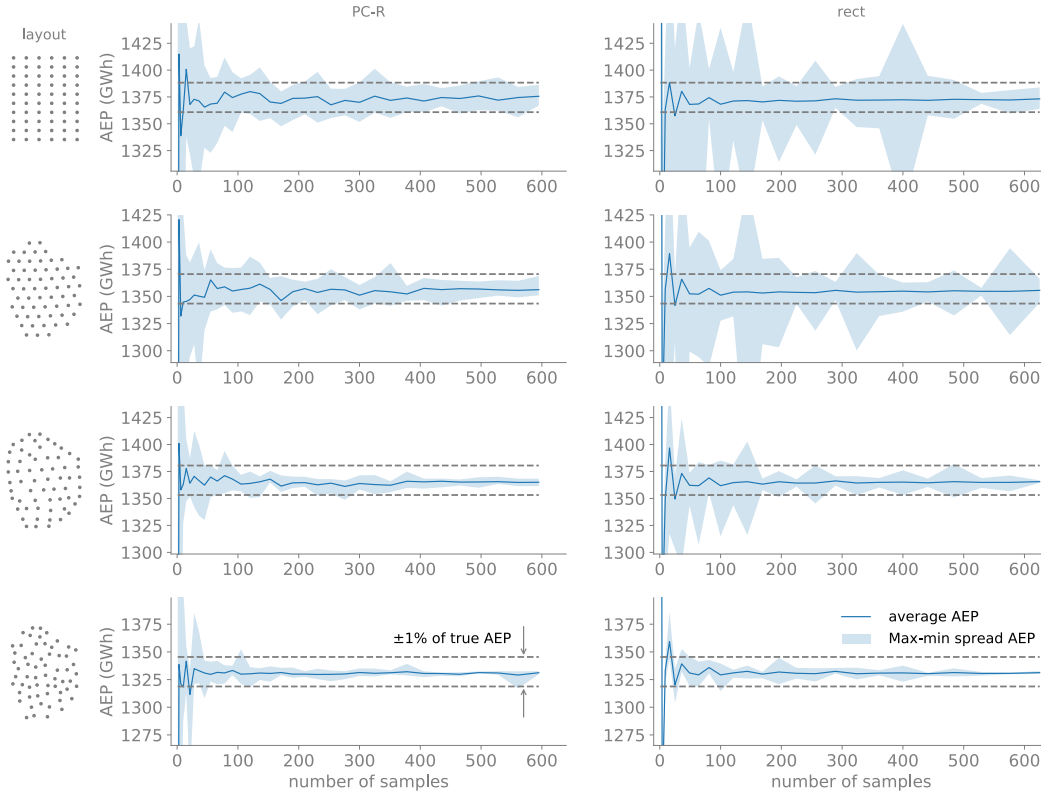


Figure 9. Variability in the convergence of the AEP. The figures of the left column are for polynomial chaos based on regression and those of the right column for the rectangle rule. The shaded area shows the spread between the maximum and minimum AEP for 10 realizations of each of the number of samples. The range in the plots corresponds to $\pm 5\%$ of the true AEP for each layout. The variability is significantly smaller for polynomial chaos, which shows that in general, it outperforms the rectangle rule. The PC-R method obtains the smaller variability by smoothing out the high-frequency oscillations present in the power response (Fig. 5).

rectangle rule, the error in the AEP is still larger than 1 % at 400 samples for the Grid layout and 600 samples for the Amalia layout.

6.3 Wind farm layout optimization

6.3.1 Optimization problem

- 5 The objective of the wind farm layout optimization is to maximize the AEP (Sect. 2.3) by changing the position of the wind turbines. We assume a fixed number of turbines, 60, of the same type (NREL 5-MW (Jonkman et al., 2009)) and constrain the turbines to stay within a given area and with a minimum separation between them. This objective and constraints result in a

nonlinear optimization under uncertainty problem with deterministic constraints:

$$\begin{aligned}
 & \underset{x,y}{\text{maximize}} && AEP(x,y,\xi) \\
 & \text{subject to} && S_{i,j} \geq 2D \quad i,j = 1 \dots nTurbines, i \neq j \\
 & && N_{i,b} \geq 0 \quad i = 1 \dots nTurbines, b = 1 \dots nBoundaries,
 \end{aligned} \tag{37}$$

where $S_{i,j}$ is the distance between each pair of turbines i and j , and D is the turbine diameter. The normal distance, $N_{i,b}$, from each turbine i to each boundary b is defined as positive when a turbine is inside the boundary, and negative when it is outside of the boundary. The boundary is the convex hull of the Princess Amalia layout—a fourteen-sided convex polygon (dashed-line boundary in the upper-left of Fig. 10). The design variables are the x, y coordinates of the 60 wind turbines, resulting in 120 design variables. The uncertain variables ξ in the objective are the wind direction and wind speed.

We solve the optimization problem with the gradient-based sequential quadratic programming optimizer SNOPT (Gill et al., 2005). We use OpenMDAO (Gray et al., 2010) and its wrapper for pyOptSparse (Perez et al., 2012) to call SNOPT from Python. We scale the variables, constraints and the objective to make them of order one, and set the tolerances to 1×10^{-4} per the function (AEP) precision.

6.3.2 Optimization results

We solve the optimization under uncertainty problem, Eq. (37), using the rectangle rule and polynomial chaos based on regression methods to compute the AEP. For each method, we consider a coarse and fine number of samples to compute the AEP. Also, for each method, we run 10 optimizations, where each optimization uses different sample points to compute the AEP (see Sect. 5.3). The 10 optimizations enable us to get a better understanding of which method is better at finding layouts with high AEP and to avoid drawing conclusions from one-off local optima. Furthermore, we consider three different starting layouts for the optimization: Amalia, Grid and Random (Sect. 5.2). The results of the optimizations are reported in Table 2.

The optimum layouts obtained with polynomial chaos based on regression are better than those obtained with the rectangle rule. For a similar number of samples to compute the AEP, the optimizations with polynomial chaos produce on average optimal layouts with 0.4 % higher AEP than the optimizations with the rectangle rule. Also, using about one-third of the samples, the optima obtained with polynomial chaos are comparable to those obtained with the rectangle rule. Furthermore, polynomial chaos finds better local optima than the rectangle rule for all the different starting layouts considered (Table 2).

The optimal layouts with the maximum AEP for each method and starting layout are shown in Fig. 10. As can be seen in that figure, the wind farm layout optimization problem contains many local optima. In general, the turbines—in the local optima—position themselves in non-grid-like patterns to minimize wake losses. Also, to minimize wake losses, the turbines try to spread out as much as possible and position themselves at the boundary of the layout (Fig. 11). For the optima obtained starting from the Amalia layout, we see that the distribution of the turbines is somewhat similar to the starting layout. As the Amalia layout is already a good layout⁸, it is not surprising that the local optima somewhat resemble the Amalia layout.

⁸[The Amalia layout is an optimal layout for a different optimization problem: the designers of the Princess Amalia wind farm likely considered different optimization constraints and objective, used a different wake model in the analysis and as mentioned in Sect. 5.2 used a different wind turbine.](#)

Table 2. The AEP of the optimized layouts. We generate the AEP statistics for each method from a set of 10 different optimizations. For a similar number of samples, both coarse and fine, polynomial chaos based on regression produces better optimal layouts than the rectangle rule. Furthermore, using about one-third of the samples, PC-R optima are comparable with the optima obtained with the rectangle rule that used 625 samples to compute the AEP. The values of the AEP reported are computed with 200,000 Monte Carlo samples.

Method	# Samples	Starting layout: Amalia		Starting layout: Grid		Starting layout: Random	
		AEP (GWh)		AEP (GWh)		AEP (GWh)	
		mean \pm std	max	mean \pm std	max	mean \pm std	max
PC-R-coarse	231	1359 \pm 1.5	1362	1362 \pm 2.0	1365	1357 \pm 2.0	1360
PC-R-fine	630	1362 \pm 1.8	1367	1367 \pm 1.2	1369	1363 \pm 1.8	1367
rect-coarse	225	1356 \pm 1.1	1358	1354 \pm 2.2	1357	1349 \pm 1.4	1351
rect-fine	625	1360 \pm 1.0	1361	1365 \pm 1.3	1367	1358 \pm 1.0	1360

The optima found starting from the Grid layout have on average the highest AEP (Table 2 and Fig. 10). When considering all methods, on average, the optimal layouts found starting from the Grid layout are 0.2 % higher than the optimal layouts starting from the Amalia layout and 0.4 % higher than the optimal layouts starting from the Random layout. The Grid layout is infeasible, as many of the turbines do not satisfy the boundary constraint (Fig. 11). However, starting from the infeasible layout is beneficial as it nudges the turbines to position themselves on the boundary. The optima found starting from the Random layout have on average the lowest AEP, but they have the most improvement in AEP over the starting layout showing that we can find a good layout even from a poor starting position (Fig. 12). Due to the presence of many local optima, starting from a good layout will usually find layouts that are better than those starting from a bad layout, such as the Random.

To properly compare the results obtained by the different methods used to compute the AEP in the optimization, we should compute the AEP of the optimal layout with a method that was not used in the optimization. The values of the AEP reported in Table 2 and Fig. 10 are computed with 200,000 Monte Carlo samples. For consistency, and to avoid potential uncertainty in the results introduced by the sample locations, we use the same 200,000 samples—wind direction, wind speed pairs—to evaluate the optimal layouts obtained by each method. Using the AEP computed with Monte Carlo to measure the improvement of the optimal layout over the starting layout will in general be lower than the improvement in AEP as measured by the method used in the optimization (Table 3). For example, for the optimal layouts in Fig. 10, with the Amalia starting layout, the average improvement in AEP over the starting layout for all the methods is 0.36 % when measured with the Monte Carlo samples and 1.50 % when measured by the method used in the optimization. For the optimizations starting from the Random layout, the average AEP improvements are 2.06 % for Monte Carlo and 3.20 % for the method used in the optimization. As the optimizations starting from the Grid layout are infeasible, we do not include the improvement in AEP for the Grid layout in Table 3. The improvements we have found starting from the Amalia layout are similar to those found by Gebraad et al. (2017) for turbine position optimization.

Most optimizations required on the order of one hundred ~~function calls to compute AEP~~ AEP computations. The optimization history for the layouts in Fig. 11 are shown in Fig. 12. Note that each AEP computation requires hundreds (the number

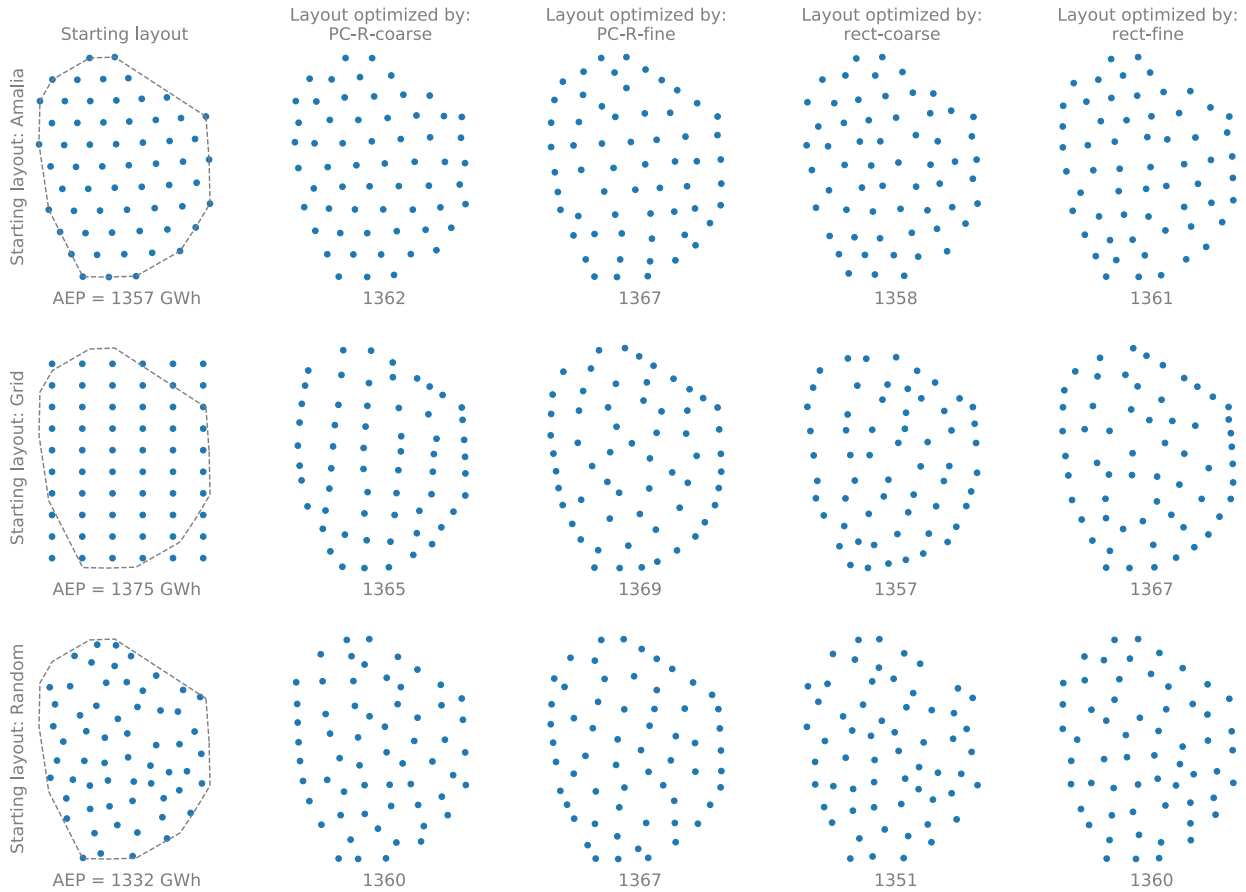


Figure 10. Optimal wind farm layouts achieved for each method to compute the AEP. The layouts in the first column show the turbines starting position for the optimization along with the boundary constraint (dashed-line). The optimal layouts correspond to those with the maximum AEP from Table 2.

of samples specified for each method in Table 2) of calls to the wake model resulting in tens of thousands of calls to the wind farm wake model per optimization. If the gradients of the AEP are computed with a first-order finite difference, 120 (the number of design variables) times as many wind farm wake model calls would be needed per optimization. We showed that using polynomial chaos finds comparable optima to the rectangle rule but using only about one-third of the samples per AEP computation. Thus, polynomial chaos finds comparable optima for roughly one-third of the optimization computational cost. And, for the same computational cost, polynomial chaos produces optimal layouts with 0.4 % higher AEP on average.

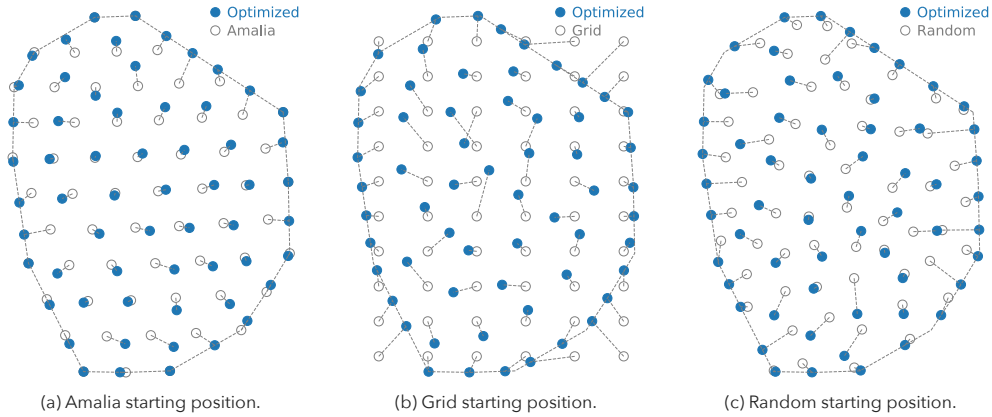


Figure 11. The initial and optimized position of the turbines for the layouts obtained with the PC-R-fine method from Fig. 10. To minimize wake losses, the turbines try to spread out as much as possible and position themselves at the boundary of the layout.

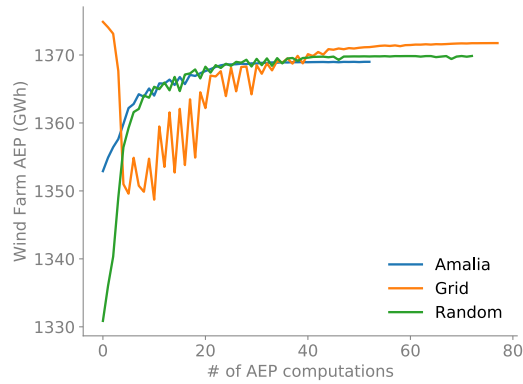


Figure 12. Convergence history of the optimization for the different starting layouts. The convergence history is for the layouts in Fig. 11.

Table 3. The improvement in AEP of the optimized layouts from Fig. 10 over their starting layouts. The improvements are smaller when measured by 200,000 Monte Carlo samples rather than directly by the method used to compute the AEP in the optimization.

Method	Starting layout: Amalia		Starting layout: Random	
	AEP % improvement as measured by Monte Carlo	AEP % improvement as measured by method	AEP % improvement as measured by Monte Carlo	AEP % improvement as measured by method
PC-R-coarse	0.36	0.75	2.08	2.86
PC-R-fine	0.73	1.19	2.62	2.93
rect-coarse	0.07	2.45	1.46	3.93
rect-fine	0.30	1.62	2.09	3.07
average	0.36	1.50	2.06	3.20

7 Discussion and conclusions

A single wind farm layout optimization requires tens to hundreds of thousands of model evaluations. During the design phase of a new wind farm, designers would need to perform many optimizations. The designers may explore scenarios with different turbine types, different sites, larger farms with a different number of turbines and possibly even systems of wind farms. Also, the presence of local optima would require many optimizations with different restarts to find the best layout. And, furthermore, there is a desire to increase the fidelity of the models used to simulate the wind farm, which will increase the time and computational cost of the optimizations⁹. Thus, to facilitate the design phase of new wind farms and to incorporate the use of higher fidelity models in the design phase requires significant improvements to reduce the simulation requirements (computational cost) of performing the wind farm optimization under uncertainty (OUU) problem (Fig. 1c).

We proposed the use of polynomial chaos (PC) to improve the uncertainty quantification step of the OUU problem. We computed AEP convergence plots and showed that polynomial chaos based on regression (PC-R) could compute the AEP accurately using on average one-fifth of the simulations required using the rectangle rule, the method currently used in practice in the wind industry. PC-R computes the AEP accurately (error less than 1 %) usually with only two hundred simulations or less depending on the wind farm layout—a significant improvement over the current industry practice of using more than a thousand model evaluations to compute the AEP.

The layout of the wind farm influences the convergence of the AEP because the layout has a significant effect on the power output of the farm as the wind conditions vary (Fig. 5). We considered four representative layouts: Grid, Amalia, Optimized and Random. The power response of the grid-like layouts (Grid and Amalia) has large oscillations caused by the large drops in power that occur when rows of wind turbines are aligned with particular wind directions. Because of the larger variability in the power response, the grid-like layouts require more simulations than the non-grid-like layouts (Optimized and Random) to converge the AEP. An extension of this work that could potentially further improve the convergence of the AEP, especially for the grid-like layouts, would be to build an approximation to the power output that takes into consideration the oscillatory and periodic behavior of the wind farm response with respect to the wind direction (Fig. 5). For instance, by using Fourier series to approximate the wind direction response and polynomials for the wind speed response.

In addition to computing the AEP efficiently, a benefit of using polynomial chaos is that it can also compute the gradient of the AEP efficiently. The use of gradients is essential to enable large-scale wind farm optimization. We described how to compute the gradients of the AEP in (Sect. 4.3). A benefit of using the rectangle rule is that no special consideration is necessary for correlated uncertain variables; whereas, for polynomial chaos, some extra considerations are necessary as discussed in Sect. 4.4. Another benefit of using PC is that it can easily incorporate multiple fidelity models to accelerate the convergence of statistics (Ng and Eldred, 2012; Palar et al., 2016), such as the AEP (Padrón, 2017). Furthermore, the effect of other uncertain variables (such as wake model parameters) on the AEP, can easily be incorporated in the polynomial chaos framework. As the number of uncertain variables increases, polynomial chaos based on regression will continue to outperform the rectangle rule

⁹If instead of using the Floris wake model in the optimization, we had used a high-fidelity fully resolved Large Eddy Simulation to model the wind farm, a back of the envelope calculation shows that the optimization would require the total Annual Energy Production (AEP) of the wind farm to power the computers used in the simulation.

(see Padrón (2017) for an example with three uncertain variables—wind speed, wind direction, and wake model parameter). For a large number (5-10) of uncertain variables, the Monte Carlo method will start to outperform polynomial chaos because it does not suffer from the *curse of dimensionality* (Smith, 2014).

We performed multiple gradient-based wind farm layout optimizations to compare the optimization results obtained with the different methods to compute the AEP: polynomial chaos based on regression and the rectangle rule. The goal of the optimization was to maximize the Annual Energy Production (AEP) of the wind farm. The layout optimization is an optimization under uncertainty (OUU) problem (Fig. 1c) with the wind speed and wind direction as uncertain variables and the positions of the wind turbines as design variables. In the resulting optimal layouts, the turbines position themselves in non-grid-like patterns and spread toward the boundary to minimize wake losses. For the optimization problem posed, ~~the~~ there are many local optima with different patterns and similar AEP (Fig. 10). To search for the best ~~optima~~, optimum, we considered different starting layouts ~~should be considered~~ (Grid, Amalia and Random). We observed that starting from a good layout (Grid, Amalia) will, in general, find better optima than starting from a bad layout (Random) ~~independent of the method used to compute the AEP~~.

An ensemble of optimizations should be used to properly compare the optimal layouts obtained using the different methods to compute the AEP. Also, for proper comparison, the AEP of the optimized layouts should be evaluated in the same way and with a method different from the one used in the optimization. To be confident about the differences between the optimal layouts, we evaluated the AEP of each optimized layout using the same 200,000 Monte Carlo samples. We found that the benefits of being able to efficiently compute the AEP with PC-R translate to being able to find better optima than those obtained when computing the AEP with the rectangle rule when using the same number of simulations. ~~On~~ For the cases considered, on average, with PC-R we find optima that are 0.4 % better than those found with the rectangle rule. This is a significant improvement as a 1 % increase in the AEP for a modern large wind farm can increase its annual revenue by millions of dollars. Also, using about one-third of the simulations, coarse PC-R finds comparable optima to those found by the ~~rectangle rule~~ fine rectangle rule for the optimization cases considered.

An interesting extension of this work, that could further improve the AEP, would be to allow yaw-based control in the optimization problem to take advantage of wake deflection. Gebraad et al. (2017) found that including yaw for each wind direction in a secondary layout optimization, after layout optimization without yaw, resulted in an additional 3.7 % increase in AEP. Including wake deflection in the wind farm layout optimization increases the complexity of the problem in terms of the number of design variables because every wind direction (and to a lesser extent, wind speed) calls for different yaw values for each turbine. The increased number of design variables makes the problem much more difficult for gradient-free optimization methods and can slow down gradient-based optimization methods as well (Thomas et al., 2017). As we have seen, polynomial chaos reduces the number of simulations (wind directions and wind speeds considered). Thus, it could also greatly reduce the number of simulations in optimization with yaw control.

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

Acknowledgements. Funding for this research was provided by the National Science Foundation under a collaborative research grant No. 1539384 and 1539388.

References

- Adams, B. M., Ebeida, M. S., Eldred, M. S., Jakeman, J. D., Swiler, L. P., Stephens, J. A., Vigil, D. M., Wildey, T. M., Bohnhoff, W. J., Dalbey, K. R., Eddy, J. P., Hu, K. T., Bauman, L. E., and Hough, P. D.: Dakota, a multilevel parallel object-oriented framework for design optimization, parameter estimation, uncertainty quantification, and sensitivity analysis: version 6.6 user's manual, Tech. rep., Sandia National Laboratories, 2017.
- Ascher, U. M. and Greif, C.: Chapter 15: numerical integration, in: *A First Course Numer. Methods*, pp. 441–479, SIAM, Philadelphia, PA, <https://doi.org/10.1137/9780898719987.ch15>, 2011.
- AWEA: AWEA U.S wind industry annual market report year ending 2015, Tech. rep., American Wind Energy Association, <http://awea.files.cms-plus.com/AnnualReportTopFacts2015.pdf>, 2016.
- 10 Barthelmie, R. J., Frandsen, S. T., Nielsen, M. N., Pryor, S. C., Réthoré, P. E., and Jørgensen, H. E.: Modelling and measurements of power losses and turbulence intensity in wind turbine wakes at middelgrunden offshore wind farm, *Wind Energy*, 10, 517–528, <https://doi.org/10.1002/we.238>, 2007.
- Barthelmie, R. J., Hansen, K., Frandsen, S. T., Rathmann, O., Schepers, J. G., Schlez, W., Phillips, J., Rados, K., Zervos, A., Politis, E. S., and Chaviaropoulos, P. K.: Modelling and measuring flow and wind turbine wakes in large wind farms offshore, *Wind Energy*, 12, 431–444, <https://doi.org/10.1002/we.348>, 2009.
- 15 Belu, R. and Koracin, D.: Statistical and spectral analysis of wind characteristics relevant to wind energy assessment using tower measurements in complex terrain, *J. Wind Energy*, 2013, 1–12, <https://doi.org/10.1155/2013/739162>, 2013.
- Beyer, H. G. and Sendhoff, B.: Robust optimization - a comprehensive survey, *Comput. Methods Appl. Mech. Eng.*, 196, 3190–3218, <https://doi.org/10.1016/j.cma.2007.03.003>, 2007.
- 20 Brand, A., Wagenaar, J., Eecen, P., and Holtslag, M.: Database of measurements on the offshore wind farm Egmond aan Zee, *Wind Energy*, 14, 1–6, <ftp://ecm.nl/pub/www/library/report/2012/m12012.pdf>, 2012.
- Briggs, K.: Navigating the complexities of wake losses, *North Am. Wind.*, 10, 2013.
- Caffisch, R. E.: Monte Carlo and quasi-Monte Carlo methods, *Acta Numer.*, pp. 1 – 49, 1998.
- Chowdhury, S., Zhang, J., Messac, A., and Castillo, L.: Optimizing the arrangement and the selection of turbines for wind farms subject to varying wind conditions, *Renew. Energy*, 52, 273–282, <https://doi.org/10.1016/j.renene.2012.10.017>, 2013.
- 25 Doostan, A. and Owhadi, H.: A non-adapted sparse approximation of PDEs with stochastic inputs, *J. Comput. Phys.*, 230, 3015–3034, <https://doi.org/10.1016/j.jcp.2011.01.002>, 2011.
- Eldred, M. S.: Design under uncertainty employing stochastic expansion methods, *Int. J. Uncertain. Quantif.*, 1, 119–146, <https://doi.org/10.1615/IntJUncertaintyQuantification.v1.i2.20>, 2011.
- 30 Eldred, M. S. and Burkardt, J.: Comparison of non-intrusive polynomial chaos and stochastic collocation methods for uncertainty quantification, in: *47th AIAA Aerosp. Sci. Meet. Incl. New Horizons Forum Aerosp. Expo.*, pp. 1–20, AIAA, <https://doi.org/10.2514/6.2009-976>, 2009.
- Eldred, M. S., Webster, C. G., and Constantine, P. G.: Evaluation of non-intrusive approaches for wiener-asky generalized polynomial chaos, in: *49th AIAA/ASME/ASCE/AHS/ASC Struct. Struct. Dyn. Mater. Conf.*, pp. 1–22, AIAA, <https://doi.org/10.2514/6.2008-1892>, 2008.
- 35 Feinberg, J. and Langtangen, H. P.: Chaospy: An open source tool for designing methods of uncertainty quantification, *J. Comput. Sci.*, 11, 46–57, <https://doi.org/10.1016/j.jocs.2015.08.008>, 2015.

- Fleming, P. A., Ning, A., Gebraad, P. M. O., and Dykes, K.: Wind plant system engineering through optimization of layout and yaw control, *Wind Energy*, 19, 329–344, <https://doi.org/10.1002/we.1836>, 2016.
- Gautschi, W.: *Orthogonal polynomials computation and approximation*, Oxford University Press Inc., New York, 2004.
- Gebraad, P., Thomas, J. J., Ning, A., Fleming, P., and Dykes, K.: Maximization of the annual energy production of wind power plants by optimization of layout and yaw-based wake control, *Wind Energy*, 20, 97–107, <https://doi.org/10.1002/we.1993>, 2017.
- Gebraad, P. M. O., Teeuwisse, F. W., van Wingerden, J. W., Fleming, P. A., Ruben, S. D., Marden, J. R., and Pao, L. Y.: Wind plant power optimization through yaw control using a parametric model for wake effects-a CFD simulation study, *Wind Energy*, 19, 95–114, <https://doi.org/10.1002/we.1822>, 2016.
- Giles, M. B. and Pierce, N. A.: An Introduction to the adjoint approach to design, *Flow, Turbul. Combust.*, 65, 393–415, <https://doi.org/10.1023/A:1011430410075>, 2000.
- Gill, P. E., Murray, W., and Saunders, M. A.: SNOPT: an SQP algorithm for large-scale constrained optimization, *SIAM Rev.*, 47, 99–131, <https://doi.org/10.1137/S0036144504446096>, 2005.
- Gray, J., Moore, K., and Naylor, B.: OpenMDAO: an open source framework for multidisciplinary analysis and optimization, in: 13th AIAA/ISSMO Multidiscip. Anal. Optim. Conf., AIAA, <https://doi.org/10.2514/6.2010-9101>, 2010.
- Griewank, A. and Walther, A.: *Evaluating derivatives: principles and techniques of algorithmic differentiation*, SIAM, second edn., <https://doi.org/10.1137/1.9780898717761>, 2008.
- GWEC: Global wind report: annual market update, Tech. rep., Global Wind Energy Council, <http://www.gwec.net/wp-content/uploads/vip/GWEC-Global-Wind-2015-Report{ }April-2016{ }22{ }04.pdf>, 2016.
- Hastie, T., Tibshirani, R., and Friedman, J.: *The elements of statistical learning*, Springer Series in Statistics, Springer New York, New York, NY, <https://doi.org/10.1007/b94608>, 2009.
- Herbert-Acero, J. F., Probst, O., Rethore, P. E., Larsen, G. C., and Castillo-Villar, K. K.: A review of methodological approaches for the design and optimization of wind farms, *Energies*, 7, 6930–7016, <https://doi.org/10.3390/en7116930>, 2014.
- Hosder, S., Walters, R. W., and Balch, M.: Efficient sampling for non-intrusive polynomial chaos applications with multiple uncertain input variables, in: Collect. Tech. Pap. - AIAA/ASME/ASCE/AHS/ASC Struct. Struct. Dyn. Mater. Conf., vol. 3, pp. 2946–2961, AIAA, <https://doi.org/doi:10.2514/6.2007-1939>, 2007.
- Jensen, N. O.: A note on wind generator interaction, Risø-M-2411 Risø Natl. Lab. Roskilde, pp. 1–16, <http://orbit.dtu.dk/files/55857682/ris{ }m{ }2411.pdf>, 1983.
- Jiménez, Á., Crespo, A., and Migoya, E.: Application of a LES technique to characterize the wake deflection of a wind turbine in yaw, *Wind Energy*, 13, 559–572, <https://doi.org/10.1002/we.380>, 2009.
- Jonkman, J., Butterfield, S., Musial, W., and Scott, G.: Definition of a 5-MW Reference Wind Turbine for Offshore System Development, Tech. Rep. TP-500-38060, National Renewable Energy Laboratory (NREL), Golden, CO, <https://doi.org/10.2172/947422>, 2009.
- Kennedy, M. C. and O’Hagan, A.: Bayesian calibration of computer models, *J. R. Stat. Soc. Ser. B (Statistical Methodol.)*, 63, 425–464, <https://doi.org/10.1111/1467-9868.00294>, 2001.
- Kusiak, A. and Song, Z.: Design of wind farm layout for maximum wind energy capture, *Renew. Energy*, 35, 685–694, <https://doi.org/10.1016/j.renene.2009.08.019>, 2010.
- Kwong, W. Y., Zhang, P. Y., Romero, D., Moran, J., Morgenroth, M., and Amon, C.: Wind Farm Layout Optimization Considering Energy Generation and Noise Propagation, in: Vol. 3 38th Des. Autom. Conf. Parts A B, p. 323, ASME, <https://doi.org/10.1115/DETC2012-71478>, 2012.

- Le Maître, O. P. and Knio, O. M.: Spectral methods for uncertainty quantification, Scientific Computation, Springer Netherlands, Dordrecht, <https://doi.org/10.1007/978-90-481-3520-2>, 2010.
- McKay, M. D., Beckman, R. J., and Conover, W. J.: Comparison of three methods for selecting values of input variables in the analysis of output from a computer code, *Technometrics*, 21, 239–245, <https://doi.org/10.1080/00401706.1979.10489755>, 1979.
- 5 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?, in: *J. Phys. Conf. Ser.*, vol. 625, pp. 1–11, <https://doi.org/10.1088/1742-6596/625/1/012030>, 2015.
- Navarro, M., Witteveen, J., and Blom, J.: Polynomial chaos expansion for general multivariate distributions with correlated variables, *arXiv Prepr. arXiv1406.5483*, pp. 1–23, 2014.
- Ng, L. and Eldred, M. S.: Multifidelity uncertainty quantification using non-intrusive polynomial chaos and stochastic collocation, in: 53rd AIAA/ASME/ASCE/AHS/ASC Struct. Struct. Dyn. Mater. Conf. 20th AIAA/ASME/AHS Adapt. Struct. Conf. 14th AIAA, pp. 1–17, AIAA, <https://doi.org/10.2514/6.2012-1852>, 2012.
- 10 Oberkampf, W., Helton, J., and Sentz, K.: Mathematical representation of uncertainty, in: 19th AIAA Appl. Aerodyn. Conf., AIAA, <https://doi.org/10.2514/6.2001-1645>, 2001.
- Oladyshkin, S. and Nowak, W.: Data-driven uncertainty quantification using the arbitrary polynomial chaos expansion, *Reliab. Eng. Syst. Saf.*, 106, 179–190, <https://doi.org/10.1016/j.res.2012.05.002>, 2012.
- 15 Padrón, A. S.: Polynomial chaos and multi-fidelity approximations to efficiently compute the annual energy production in wind farm layout optimization, Phd thesis, Stanford University, <https://purl.stanford.edu/sc655fd3999>, 2017.
- Padrón, A. S., Stanley, A. P. J., Thomas, J. J., Alonso, J. J., and Ning, A.: Polynomial chaos for the computation of annual energy production in wind farm layout optimization, *J. Phys. Conf. Ser.*, 753, <https://doi.org/10.1088/1742-6596/753/3/032021>, 2016.
- 20 Palar, P. S., Tsuchiya, T., and Parks, G. T.: Multi-fidelity non-intrusive polynomial chaos based on regression, *Comput. Methods Appl. Mech. Eng.*, 305, 579–606, <https://doi.org/10.1016/j.cma.2016.03.022>, 2016.
- Perez, R. E., Jansen, P. W., and Martins, J. R. R. A.: pyOpt: a python-based object-oriented framework for nonlinear constrained optimization, *Struct. Multidiscip. Optim.*, 45, 101–118, <https://doi.org/10.1007/s00158-011-0666-3>, 2012.
- Sanderse, B.: Aerodynamics of wind turbine wakes: Literature review, *Energy Res. Cent. Netherlands*, pp. 1–46, 2009.
- 25 Smith, R. C.: Uncertainty quantification theory, implementation, and applications, SIAM, <http://bookstore.siam.org/cs12/>, 2014.
- Thomas, J. J., Gebraad, P. M., and Ning, A.: Improving the FLORIS wind plant model for compatibility with gradient-based optimization, *Wind Eng.*, pp. 1–16, <https://doi.org/10.1177/0309524X17722000>, 2017.
- U.S. Department of Energy: Wind vision: A new era for wind power in the United States, Tech. rep., U.S. Department of Energy, Washington, DC, <http://www.energy.gov/sites/prod/files/WindVision{ }Report{ }final.pdf>, 2015.
- 30 Wan, X. and Karniadakis, G. E.: Multi-element generalized polynomial chaos for arbitrary probability measures, *SIAM J. Sci. Comput.*, 28, 901–928, <https://doi.org/10.1137/050627630>, 2006.