



A numerical framework for optimal trade-offs between land use and LCOE using efficient, blockage-aware multi-fidelity methods

Cory Frontin, Jeffrey Allen, Christopher J. Bay, Jared Thomas, Ethan Young, and Pietro Bortolotti

National Renewable Energy Laboratory, 15013 Denver West Parkway, Golden, CO 80401, USA

Correspondence: Cory Frontin (cory.frontin@nrel.gov)

Abstract. This paper introduces a novel approach to efficiently estimate the annual energy production (AEP) of a wind farm. The numerical predictions are generated thanks to a multi-fidelity model that combines a classical low-fidelity wake engineering solver with a mid-fidelity computational fluid dynamic solver. The novel setup is not only faster than conventional approaches but is also capable of estimating the AEP of tightly spaced wind farms. Using this approach, we explore the trade-off between land use and levelized cost of energy (LCOE) for a wind farm made of 25 turbines. The results of this study, which ignore impact the layout sensitivity of fatigue loads and their incumbent effect on costs, quantify the penalty on LCOE performance that can be paid to restrict the land use of a wind farm. The results also exhibit the novel capabilities of our approach for multi-fidelity wind farm design to avoid false optima according to incomplete representations of the relevant physical phenomena.

10 1 Background and motivation

The most prominent benefits of wind energy as a source of electricity are that it is inexhaustible, inexpensive, and secure. However, wind energy may also cause unwanted negative effects such as visual impact, shadow flicker, acoustic emissions, and mortality of wildlife. These negative effects not only impact communities and wildlife near wind farms but also harm the wind energy industry as a whole because they may delay, disrupt, and diminish the success of wind energy projects across project life cycles.

Veers et al. (2023) highlighted the need for holistic design approaches to incorporate social and environmental metrics into the design of the next generation of wind energy systems. A major challenge to achieving this goal is that social and environmental factors are usually harder to pose as straightforward optimization targets. In this work, we take a first step in this direction by studying the trade-off between the levelized cost of energy (LCOE) and land use.

20 We elect to study land use because it represents a simple and immediate metric that interfaces with both social and environmental factors. On the social side, land use (on a per-turbine or per-capacity basis) dictates the likelihood that a wind farm creates an unwanted negative impact. We work under the assumption that by minimizing the use of land, a wind farm minimizes its impact on neighbors. The assumption is supported by setback ordinances, which are used by government bodies to limit the proximity of wind turbines to structures or critical infrastructure. Setback ordinances significantly affect the deployment potential of wind energy (Lopez et al., 2023).



On the environmental side, the correlation between land use and impact on wildlife is harder to establish. Emerging research has investigated the impact of wind turbine designs and wind farm layouts on wildlife, specifically on raptors (Quon et al., 2022), bats (Hein and Straw, 2021), and sage grouse (LeBeau et al., 2017). In this work we assume that land use is correlated to the likelihood that a farm intersects with migratory pathways or habitats that put wildlife at risk of disruption or injury.

30 In terms of costs, reducing the land use of wind farms has a direct impact on energy losses and balance-of-system (BOS) costs. These two metrics tend to drive layout optimization in different directions: To minimize energy losses, which are caused by turbine wakes, turbines in a wind farm should be spaced as far apart as possible. Conversely, to minimize BOS costs, turbines should be installed closer to each other in a wind farm, as tighter spacing limits costs associated with new roads and cabling.

35 In this work, we optimize wind farm layouts for LCOE, the most common metric for understanding the balance between the production of energy and the lifetime costs associated with energy production. This is not the first work in this direction. Fleming et al. (2016) developed a method for optimizing farm layout and controls with respect to expected farm power density or power given cable length constraints, but stopped short of full cost modeling. Other work considered trade-offs among noise, production, and land use (Yamani Douzi Sorkhabi et al., 2016) but without fully considering the physical impacts of the land use
40 constraint. As the spacing of wind farms becomes tighter, technical challenges in the estimation of farm performance increase. Power estimates from engineering wake models are known to be inaccurate when turbines are in close proximity (Göçmen et al., 2016), though progress has been made in developing wake models with more accurate shape (Keane, 2021) and deficit (Zhang et al., 2023) in the near wake. These advancements can help improve wind farm performance estimation.

Unfortunately, wake effects are the leading loss mechanism in wind farms, but they are not the only ones. Blockage effects
45 must also be considered. The precise definition of blockage varies slightly in the literature, but the term refers to the deflection of airflow around an obstruction due to the pressures exerted by wind turbines and farms on the surrounding flow (Strickland and Stevens, 2022). Blockage effects can occur both upstream and downstream of the obstruction and can result in speedups and deficits in the flow. According to Sanchez Gomez et al. (2023), blockage effects are estimated to reduce the wind speed by 1 % to 5 %, and sometimes as much as 10 %, at a distance of 1.5 to 3 rotor diameters upstream of the first row of wind
50 turbines. While the focus of blockage study in wind energy is typically the negative upstream effects that slow the airflow into the front row of wind turbines, ignoring the downstream speedups and energy redistributing effects of blockage could also bias AEP predictions (Meyer Forsting et al., 2023). Wind farm blockage effects are typically neglected in engineering design tools, though progress is being made in the development of engineering design tools for estimating blockage effects. Nonetheless, all known state-of-the-art engineering-level blockage models neglect some effects of blockage and slowdown (Nygaard et al.,
55 2020).

Blockage effects are caused by the interactions between the velocity and pressure fields, which cause incoming flow to stagnate at the front of a farm, reducing power yield. These effects fundamentally require flow-resolving computational fluid dynamics (CFD) estimates to capture accurate power estimates. While engineering tools seek to decouple the wake and blockage effects, these phenomena are inherently coupled. The combination of wake and blockage effects in a wind farm can be



60 captured more effectively through CFD than through engineering wake models; however, CFD approaches are much more computationally intensive than engineering models.

To balance the benefits and costs of both CFD and engineering model approaches, this paper applies a novel numerical framework to minimize the land usage for a structured five-by-five wind farm while efficiently incorporating engineering and CFD simulations to estimate annual energy production (AEP). In turn, LCOE estimates can be made more efficiently than
65 would be possible with CFD alone and more accurately than is possible with current engineering models. We show Pareto fronts describing the possible trade-offs in terms of land usage, power generation, and costs for a family of structured farms. These studies can inform both the policy decisions made by key stakeholder institutions and the design procedures adopted by wind farm developers, resulting in the maximization of wind energy deployment while limiting the impact on communities and wildlife.

70 2 Numerical framework

2.1 Optimization framework

The toolchain described in this manuscript aims to optimize the layout of wind farms by minimizing LCOE. LCOE is defined as the ratio of lifetime-averaged annual costs (C) and lifetime-averaged AEP and can be expressed in terms of U.S. dollars (USD) per megawatt-hour:

$$75 \text{ LCOE} = \frac{C}{\text{AEP}} = \frac{r_{fc} \times (\text{CapEx} + \text{BOS}) + \text{OpEx}}{\text{AEP}}. \quad (1)$$

In this equation, C is composed of operational expenditures (OpEx), including land lease costs, maintenance costs, and other recurring operational expenses; initial capital costs (CapEx), or the capital expenditures necessary to acquire the wind turbines and related equipment; and the BOS costs, including permits, transportation, roads, foundations, electrical connections, assembly, and installation. The overall capital costs, including the BOS costs, are multiplied by the fixed charge rate (r_{fc}) to
80 determine their annualized cost; r_{fc} represents the annual amount per dollar of the capital costs needed to cover all project financing, including taxes, insurance, depreciation, return on debt and equity, and related fees¹.

In Fig. 1, the information flow for the optimization process is shown in an extended design structure matrix (XDSM) diagram (Lambe and Martins, 2012). Each of the components, indicated by red rectangles, is interfaced and controlled by the OpenMDAO library (Gray et al., 2019). Flow solvers used in this work, the National Renewable Energy Laboratory's (NREL's) framework FLOW Redirection and Induction in Steady State (FLORIS) and NREL's finite-element solver WindSE, are described in Sect. 2.2.1 and 2.2.2, respectively. The multi-fidelity AEP integrator used in this work is developed in Sect. 2.3, and the land use calculator implements straightforward area calculations for a parallelogram-shaped farm. In this work, CapEx and OpEx are modeled as constant inputs. The BOS and LCOE calculators are from the Wind-Plant Integrated System Design and Engineering Model (WISDEM®), a Python framework also built around the OpenMDAO library. The wind farm parameteri-
90 zation used to create the wind farm configuration design variables is described in Sect. 3.1. In this project, only some elements

¹If a project is realized by internal financing, then these monthly contributions can represent the equivalent internal costs of capital.



of WISDEM are used (namely, PlantFinanceSE) to compute LCOE using the terms in Eq. (1), and LandBOSSE (Land-Based Balance-of-System Systems Engineering Model) is used to compute BOS as a function of the farm design. Future extensions of this framework can include the integration of more elements of WISDEM to bring CapEx and OpEx metrics into the optimization process, especially for farm/turbine co-design. This aspect is discussed in Sect. 5.2. Additionally, the LCOE metric
95 can be easily replaced by other value-based metrics, such as the ones described in Mowers and Mai (2021). Similar metrics are already available in PlantFinanceSE for a given set of inputs describing the variations in space and time of electricity price.

From Eq. (1) and Fig. 1, we can understand at a high level how various stages of wind energy system design impact LCOE. The design of the wind turbine itself has a direct impact on CapEx, BOS, OpEx, and AEP. The design of the wind farm, holding the turbine constant, has no impact on CapEx, a direct impact on BOS and AEP, and a nominally weak, indirect impact
100 on OpEx. In this work, we neglect the indirect OpEx and CapEx dependencies due to turbine redesign and consider how the impacts of farm design on BOS and AEP are passed through to LCOE.

2.2 Wind farm solvers

The denominator of Eq. (1) is the mean AEP of the wind farm. In our framework, this is the most computationally intensive term in Eq. (1). It is typical to treat AEP as a probabilistic integral (Herbert-Acero et al., 2014):

$$105 \quad \text{AEP} = \int_0^{2\pi} \int_0^{\infty} \underbrace{T_{\text{year}} f_{\text{power}}(\psi, V) p(\psi, V)}_{\mathcal{I}_{\text{AEP}}(\psi, V)} dV d\psi \quad (2)$$

where we use $T_{\text{year}} = 8760$ hours. Here, f_{power} is the farm power response at wind direction ψ and velocity V , and $p(\psi, V)$ is the wind rose joint probability density function (pdf), i.e., the likelihood that the wind is blowing from direction ψ with speed V . In this work, we compute f_{power} to evaluate the integrand, $\mathcal{I}_{\text{AEP}}(\psi, V)$, using three different approaches:

- Summing the power curve-indicated powers for the turbines in the farm at the hub-height freestream wind speed, ignoring
110 all inter-turbine aerodynamic effects
- An engineering (analytical) wake model implemented in FLORIS, described in Sect. 2.2.1
- A Reynolds-averaged Navier-Stokes/Actuator Disk (RANS-AD) method implemented in WindSE, described in Sect. 2.2.2.

These methods of calculating f_{power} , each at a drastically different level of fidelity, can be used individually or combined in estimating AEP using various strategies for numerical integration.

115 We consider three numerical integration strategies for evaluating the integral in Eq. (2) in this paper:

- Rectangle rule integration, the standard approach, which discretizes the domain of Eq. (2) to approximate the integral as a sum over many samples of the integrand. The rectangle rule approach and other simple numerical integration techniques are typical for AEP computation with engineering design tools (Herbert-Acero et al., 2014), but are generally not tractable with higher-fidelity tools.



- 120 – The Gaussian process-based approach of King et al. (2020), which approximates the integral using sets of dozens-to-hundreds of evaluations of the power at so-called Bayesian quadrature points in the space of wind conditions that are optimized with respect to p and assumptions about the smoothness of f_{power} .
- A novel multi-fidelity Gaussian process-based approach, detailed in Sect. 2.3, that extends the approach of King et al. (2020) to fuse thousands of FLORIS (NREL, 2024) evaluations with small numbers of WindSE evaluations to achieve
125 accurate AEP estimates with significantly lower computational costs than would be possible using WindSE alone.

2.2.1 FLORIS

NREL's wind farm solver FLORIS is an engineering design tool for modeling and optimizing wind farm performance. It includes implementations of several different engineering wake models. FLORIS is particularly useful in optimization scenarios because of its low computational cost. FLORIS approximates mean wake effects to simulate the impact of turbine
130 wakes on downstream turbines, although it does not capture blockage effects. We used the following models from the FLORIS v4.1 (NREL, 2024) implementation:

- The Gaussian velocity model, based on the works of Bastankhah and Porté-Agel (2016) and Niayifar and Porté-Agel (2016), which represents the velocity deficit as a Gaussian distribution in the spanwise direction, with parameters adjustable by the user; this model distinguishes between near and far wake zones, both maintaining a Gaussian profile but
135 differing in their wake recovery models
- The Gauss deflection model, which blends the approaches of Bastankhah and Porté-Agel (2016) and King et al. (2021) to calculate the wake deflection
- The Crespo-Hernández turbulence model (Crespo and Hernández, 1996)
- The freestream square root of the sum of the squares wake combination model (Katic et al., 1987).

140 Turbulence intensities are set using FLORIS, and wind shear settings were adjusted to match the power-law wind profile used in WindSE simulations.

2.2.2 WindSE

For the higher-fidelity simulations we used the steady-state RANS solver from the open-source wind farm optimization tool WindSE. WindSE is a Python package built on the finite-element package FEniCS (Logg et al., 2012). WindSE uses a mixing-
145 length model for eddy viscosity for stabilization and models turbines using actuator disks (King et al., 2017). WindSE resolves both the velocity and pressure fields so that it can capture realistic blockage effects, which can help improve the power and AEP estimation. WindSE also uses the package dolfin-adjoint to provide automatic differentiation using the adjoint method to pass gradient information through the solution of the nonlinear partial differential equations (Mitusch et al., 2019). Although not used in this work, this allows WindSE to return the gradient of the design variables (e.g., turbine location) with respect to



150 the objective functions (e.g., power production). WindSE is significantly more computationally expensive than FLORIS, with the cases in this work requiring around 20 minutes while the FLORIS runs require less than a second. The trade-off between fidelity and computation time dictated much of our optimization approach; we ran WindSE only where it could significantly improve power predictions compared to FLORIS and ran FLORIS to quickly fill in the bulk of the power response across the space of wind conditions.

155 The computational expense and the desire to keep the multi-fidelity optimization framework as computationally tractable as possible guided much of the domain sizing for the limited number of WindSE runs. Specifically, we targeted a single WindSE simulation wall-clock time under 20 minutes to ensure that multiple function evaluations, $\mathcal{O}(10)$, could be performed within an optimization loop while maintaining an overall simulation time under 24 hours.

We identified that an unstructured base mesh with a length scale equivalent to a $28 \times 28 \times 18$ box mesh in the x -, y -, and
160 z -directions, respectively, provided a solution that gave an appropriate level of accuracy and provided for the stability of the RANS solver while balancing computational limitations. Starting from this base mesh, we shifted mesh nodes vertically, such that 65 % of the nodes reside below $h_{\text{hub}} + D_{\text{rotor}}$ (where h_{hub} is hub height and D_{rotor} is rotor diameter) to adequately refine the near-ground region. We then defined four levels of nested localized refinements surrounding each turbine: three cylindrical refinements with axes aligned with the streamwise direction extending 14.0, 6.0, and 3.75 D_{rotor} downstream with radii of 2.0,
165 2.0, and 1.25 D_{rotor} , respectively; and a final cylindrical refinement with axes aligned with the z -direction extending to the tip height with a radius of 1.25 D_{rotor} . We acknowledge that a rigorous mesh convergence study should be conducted in the future to quantify the accuracy of the present setup.

For the boundary conditions, WindSE prescribes a power-law wind profile with the power coefficient tuned to achieve an equilibrium boundary layer with the actuator disks inactive. This solution is applied as a Dirichlet condition on upstream non-
170 horizontal boundaries with scaling to achieve a prescribed hub-height velocity. A no-stress boundary condition was applied on downstream non-horizontal boundaries. The ground was given a no-slip boundary condition, and the remaining boundaries were set to free-slip conditions. We neglected more complex atmospheric conditions to align with the goal of creating an enhanced engineering design tool rather than a high-fidelity CFD solver.

2.3 Multi-fidelity Gaussian process integration for AEP

175 In this section, we propose a hierarchical multi-fidelity optimization approach for computationally tractable AEP estimation using sparse RANS-AD estimates of farm power. We take an approach that is similar to that of King et al. (2020), in which the power response from Eq. (2), $f_{\text{power}}(\psi, V)$, is approximated by a Gaussian process, $f_{\text{power}}^{\text{GP}}(\psi, V)$. This strategy takes a set of power response evaluations $\{f_{\text{power}}^{(k)}\}_{k=1}^K$ that are sampled at a set of points in the space of wind conditions, $\{(\psi^{(k)}, V^{(k)})\}_{k=1}^K$, and the weak assumption about the smoothness of $f_{\text{power}}(\psi, V)$ inherent in Gaussian process regression. From these, the
180 approach yields an uncertainty-imbued estimate of $\mathcal{I}_{\text{AEP}}(\psi, V)$, which in turn can be used to generate estimates of the mean and variance of AEP, α_{AEP} and β_{AEP}^2 , respectively, such that

$$\left(\text{AEP} \mid \{f_{\text{power}}^{(k)}\}_{k=1}^K\right) \sim \mathcal{N}(\alpha_{\text{AEP}}, \beta_{\text{AEP}}^2). \quad (3)$$



Moreover, this approach allows for analytical expressions for the derivatives $\partial\alpha_{AEP}/\partial f_{\text{power}}^{(k)}$, which can be beneficial for efficient optimization. Finally, King et al. (2020) show that, given the joint pdf that characterizes a wind resource, $p(\psi, V)$, and independent of f_{power} , there is a set of K wind conditions, $\{(\psi^{(k)}, V^{(k)})\}_{k=1}^K$, that minimize β_{AEP}^2 and thus give the best possible AEP estimate for a given K . Because these points are independent of the values of f_{power} , the resulting method is known as Bayesian quadrature because it allows for the recycling of the length- K set of points and a set of associated quadrature-like “weights” as long as the wind resource pdf p is unchanged.

This method offers a more efficient and feature-rich means for estimating Eq. (2), but it has some critical flaws. A key issue is that the estimates $f_{\text{power}}^{\text{GP}}(\psi, V)$ tend to zero away from the region in which f_{power} is sampled, contributing to error in the AEP estimates originating from low-probability wind conditions. The insight of this work is that when we use CFD estimates of f_{power} , we also have the ability to sample using engineering design tool estimates of f_{power} at virtually zero additional cost. Kennedy (2000) developed an auto-regressive procedure for hierarchical multi-fidelity Gaussian process regression, such that a multi-fidelity surrogate model of the form $f_{\text{multi}}(\psi, V) = \rho f_{\text{lofi}}(\psi, V) + \Delta f_{\text{corr}}(\psi, V)$ can be created, where f_{lofi} is a Gaussian process fit to low-fidelity data. These couple with ρ , a regression parameter, and Δf_{corr} , a corrective Gaussian process, which can be fit simultaneously based on a small number of high-fidelity data that are co-located with low-fidelity data.

This multi-fidelity technique has been previously demonstrated to be effective for outer-loop multi-fidelity optimization in the wind industry, including the work of Jasa et al. (2022). In this work, we implement the multi-fidelity Gaussian process approach following the efficient recursive formulation of Le Gratiet (2013) and integrate it into the AEP integration methodology of King et al. (2020) as a multi-fidelity inner-loop model. Complete details of the implementation can be found in Appendix A. The result is that we can write an estimate of the AEP based on a fusion of low- and high-fidelity results:

$$\left(\text{AEP} \left| \{f_{\text{hifi}}^{(k)}\}_{k=1}^{K_{\text{hifi}}}, \{f_{\text{lofi}}^{(k)}\}_{k=1}^{K_{\text{lofi}}}\right.\right) \sim \mathcal{N}(\alpha_{\text{MF}}, \beta_{\text{MF}}^2). \quad (4)$$

This approach allows the low-fidelity solver to cover regions in the wind condition joint pdf with lower likelihood with good-enough power estimates while delivering the accuracy of the high-fidelity solver in regions with high likelihood of wind (and therefore the highest contribution to the AEP).

Implementation

We have re-implemented the single-fidelity estimates used in King et al. (2020). For consistent fitting in the small-sample limit, we fix the length scale parameters of the Gaussian process regression at 22.5° in the wind direction axis and 2.5 m s^{-1} in the wind speed axis; these fixed length scales are used in the single-fidelity and later in the multi-fidelity fits. The Bayesian quadrature optimization in King et al. (2020) is also implemented using the COBYLA optimization solver. The use of Bayesian quadrature points in this work refers to the set of single-fidelity optimal points using this optimization.

The multi-fidelity methods of Le Gratiet (2013) are also re-implemented and extended, where the lowest-fidelity method is set using our single-fidelity implementation. Our implementation then allows the evaluation of the AEP integral, thus extending the single-fidelity methodology of King et al. (2020) to multi-fidelity Gaussian processes and allowing for the computation of derivatives and uncertainties with respect to the AEP integral. Finally, this work incorporates the multi-fidelity AEP integration



capability into an OpenMDAO framework for model coordination and derivative propagation. The authors note that the results in this manuscript use derivative-free optimization methods exclusively, and the derivative propagation capabilities described are not featured in the results in this manuscript.

3 Demonstration of multi-fidelity AEP estimation

220 3.1 Case setup

We now apply the framework to the minimization of land usage for a five-by-five structured wind farm composed of 3.4 MW wind turbines. The model of the turbine was defined for the International Energy Agency Wind Technology Collaboration Programme Task 37 on Systems Engineering and has a rotor diameter of 130 m, a hub height of 110 m, and an aerodynamic power of 3.62 MW (Bortolotti et al., 2019). In this work, we demonstrate the optimization framework using wind resource data for a site in Oklahoma, USA, used within the American Wake Experiment (AWAKEN; see Moriarty et al. (2023)). Data were provided by Engie as part of the AWAKEN project and accessed through the U.S. Department of Energy Atmosphere to Electrons wind data hub (U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, 2023). In Fig. 2, the joint pdf of wind direction and speed for the wind resource is shown along with the corresponding wind rose for the site.

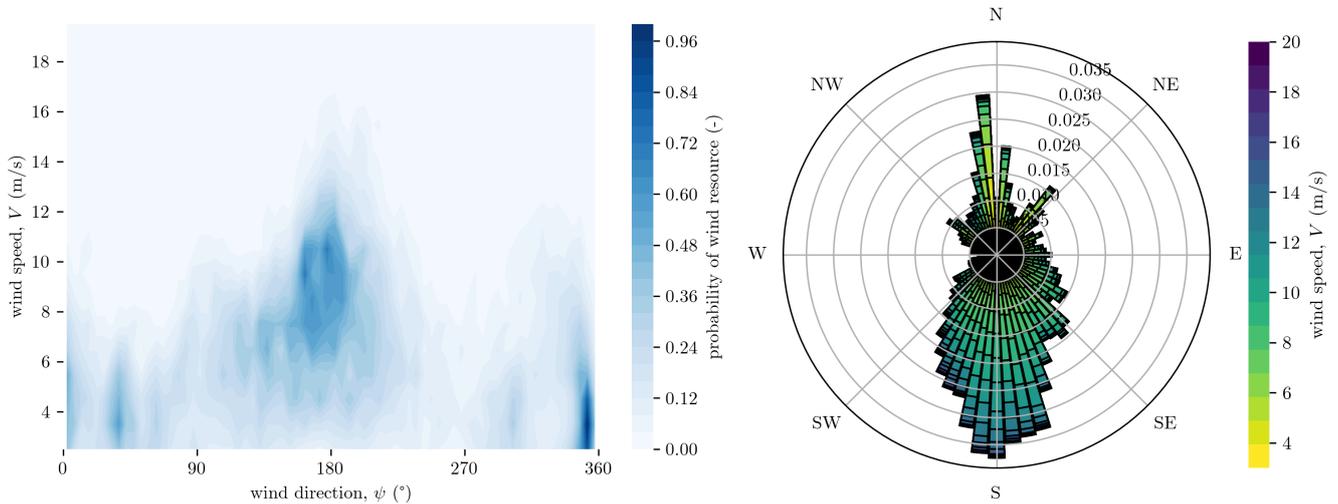


Figure 2. Wind resource joint pdf (left) and equivalent wind rose (right) for the AWAKEN site.

In this work, we consider row- and column-structured wind farms that can be characterized by four variables: the orientation angle, θ ; the equidistant spacing of turbine rows within the farm, L_1 ; the equidistant spacing of turbine columns within the farm perpendicular to the orientation axis, L_2 ; and the angle of clockwise skew of the farm grid, ϕ . These variables are grounded in a static global coordinate system defined by x and y . The reference frame is rotated counterclockwise by θ from the global reference frame and defined by v_1 and v_1^\perp . The final local reference frame is skewed from the rotated frame by rotating v_2



from v_1^\perp counterclockwise about the origin by the angle ϕ so that the local reference frame of the rotated and skewed farm is given by v_1 and v_2 . A graphical representation of the layout variables can be seen in Fig. 3.

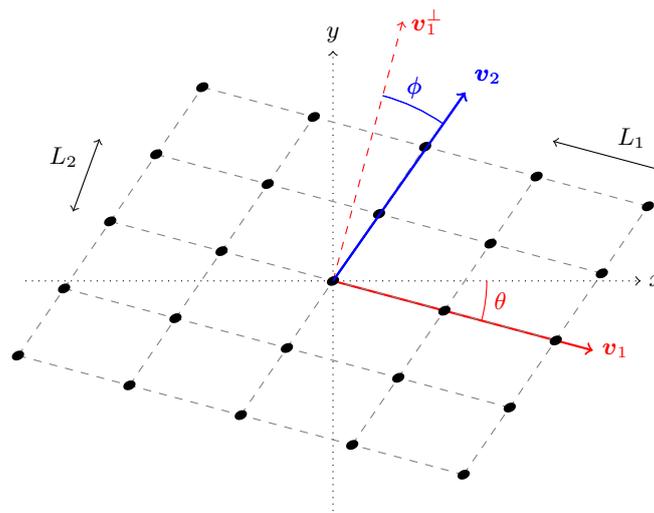


Figure 3. Four-variable parameterization of grid farm layout.

235

3.2 Baseline power response and AEP estimation across fidelities

We begin with a study to demonstrate the difference between the power response surfaces that result from FLORIS and WindSE. For this effort, we use a rectangular farm, for which $L_1 = 3.5D_{\text{rotor}}$, $L_2 = 7.0D_{\text{rotor}}$, and $\theta = \phi = 0$. This farm, with close row (streamwise) spacing with wind from 0° and 180° and close column (perpendicular to the wind direction) spacing with wind from 90° and 270° should demonstrate blockage when the column spacing is tight while showing high wake losses when the row spacing is tight. We expect that both WindSE and FLORIS capture the wake losses with accuracy for tight row spacing, but only WindSE should capture the blockage effects. Figure 4 shows a comparison of the FLORIS and WindSE power responses for this case.

A preliminary note on these results is that tuning the WindSE power curve model for datasheet accuracy has not been a primary focus of this work; we take the results of WindSE at face value in this manuscript while acknowledging a likely overprediction of the Region II power for this tuning. A particular form of tuning mismatches can be seen where the power curve implementation in WindSE has an exaggerated thrust peak-shaving effect compared to FLORIS; this effect can be seen near the rated velocity, where the transition from Region II to Region III is more significantly smoothed in the WindSE result. After tuning mismatches, which are ultimately incidental and reducible, we move on to disagreements caused by physical modeling. To begin, vertical streaks in the difference plotted in Fig. 4, the most pronounced at 0° and 180° , exemplify regions where diffusion and meandering cause “smoothing” of the directional content of the power response surface. This occurs because WindSE resolves the pressure fields, which allow the wakes to interact and mix according to physical principles and

250

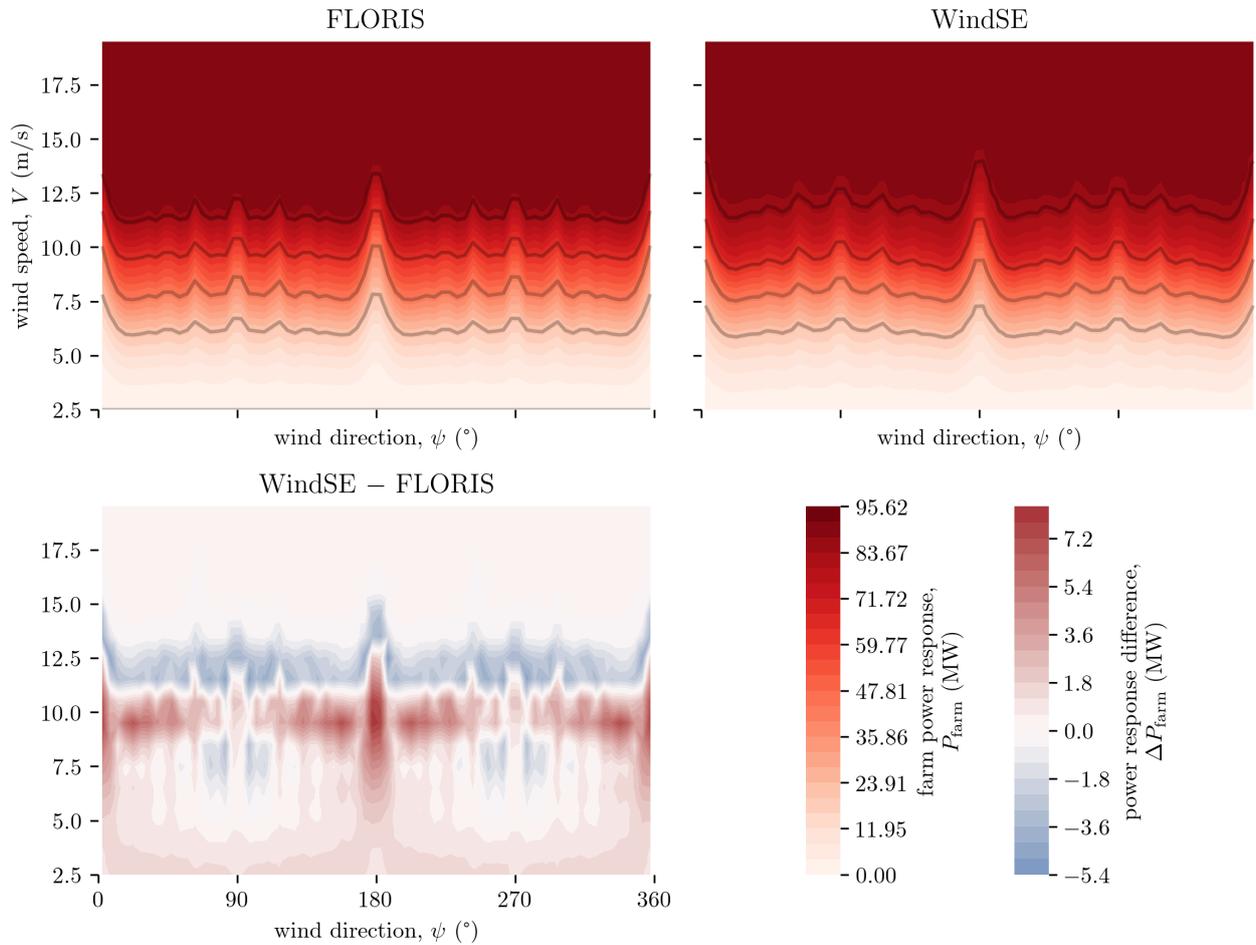


Figure 4. Resolved FLORIS (low-fidelity) and WindSE (high-fidelity) power response surfaces for reference farm; 25 % increments of rated power shown as gray contours.

incorporate blockage and stagnation effects on both the turbine and farm level. We can observe two major differences in the data; the first is the modeling mismatch between WindSE and FLORIS, which amounts to a calibration error in the FLORIS layer. This difference is effectively a reducible modeling error. On the other hand, we observe a second major difference, from the blockage effects, which can be seen in the region II operations, where WindSE reports a direction-dependent depression in power compared to FLORIS between 45° and 135° (and mirrored again between 225° and 315°) in direction and between 7.5 and 12.5 m s^{-1} in speed. While FLORIS could be tuned to match the exaggerated thrust peak shaving or smoothed to mimic diffusion, accurate capture of the blockage effects is fundamentally elusive with engineering design tools, as mentioned in Sect. 1. This study clearly demonstrates the need for some ability to capture effects that engineering tools like FLORIS do not capture. This allows us to avoid optimizations that converge toward regions in the design space that exploit inefficiencies



in the engineering model rather than real physical minimizers, which is especially important when investigating effects in the small-spacing limit.

Finally, we can use these power response surfaces to compute reference rectangle rule AEP estimates for the wind joint pdf from Fig. 2. These computations require 1,296 evaluations, with samples from FLORIS or WindSE made at every 5° and 1 m s^{-1} . The resulting FLORIS AEP estimate is 260.74 GWh; this compares to the WindSE AEP estimate of 272.36 GWh. The causality of this error is, of course, due to a mix of blockage, unmodeled physics, and tuning distinctions, which are hard if not impossible to disentangle. In practice, we will want to capture the effects present at the highest fidelity, and we will reduce the modeling discrepancies as much as possible in the lower-fidelity estimates². Thus, if we take the computationally intensive reference WindSE result to be a truth value, it implies an error of 11.62 GWh in the FLORIS AEP result, about 4.3 % error.

3.3 Multi-fidelity demonstration: power curve + FLORIS + WindSE

Now, we seek to demonstrate the effectiveness of multi-fidelity Gaussian process-based power and AEP approximation. The goal of using the multi-fidelity method is to correct the low-fidelity FLORIS data, which can not resolve all of the effects of the physics-resolving high-fidelity method but can give a reasonable prior for the true behavior. To demonstrate its performance, we reuse the same layout setup as in Sect. 3.2, with $L_1 = 7D_{\text{rotor}}$, $L_2 = 3.5D_{\text{rotor}}$, and $\theta = \phi = 0^\circ$. We take samples of farm power from WindSE and FLORIS at 64 co-located points in the space of wind resources, with the points chosen to be the Bayesian quadrature points that minimize a single-fidelity Gaussian process-based AEP estimate. These are supplemented by 960 additional stand-alone evaluations of FLORIS. The FLORIS-only sample points are selected using a uniform distribution in the wind direction dimension and a Weibull distribution in the wind speed dimension. Finally, we also include 2,048 samples of the trivial power curve method. Half of the power curve samples are co-located with the FLORIS evaluations (64 of which, in turn, are co-located with both WindSE and FLORIS evaluations), and the remainder are uniformly sampled in both the wind direction and wind speed dimensions. The sampling locations for WindSE and FLORIS can be seen in Fig. 5 (samples at the trivial fidelity are taken uniformly in the range of direction and speed and omitted for brevity).

The multi-fidelity Gaussian process fuses these three layers of observations to approximate the power response and AEP integrand. The multi-fidelity process fuses Gaussian process-based response surfaces generated at each level, the mean surfaces of which can be found in Fig. 6, with the corresponding standard deviation in Fig. 7. The standard deviation surfaces demonstrate the regions of relative trust for each Gaussian process surface, which are leveraged by the multi-fidelity fitting process.

This procedure sequentially incorporates higher-fidelity data, and the resulting sequential power responses are depicted in Fig. 8 alongside the corrections that link a given level to the next lower level of fidelity. At the lowest level, the multi-fidelity response and the trivial, power curve Gaussian process response are identical and given by the trivial response in Fig. 6.

In Fig. 8, we can see the correction that is generated by the multi-fidelity process to correct the power curve response to account for the FLORIS results and the resulting power response in the left-hand column. The correction in the low-fidelity

²In this work, we accept a level of otherwise reducible modeling errors in the interest of proving out the multi-fidelity approach and its ability to correct them.

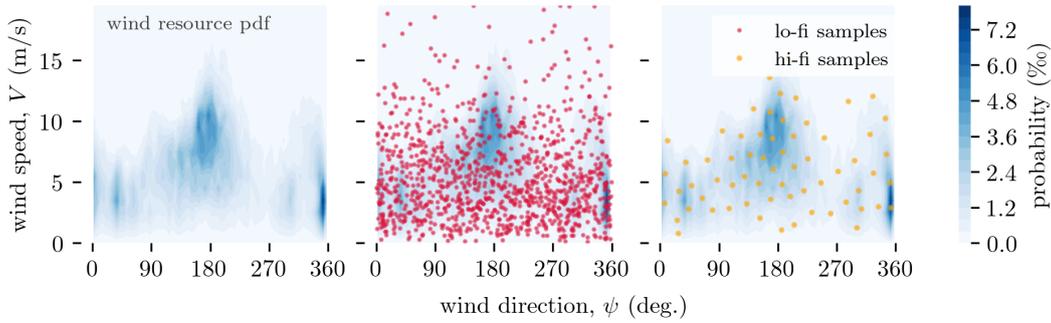


Figure 5. Sampling locations for multi-fidelity AEP estimation with respect to wind resource joint pdf. Farm power sampled from FLORIS for low-fidelity and WindSE for high-fidelity.

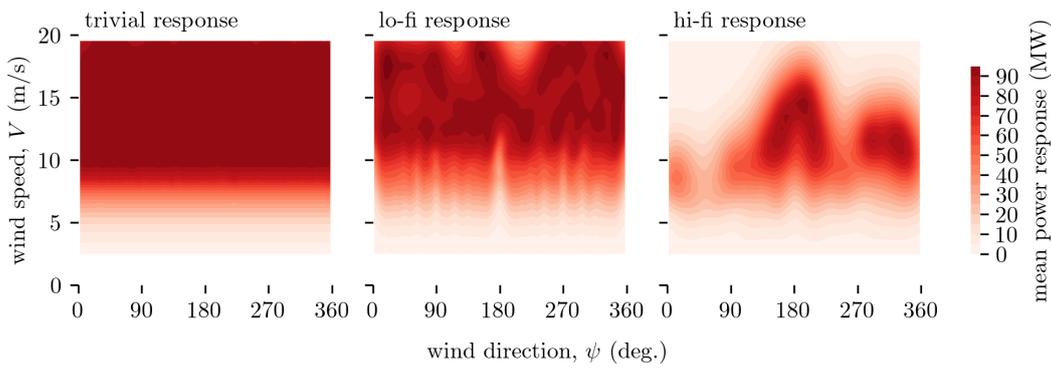


Figure 6. Mean single-fidelity Gaussian process response surfaces for each level of the multi-fidelity AEP integrand.

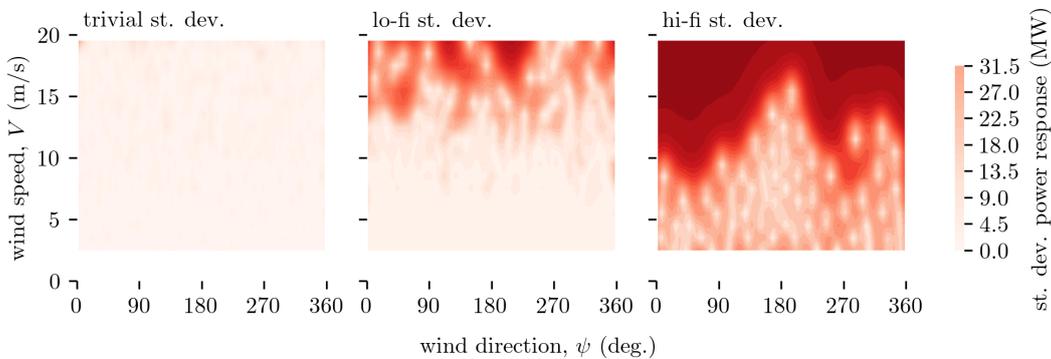


Figure 7. Standard deviation of single-fidelity Gaussian process-based response for each level of the multi-fidelity AEP integrand.

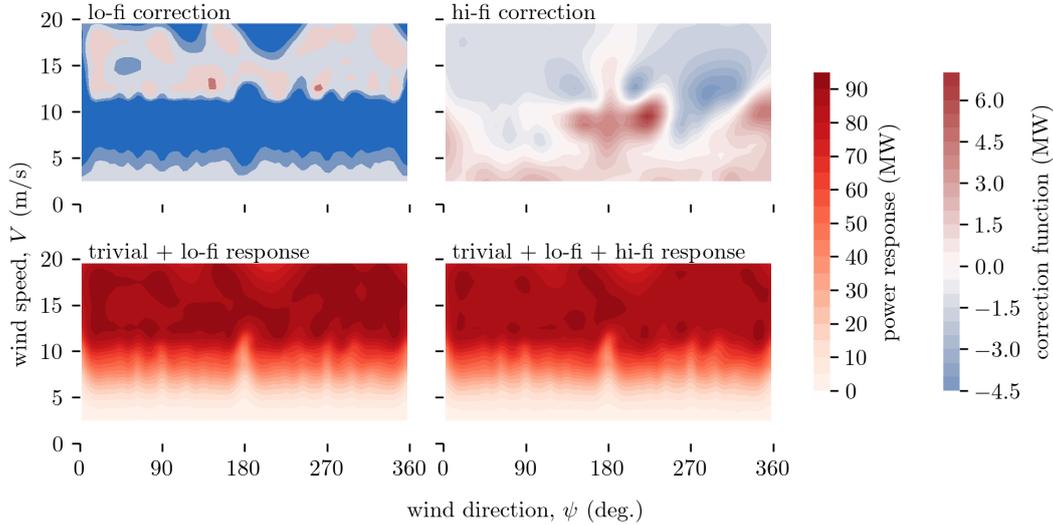


Figure 8. Multi-fidelity corrections and resulting responses after incorporating low-fidelity and high-fidelity data to create a multi-fidelity power response. Correction color bar scale truncated for low-fidelity result for clarity.

case bridges the gap – namely, wake interactions – between the trivial power curve model and low-fidelity FLORIS. The benefit
 295 conferred by the multi-fidelity method at this stage is relatively small and confined to the upper velocities of the rated power
 shelf in the low-fidelity response in Fig. 6, which is sparsely sampled by FLORIS.

The right-hand column of Fig. 8, however, demonstrates how the multi-fidelity process captures the power depression due to
 blockage (in blue, around 90° and 270° wind direction) using a small number of samples. The result of using the multi-fidelity
 Gaussian process-based power estimate is an AEP of 270.25 GWh. In comparison to the reference WindSE result, the error is
 300 2.1 GWh, about 0.8 % error using only 64 samples of the high-fidelity solver.

4 Wind farm layout optimization

At this point, it has been demonstrated that we can approximate blockage-aware AEP using the multi-fidelity solver to capture
 key features from the high-fidelity solver at significantly reduced cost. Now, we can attempt to perform an unconstrained
 optimization using the AEP approximated by the multi-fidelity method. We now pose a minimization problem of the form

$$305 \quad \underset{\theta, L_1, \phi, L_2}{\text{minimize}} \text{ LCOE}(\theta, L_1, \phi, L_2). \quad (5)$$

In addition to the unconstrained flavor, we can also pose a constrained version of this problem:

$$\begin{aligned}
 &\underset{\theta, L_1, \phi, L_2}{\text{minimize}} \text{ LCOE}(\theta, L_1, \phi, L_2) \\
 &\text{subject to } A = A_{\text{limit}},
 \end{aligned} \quad (6)$$

These minimizations are implemented using OpenMDAO as described above, and a 20-instance multi-start approach is used to approximate a global optimizer using the COBYLA algorithm. Pareto optimality fronts are computed using the epsilon-constraint method (Martins and Ning, 2021).

4.1 FLORIS-only optimization

We first use FLORIS in isolation (i.e., without the multi-fidelity AEP estimation) to generate a Pareto front of LCOE with respect to land use. In Fig. 9, we can see the resulting Pareto front that is generated by optimizing the FLORIS LCOE estimates at fixed values of A_{limit} . In addition to the Pareto front, we have also taken the Pareto set generated using FLORIS and run it

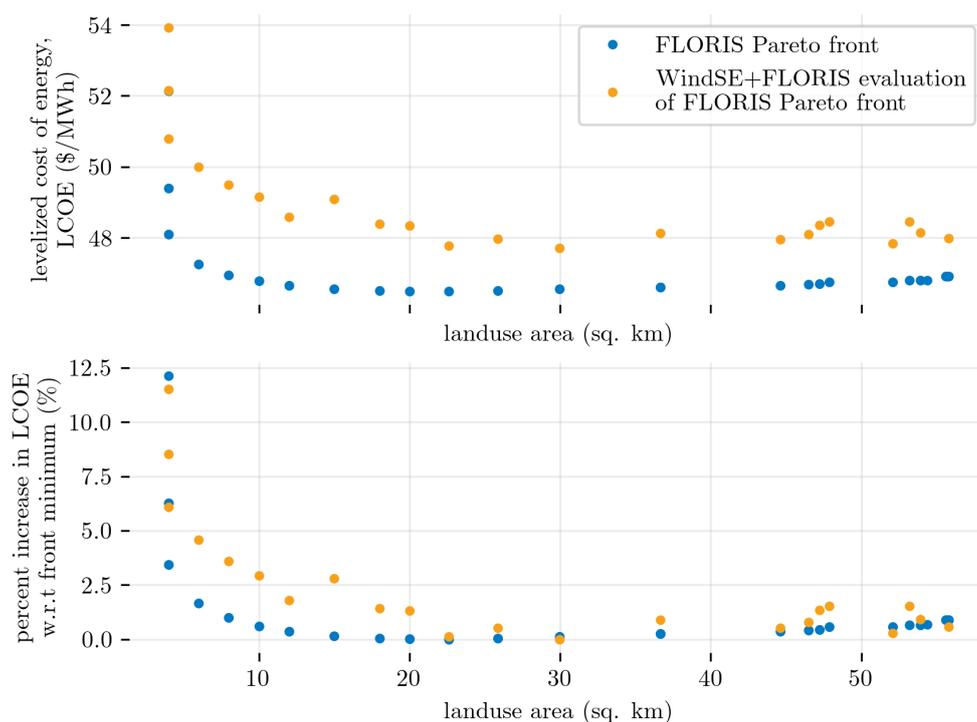


Figure 9. FLORIS-only land use vs. LCOE Pareto front. Pareto-optimal designs are found using FLORIS, and these designs are additionally simulated using the multi-fidelity approach.

with the full multi-fidelity AEP estimation capability to assess the differences between the results computed with and without WindSE evaluations incorporated.

The result in Fig. 9 demonstrates the necessity of accounting for effects that are not captured by low-fidelity methods. The multi-fidelity result, which was shown in Sect. 3.3 to correspond closely to the AEP that is predicted with WindSE in isolation, results in a quantitatively and, perhaps more importantly, qualitatively different front. WindSE predicts a higher



320 LCOE across the entire Pareto set but demonstrates significantly higher-percentage increases in the LCOE as land use is increasingly constrained, where blockage is an important aspect of the farm aerodynamics.

Even in the limit of large land use, LCOE tends to be higher with the multi-fidelity method than FLORIS predicts, which can likely be explained by the limited extent of FLORIS's wakes and the ensuing insensitivity to wake interactions for the FLORIS optimization process. These LCOE sub-optimality are made even more notable given the fact that WindSE – and
325 therefore the multi-fidelity approach – tends to overpredict the AEP on the farm in Sect. 3.2 and therefore tends to underpredict LCOE. Given this result, we now endeavor to design LCOE-optimal farms using the multi-fidelity AEP estimation capabilities exhibited in Sect. 3.

4.2 Multi-fidelity Pareto front with land use and LCOE

Now, we can find the Pareto set using the multi-fidelity AEP estimation capability to solve the optimization problem in Eq. (5) and Eq. (6). Figure 10 shows the resulting Pareto front that emerges from the optimizer termination points. Immediately, it

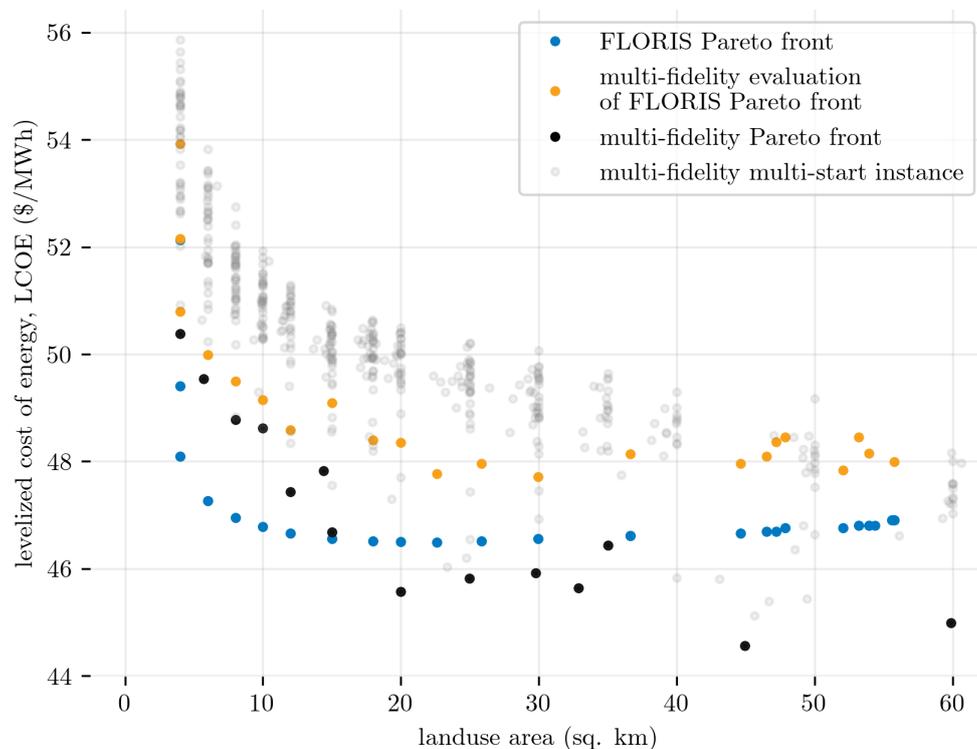


Figure 10. Epsilon-method estimate of the land use/LCOE Pareto front. All multi-fidelity multi-start instances shown with transparency, black indicates land-use-constrained minimizers generated with the epsilon method.

330



becomes clear that the multi-fidelity optimization discovers a new and distinct Pareto set compared to the FLORIS Pareto set, whether that set is analyzed with FLORIS or reanalyzed using the multi-fidelity approach.

By comparison to the pure FLORIS result, we see that the multi-fidelity minimum LCOE is about 4 % lower than the FLORIS minimum LCOE, which is consistent with the overprediction of AEP observed on the reference farm in Sect. 3.2. The multi-fidelity Pareto front suggests that wake and aerodynamic effects have an impact that extends beyond around 20–30 km² of land use, albeit a small one, which is not captured by FLORIS since the wake deficit models have finite lengths. Now, we can compare the LCOE produced by reanalyzing the FLORIS Pareto set using the multi-fidelity toolchain with the LCOE Pareto front generated entirely with the multi-fidelity toolchain. From this comparison, we can see the importance of using the multi-fidelity toolset. In the small land use limit, multi-fidelity designs outperform the FLORIS-optimized farms by a modest gain around \$0.25 MWh⁻¹, which we can expect to be driven by blockage. Somewhat unexpectedly, a more pronounced excess of \$2 MWh⁻¹ in LCOE emerges between the FLORIS-optimized designs under multi-fidelity reanalysis and the multi-fidelity-optimized designs. Taken together, these results strongly underscore the importance of incorporating higher-fidelity analyses in optimization.

Before concluding, we remark on the quality of the epsilon method optimization for the notably expensive multi-fidelity toolchain. Figure 11 shows the optimization trends of the constrained problems. The optimization traces in Fig. 11 reveal that the multi-start optimization is quite necessary to overcome strong local optima in the orientation angle. Moreover, despite the multi-start optimization approach, it is not necessarily guaranteed that we are approaching the globally optimal farm design at any given amount of land use, even in this highly simplified design space. Nonetheless, from the emerging Pareto set, we can understand the trade-off between land use and LCOE, which significantly penalizes LCOE as land use shrinks – for this farm of twenty-five 3.4 MW turbines – below 20 km². Above 20 km², constrained land use continues to induce a moderate penalty compared to the lowest observed LCOE in the set of experiments here. Importantly, across most of the land use values evaluated here and for this farm in particular, the penalty LCOE to shrink a farm from the unconstrained optimal LCOE to the best farm we have at 4 km² amounts to around \$5 MWh⁻¹. This penalty for constrained land use is steep compared to that predicted by FLORIS and must be handled appropriately, but it is nonetheless possible that trades of up to \$5 MWh⁻¹ to reduce the land use constraints of a farm might be useful to developers.

5 Conclusions

5.1 Summary

In this work we proposed a numerical framework for blockage-aware layout optimization of wind farms using a multi-fidelity approach. We demonstrated the framework by applying it to minimize LCOE subject to land use constraints. The goal of the framework is to provide a tool that can be used to optimize wind farms while accounting for blockage and other effects that may not be fully captured in low-fidelity tools but become increasingly important when wind farms are subject to land use restrictions through environmental and social constraints. The framework efficiently combines CFD-based and engineering model

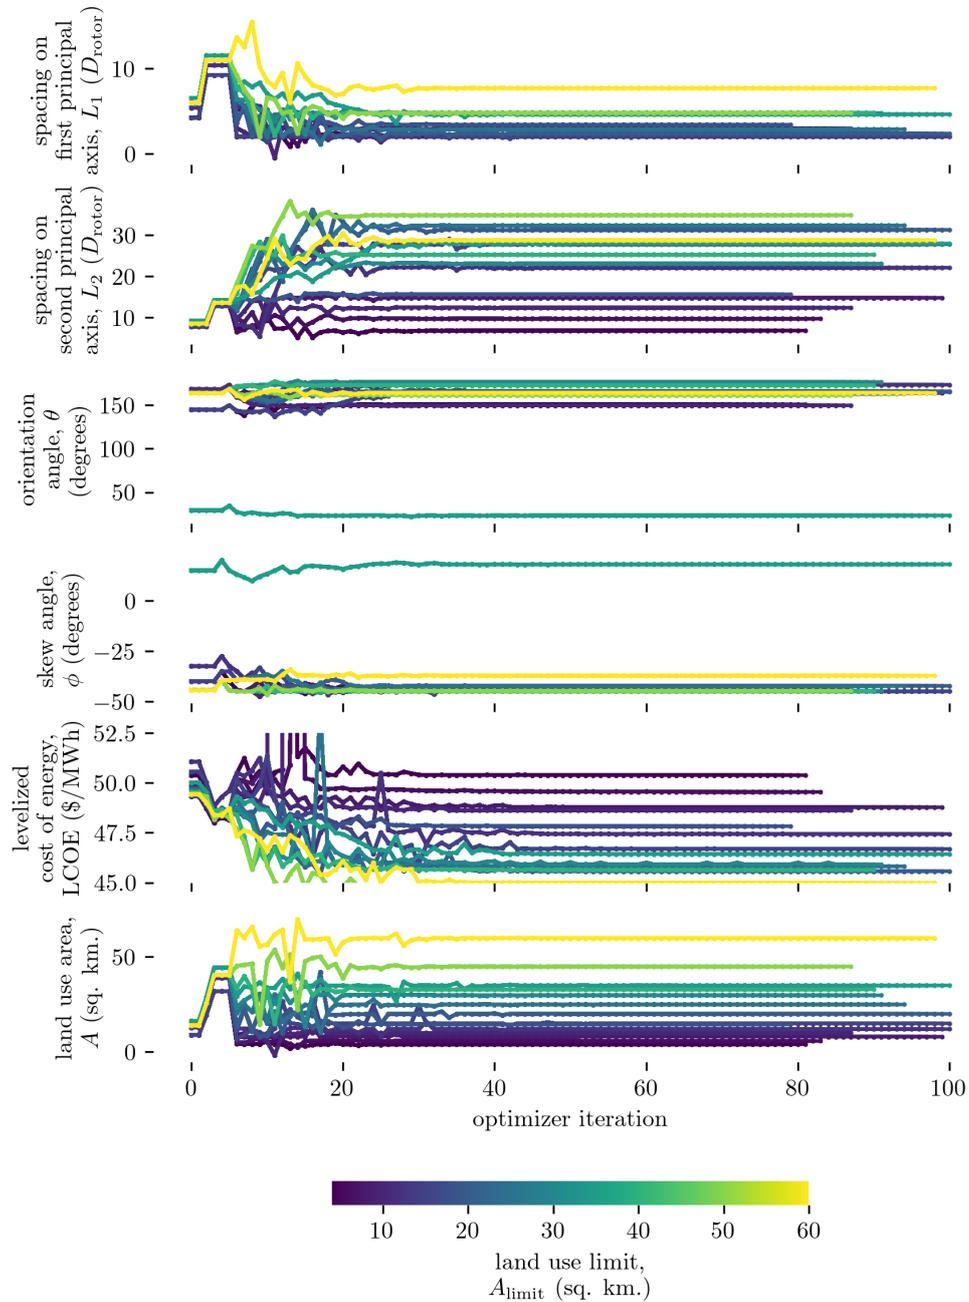


Figure 11. Convergence trends for constrained layout optimization problems. Traces colored by land use area constraint values of 4, 6, 8, 10, 12, 15, 18, 20, 25, and 30 square kilometers. Only the LCOE-minimizing multi-start instance is shown for a given constraint value.

estimates of the AEP. We found that our method can achieve around a 2 order of magnitude reduction in the computational effort necessary to realize an AEP estimate with under 1 % error as compared to using a RANS model alone.



365 To help us better understand the trade-offs between the cost of energy and the environmental and social constraints placed
on new wind farm developments, we used the new framework to minimize the LCOE of a five-by-five wind farm made of 3.4
MW wind turbines with varying land use constraints. The results provide a Pareto front that gives the best-case trade between
land use and LCOE for this site, which indicates that it is possible to allow for a moderate increase in LCOE to significantly
reduce the land use of a wind farm. Specifically, for the cost of about a 2 % increase in LCOE, we can shift from a 20 km²
370 farm to a 12 km² farm. Further area reductions are available for greater increases in LCOE, but the area reductions become
progressively more expensive.

5.2 Future Work

This work demonstrates a proof of concept for a multi-fidelity, multi-disciplinary farm optimization framework. The first
and most self-evident next steps are to improve and realize the full computational efficiency benefits of the multi-fidelity
375 capabilities. In this work, the multi-fidelity Gaussian process was developed as a two-step fit, in which fixed length scales
were used to simplify the parameter fitting problem; the framework offers the possibility of simultaneous fitting of the full
multi-fidelity Gaussian process stack, which can allow improved convergence. Additionally, the sampling points used for
the high-fidelity method are generated using the single-fidelity Bayesian quadrature points. It was beyond the scope of this
manuscript to fully develop a method for multi-fidelity-optimal Bayesian quadrature points, which would improve the ability
380 of the method to maximize the utility of the high-fidelity evaluations in the context of the availability of lower-fidelity data.

A second area of improvement is to minimize the modeling differences between FLORIS and WindSE. While we worked
to do so shrink these differences, key gaps remained. For example, a divergence in the amount of peak thrust shaving around
rated wind speed continued to persist and creates a source of modeling error between the underlying data that drives our
AEP estimates. Before application for a production farm, achieving best-possible alignment between the low- and high-fidelity
385 solvers would be necessary.

In addition to the improvement to multi-fidelity AEP estimation, expanded optimization capabilities are another key area
for future work. The intent of this work was industrially computable CFD-enabled AEP/LCOE optimization, but the necessity
of multi-start optimizations increased the computational work necessary for the results in this manuscript beyond what might
typically be available to a wind farm developer. In general, wind turbine layout optimizations involve extremely complex re-
390 lations (Ning, 2020), and while a range of optimization methods have had reasonable success in overcoming them, there are
methods that are more efficient and reliable than others for the wind farm layout optimization problem (Thomas et al., 2023).
Some potential avenues for improving our optimization efficiency and reliability would involve using the current inner-loop
multi-fidelity AEP estimation with other optimization techniques. Using a global search approach could provide more consist-
ent optimization results, provided we do not use too many design variables. Using gradient-based algorithms and leveraging
395 the adjoint derivative calculations available in WindSE, the tool could reduce the computational cost and allow us to include
more design variables in the optimization problem. We could also use a combined inner- and outer-loop multi-fidelity strategy
to effectively integrate high-fidelity data while exploring broadly in the entire design space without requiring high-fidelity
samples at every step.



400 A third, more complex, and perhaps more critical extension of this work is including the variation of CapEx and OpEx
with respect to wind farm layout. If a wind farm design significantly increases loads, leading to component fatigue, then
these variations will be important, especially for tightly spaced farms. Unfortunately, integrating load data is an outstanding
challenge; the use of mid-fidelity load data is a possible approach, but it can be challenging to incorporate into optimization
procedures. On the other hand, surrogate load modeling has shown promise, but with challenges in its applicability beyond
its training space (Shaler et al., 2022). Incorporating load estimation in the wind farm layout optimization problem is a prime
405 candidate for the application of our multi-fidelity framework. However, once load estimates are created, the load information
will still need to be incorporated into CapEx and OpEx models. The impact of loads on CapEx could be included in a fairly
straightforward manner using turbine design tools, but including loads in OpEx models is much more complicated. Successful
modeling of the connections between loads and OpEx will likely require the use of historical datasets, which tend to consist of
protected proprietary industrial data.

410 A fourth area of future work consists of expanding the framework to offshore wind farms, for both fixed and floating founda-
tion turbines. This extension will require the integration of additional models and constraints. The extension to offshore will
further complicate the solution space, making conclusive convergence to a global optimal more difficult for the optimization
algorithm, though it may not ultimately be necessary given many similar local optima.

Code availability. All elements of the numerical toolchain described in this paper are publicly available and released by NREL under the
415 permissive Apache 2.0 license. FLORIS v4.1 is available at <https://github.com/NREL/floris>. WindSE is available at <https://github.com/NREL/WindSE>. WISDEM v3.16.4 is available at <https://github.com/wisdem/wisdem>. These three numerical tools are also available in the conda-
forge channel of Anaconda <https://anaconda.org/conda-forge>. The framework implementing the Gaussian process and integrating FLORIS,
WindSE, and WISDEM into a single workflow is available at <https://github.com/wisdem/ard>.

Author contributions. CF led both the technical work and the preparation of this manuscript. PB helped define the scope of work and the
420 architecture of the numerical framework. CJB and JT helped with the integration of FLORIS and the parametrization of the wind farms. EY
and JA helped with the integration of WindSE. All authors contributed to the manuscript.

Competing interests. The authors declare that no competing interests are present.

Acknowledgements. The research was performed using computational resources sponsored by the Department of Energy's Office of Energy
Efficiency and Renewable Energy and located at the National Renewable Energy Laboratory. This work was authored by the National
425 Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract
No. DE-AC36-08GO28308. Funding provided by the U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Wind



Energy Technologies Office. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for
430 U.S. Government purposes.



References

- Bastankhah, M. and Porté-Agel, F.: Experimental and theoretical study of wind turbine wakes in yawed conditions, *Journal of Fluid Mechanics*, 806, 506–541, <https://doi.org/10.1017/jfm.2016.595>, 2016.
- 435 Bortolotti, P., Tarres, H. C., Dykes, K., Merz, K., Sethuraman, L., Verelst, D., and Zahle, F.: IEA Wind Task 37 on Systems Engineering in Wind Energy – WP2.1 Reference Wind Turbines, Tech. rep., International Energy Agency, <https://www.nrel.gov/docs/fy19osti/73492.pdf>, 2019.
- Crespo, A. and Hernández, J.: Turbulence characteristics in wind-turbine wakes, *Journal of Wind Engineering and Industrial Aerodynamics*, 61, 71–85, [https://doi.org/https://doi.org/10.1016/0167-6105\(95\)00033-X](https://doi.org/https://doi.org/10.1016/0167-6105(95)00033-X), 1996.
- 440 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.
- Göçmen, T., van der Laan, P., Réthoré, P.-E., Diaz, A. P., Larsen, G. C., and Ott, S.: Wind turbine wake models developed at the technical university of Denmark: A review, *Renewable and Sustainable Energy Reviews*, 60, 752–769, <https://doi.org/https://doi.org/10.1016/j.rser.2016.01.113>, 2016.
- 445 Gray, J. S., Hwang, J. T., Martins, J. R. R. A., Moore, K. T., and Naylor, B. A.: OpenMDAO: An open-source framework for multidisciplinary design, analysis, and optimization, *Structural and Multidisciplinary Optimization*, 59, 1075–1104, <https://doi.org/10.1007/s00158-019-02211-z>, 2019.
- Hein, C. and Straw, B.: Proceedings from the State of the Science and Technology for Minimizing Impacts to Bats from Wind Energy, Tech. Rep. TP-5000-78557, NREL, <https://doi.org/10.2172/1772994>, 2021.
- 450 Herbert-Acero, J. F., Probst, O., Réthoré, P.-E., Larsen, G. C., and Castillo-Villar, K. K.: A Review of Methodological Approaches for the Design and Optimization of Wind Farms, *Energies*, 7, 6930–7016, <https://doi.org/10.3390/en7116930>, 2014.
- Jasa, J., Bortolotti, P., Zalkind, D., and Barter, G.: Effectively Using Multifidelity Optimization for Wind Turbine Design, *Wind Energy Science*, 7, 991–1006, <https://doi.org/10.5194/wes-7-991-2022>, 2022.
- 455 Katic, I., Højstrup, J., and Jensen, N.: A Simple Model for Cluster Efficiency, in: EWEC’86. Proceedings. Vol. 1, edited by Palz, W. and Sesto, E., pp. 407–410, A. Raguzzi, european Wind Energy Association Conference and Exhibition, EWEC ’86 ; Conference date: 06-10-1986 Through 08-10-1986, 1987.
- Keane, A.: Advancement of an analytical double-Gaussian full wind turbine wake model, *Renewable Energy*, 171, 687–708, <https://doi.org/https://doi.org/10.1016/j.renene.2021.02.078>, 2021.
- Kennedy, M.: Predicting the Output from a Complex Computer Code When Fast Approximations Are Available, *Biometrika*, 87, 1–13, <https://doi.org/10.1093/biomet/87.1.1>, 2000.
- 460 King, J., Fleming, P., King, R., Martínez-Tossas, L. A., Bay, C. J., Mudafort, R., and Simley, E.: Control-oriented model for secondary effects of wake steering, *Wind Energy Science*, 6, 701–714, <https://doi.org/10.5194/wes-6-701-2021>, 2021.
- 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.
- 465 King, R. N., Dykes, K., Graf, P., and Hamlington, P. E.: Optimization of Wind Plant Layouts Using an Adjoint Approach, *Wind Energy Science*, 2, 115–131, <https://doi.org/10.5194/wes-2-115-2017>, 2017.



- Lambe, A. B. and Martins, J. R. R. A.: Extensions to the Design Structure Matrix for the Description of Multidisciplinary Design, Analysis, and Optimization Processes, *Structural and Multidisciplinary Optimization*, 46, 273–284, <https://doi.org/10.1007/s00158-012-0763-y>, 2012.
- 470 Le Gratiet, L.: Multi-Fidelity Gaussian Process Regression for Computer Experiments, Ph.D. thesis, Université de Paris-Diderot, Paris, 2013.
- LeBeau, C. W., Johnson, G. D., Holloran, M. J., Beck, J. L., Nielson, R. M., Kauffman, M. E., Rodemaker, E. J., and McDonald, T. L.: Greater Sage-grouse Habitat Selection, Survival, and Wind Energy Infrastructure, *The Journal of Wildlife Management*, 81, 690–711, <https://doi.org/10.1002/jwmg.21231>, 2017.
- Logg, A., Mardal, K.-A., Wells, G. N., et al.: *Automated Solution of Differential Equations by the Finite Element Method*, Springer, ISBN 978-3-642-23098-1, <https://doi.org/10.1007/978-3-642-23099-8>, 2012.
- 475 Lopez, A., Cole, W., Sergi, B., Levine, A., Carey, J., Mangan, C., Mai, T., Williams, T., Pinchuk, P., and Gu, J.: Impact of Siting Ordinances on Land Availability for Wind and Solar Development, *Nature Energy*, 8, 1034–1043, <https://doi.org/10.1038/s41560-023-01319-3>, 2023.
- Martins, J. R. R. A. and Ning, A.: *Engineering Design Optimization*, Cambridge University Press, 1 edn., ISBN 978-1-108-98064-7 978-1-108-83341-7, <https://doi.org/10.1017/9781108980647>, 2021.
- 480 Meyer Forsting, A. R., Navarro Diaz, G. P., Segalini, A., Andersen, S. J., and Ivanell, S.: On the accuracy of predicting wind-farm blockage, *Renewable Energy*, 214, 114–129, <https://doi.org/https://doi.org/10.1016/j.renene.2023.05.129>, 2023.
- Mitusch, S. K., Funke, S. W., and Dokken, J. S.: dolfin-adjoint 2018.1: automated adjoints for FEniCS and Firedrake, *Journal of Open Source Software*, 4, 1292, <https://doi.org/10.21105/joss.01292>, 2019.
- Moriarty, P., Bodini, N., Cheung, L., Hamilton, N., Herges, T., Kaul, C., Letizia, S., Pekour, M., and Simley, E.: Overview of Recent
485 Observations and Simulations from the American WAKE experimeNt (AWAKEN) Field Campaign, *Journal of Physics: Conference Series*, 2505, 012 049, <https://doi.org/10.1088/1742-6596/2505/1/012049>, 2023.
- Mowers, M. and Mai, T.: An evaluation of electricity system technology competitiveness metrics: The case for profitability, *The Electricity Journal*, 34, 106 931, <https://doi.org/https://doi.org/10.1016/j.tej.2021.106931>, 2021.
- Niyayifar, A. and Porté-Agel, F.: Analytical Modeling of Wind Farms: A New Approach for Power Prediction, *Energies*, 9,
490 <https://doi.org/10.3390/en9090741>, 2016.
- Ning, A.: *Wind Farm Layout Optimization Test Cases*, 2020.
- NREL: FLORIS. Version 4.1, <https://github.com/NREL/floris>, 2024.
- 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.
- 495 Quon, E., Tripp, C., Sandhu, R., Thedin, R., Hein, C., Sara, B., Thresher, B., Farmer, C., Owen, T., Miller, T., and Brandes, D.: Progress Towards a Predictive Eagle Behavior and Risk Modeling Framework: Overview and Recent Validation Efforts, Tech. Rep. PR-5000-84549, NREL, <https://www.osti.gov/biblio/1902138>, 2022.
- Sanchez Gomez, M., Lundquist, J. K., Mirocha, J. D., and Arthur, R. S.: Investigating the physical mechanisms that modify wind plant blockage in stable boundary layers, *Wind Energy Science*, 8, 1049–1069, <https://doi.org/10.5194/wes-8-1049-2023>, 2023.
- 500 Shaler, K., Jasa, J., and Barter, G. E.: Efficient Loads Surrogates for Waked Turbines in an Array, *Journal of Physics: Conference Series*, 2265, 032 095, <https://doi.org/10.1088/1742-6596/2265/3/032095>, 2022.
- Strickland, J. M. and Stevens, R. J.: Investigating wind farm blockage in a neutral boundary layer using large-eddy simulations, *European Journal of Mechanics - B/Fluids*, 95, 303–314, <https://doi.org/https://doi.org/10.1016/j.euromechflu.2022.05.004>, 2022.



- Thomas, J. J., Baker, N. F., Malisani, P., Quaeghebeur, E., Sanchez Perez-Moreno, S., Jasa, J., Bay, C., Tilli, F., Bieniek, D., Robinson, N.,
 505 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.
- U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy: Atmosphere to Electrons (A2e),
<https://doi.org/10.21947/2287629>, 2023.
- Veers, P., Dykes, K., Baranowski, R., Bay, C., Bortolotti, P., Doubrawa, P., MacDonald, S., Rooney, S., Bottasso, C. L., Fleming, P., Haupt,
 510 S. E., Hale, A., Hein, C., and Robertson, A.: Grand Challenges Revisited: Wind Energy Research Needs for a Global Energy Transition,
 Tech. Rep. TP-5000-86564, NREL, <https://doi.org/10.2172/2229554>, 2023.
- Yamani Douzi Sorkhabi, S., Romero, D. A., Yan, G. K., Gu, M. D., Moran, J., Morgenroth, M., and Amon, C. H.: The Im-
 pact of Land Use Constraints in Multi-Objective Energy-Noise Wind Farm Layout Optimization, *Renewable Energy*, 85, 359–370,
<https://doi.org/10.1016/j.renene.2015.06.026>, 2016.
- 515 Zhang, S., Gao, X., Ma, W., Lu, H., Lv, T., Xu, S., Zhu, X., Sun, H., and Wang, Y.: Derivation and verification of
 three-dimensional wake model of multiple wind turbines based on super-Gaussian function, *Renewable Energy*, 215, 118968,
<https://doi.org/https://doi.org/10.1016/j.renene.2023.118968>, 2023.

Appendix A: Formulation of multi-fidelity AEP integrator

A1 Gaussian processes

- 520 A two-dimensional Gaussian process (GP) $f = \mathcal{GP}[m, k]$ is defined by a mean function $m = m(\mathbf{x}; \theta)$ and a covariance function
 $k(\mathbf{x}, \mathbf{x}'; \theta)$, which describe the mean value of a function and the covariance of two observations \mathbf{x} and \mathbf{x}' , respectively. In this
 work, we exclusively use a zero mean function $\mu = 0$ and Matérn 5/2 covariance function given by:

$$k(\mathbf{x}, \mathbf{x}'; \theta) = \sigma^2 \left(1 + \sqrt{5 \sum_j \frac{(x_j - x'_j)^2}{l_j^2}} + \frac{5}{3} \sum_j \frac{(x_j - x'_j)^2}{l_j^2} \right) \exp \left(- \sqrt{5 \sum_j \frac{(x_j - x'_j)^2}{l_j^2}} \right), \quad (\text{A1})$$

- such that σ^2 is a variance parameter and l is a length scale parameter. The corresponding correlation function can be given by
 525 $r(\mathbf{x}, \mathbf{x}'; \theta) = k(\mathbf{x}, \mathbf{x}'; \theta) / \sigma^2$.

Given a set of observations $\{y^{(k)}\}_k$ that are made at points $\{\mathbf{x}^{(k)}\}_k$, the mean of a Gaussian process conditioned on the data
 $\mathcal{D} = \{\mathbf{x}^{(k)}, y^{(k)}\}_k$ is given by

$$\mu_{\mathcal{GP}|\mathcal{D}}(\mathbf{x}) = \mathbf{k}(\mathbf{x})^\top \mathbf{K}^{-1} \mathbf{y}, \quad (\text{A2})$$

- where $[\mathbf{y}]_k = y^{(k)}$ represents the data matrix, $[\mathbf{k}(\mathbf{x})]_k = k(\mathbf{x}, \mathbf{x}^{(k)})$ gives the covariance of a point \mathbf{x} that is queried by the
 530 Gaussian process with respect to the observed points, and $\mathbf{K}_{jk} = k(\mathbf{x}^{(j)}, \mathbf{x}^{(k)})$ is the covariance matrix between the observed
 points. Likewise, the covariance function of the Gaussian process conditioned on \mathcal{D} can be given by:

$$k_{\mathcal{GP}|\mathcal{D}}(\mathbf{x}, \mathbf{x}') = k(\mathbf{x}, \mathbf{x}') - \mathbf{k}(\mathbf{x})^\top \mathbf{K}^{-1} \mathbf{k}(\mathbf{x}'). \quad (\text{A3})$$



A2 Integral quantities of single-fidelity Gaussian processes

Following King et al. (2020), the mean and covariance functions of a Gaussian process are integrable on some domain Ω , and
 535 thus, probabilistic integral quantities can be written as

$$I = \int_{\mathbf{x}} p(\mathbf{x}) f(\mathbf{x}) d\mathbf{x} \quad (\text{A4})$$

for some probability density function p . Then we can express the quantity I in terms of a normal distribution

$$I \sim \mathcal{N}(\alpha, \beta), \quad (\text{A5})$$

with

$$540 \quad \alpha = \int_{\mathbf{x}} p(\mathbf{x}) \mu_{\mathcal{GP}|\mathcal{D}}(\mathbf{x}) d\mathbf{x} = \mathbf{w}^\top \mathbf{K}^{-1} \mathbf{y} \quad (\text{A6})$$

and

$$\beta^2 = \int_{\mathbf{x}'} \int_{\mathbf{x}} p(\mathbf{x}) k_{\mathcal{GP}|\mathcal{D}}(\mathbf{x}, \mathbf{x}') p(\mathbf{x}') d\mathbf{x} d\mathbf{x}' = Z - \mathbf{w}^\top \mathbf{K}^{-1} \mathbf{w}. \quad (\text{A7})$$

The latter expressions for α and β show how the expressions for the integrated quantity I can be decomposed into

$$[\mathbf{w}]_j = \int_{\mathbf{x}} p(\mathbf{x}) k(\mathbf{x}, \mathbf{x}^{(j)}) d\mathbf{x}, \quad (\text{A8})$$

545 and Z ,

$$Z = \int_{\mathbf{x}'} \int_{\mathbf{x}} p(\mathbf{x}) k(\mathbf{x}, \mathbf{x}') p(\mathbf{x}') d\mathbf{x} d\mathbf{x}', \quad (\text{A9})$$

which along with \mathbf{K} are independent of the data. Moreover, Z is independent of the choice of observation points $\mathbf{x}^{(k)}$, allowing for the observation points to be chosen by maximizing the $\mathbf{w}^\top \mathbf{K}^{-1} \mathbf{w}$ to minimize β , which represents the uncertainty in the estimate of α within I , without a priori knowledge of \mathbf{y} .

550 A3 Multi-fidelity Gaussian processes

The recursive multi-fidelity co-kriging method of Le Gratiet (2013), which they prove to be equivalent to the multi-fidelity co-kriging method of Kennedy (2000), seeks to approximate the output of some model $y_t^{(k)}(\mathbf{x})$ by a Gaussian process $z_t(\mathbf{x})$, where $y_1(\mathbf{x}), y_2(\mathbf{x}), \dots, y_t(\mathbf{x}), \dots, y_s(\mathbf{x})$ give the models in order from low to high fidelity. It is presumed that the cost of computing the output of y_t becomes more expensive as t gets larger, but that these computations produce higher-accuracy outputs.

555 The multi-fidelity co-kriging method attempts to approximate what the output of y_t would be at unsampled points, by leveraging a moderate number of samples of y_t and progressively larger samples of $y_{t-1}, y_{t-2}, \dots, y_1$. In order to do so, the



multi-fidelity model takes the form of:

$$\begin{aligned}
 z_t(\mathbf{x}) &= \rho_{t-1}(\mathbf{x})z_{t-1}(\mathbf{x}) + \delta_t(\mathbf{x}) \\
 z_{t-1}(\mathbf{x}) &\perp \delta_t(\mathbf{x}) \\
 \rho_{t-1}(\mathbf{x}) &= \mathbf{g}'_{t-1}(\mathbf{x})\boldsymbol{\beta}_{\rho_{t-1}},
 \end{aligned} \tag{A10}$$

560 where ρ_{t-1} is a regression field, based on the set of regression basis functions \mathbf{g}_{t-1} and coefficients $\boldsymbol{\beta}_{\rho_{t-1}}$, and δ_t is a Gaussian process that is orthogonal to the next-lower Gaussian process z_{t-1} . The correcting function δ_t has a mean function $\mathbf{f}'_t(x)\boldsymbol{\beta}_t$, based on the set of regression basis functions \mathbf{f}_t and coefficients $\boldsymbol{\beta}_t$, and a covariance kernel $\sigma_t^2 r_t(\mathbf{x}, \mathbf{x}')$, where r_t gives the covariance kernel and σ_t^2 is a parameter.

By applying these requirements and definitions, Le Gratiet (2013) show that the Gaussian process z_t is given by

$$z_t(\mathbf{x}) \sim \mathcal{N}(\mu_{z_t}(\mathbf{x}), s_{z_t}^2(\mathbf{x})) \tag{A11}$$

565 with

$$\mu_{z_t}(\mathbf{x}) = \rho_{t-1}(\mathbf{x})\mu_{z_{t-1}}(\mathbf{x}) + \mathbf{f}'_t(\mathbf{x})\boldsymbol{\beta}_t + \mathbf{r}'_t(\mathbf{x})\mathbf{R}_t^{-1}(\mathbf{y}_t - \rho_{t-1}(\mathcal{D}_t) \odot z_{t-1}(\mathcal{D}_t) - \mathbf{F}_t\boldsymbol{\beta}_t), \tag{A12}$$

$$s_{z_t}^2(\mathbf{x}) = \text{Cov}_{z_t}(\mathbf{x}, \mathbf{x}), \tag{A13}$$

and

$$570 \text{Cov}_{z_t}(\mathbf{x}, \mathbf{x}') = \rho_{t-1}(\mathbf{x})\text{Cov}_{z_{t-1}}(\mathbf{x}, \mathbf{x}')\rho_{t-1}(\mathbf{x}') + \sigma_t^2(1 - \mathbf{r}'_t(\mathbf{x})\mathbf{R}_t^{-1}\mathbf{r}_t(\mathbf{x}')). \tag{A14}$$

In the above, \mathcal{D}_t represents the points \mathbf{x} sampled on level t , \mathbf{y}_t represents the sampled output on \mathcal{D}_t , and \mathbf{F}_t represents \mathbf{f}_t sampled on \mathcal{D}_t . The \odot symbol represents a pointwise multiplication. The correlation matrix is given by $[\mathbf{R}_t]_{jk} = r(\mathbf{x}^{(j)}, \mathbf{x}^{(k)})$ and the correlation function $[\mathbf{r}_t]_j(\mathbf{x}) = r(\mathbf{x}, \mathbf{x}^{(j)})$.

A4 Integral quantities of multi-fidelity Gaussian processes

575 If we take f in Eq. (A4) to be a multi-fidelity Gaussian process, then Eq. (A5) holds with

$$\alpha = \int_{\mathbf{x}} p(\mathbf{x})\mu_{\text{MF}|\mathcal{D}}(\mathbf{x}) d\mathbf{x} \tag{A15}$$

and

$$\beta^2 = \int_{\mathbf{x}'} \int_{\mathbf{x}} p(\mathbf{x})k_{\text{MF}|\mathcal{D}}(\mathbf{x}, \mathbf{x}')p(\mathbf{x}') d\mathbf{x} d\mathbf{x}' \tag{A16}$$



where $\mu_{\text{MF}|\mathcal{D}}(\mathbf{x}) = \mu_{z_s}(\mathbf{x})$ and $k_{\text{MF}|\mathcal{D}}(\mathbf{x}, \mathbf{x}') = \text{Cov}_{z_s}(\mathbf{x}, \mathbf{x}')$. This results in α and β of the form:

$$\begin{aligned}
 \alpha = & \int_{\mathbf{x}} p(\mathbf{x}) \rho_{t-1}(\mathbf{x}) \mu_{z_{t-1}}(\mathbf{x}) d\mathbf{x} + \left(\int_{\mathbf{x}} p(\mathbf{x}) \mathbf{f}'_t(\mathbf{x}) d\mathbf{x} \right) \beta_t \\
 & + \left(\int_{\mathbf{x}} p(\mathbf{x}) \mathbf{r}'_t(\mathbf{x}) d\mathbf{x} \right) \mathbf{R}_t^{-1} (\mathbf{y}_t - \rho_{t-1}(\mathcal{D}_t) \odot z_{t-1}(\mathcal{D}_t) - \mathbf{F}_t \beta_t)
 \end{aligned} \tag{A17}$$

and

$$\begin{aligned}
 \beta^2 = & \int_{\mathbf{x}'} \int_{\mathbf{x}} p(\mathbf{x}) \rho_{t-1}(\mathbf{x}) \text{Cov}_{z_{t-1}}(\mathbf{x}, \mathbf{x}') \rho_{t-1}(\mathbf{x}') p(\mathbf{x}') d\mathbf{x} d\mathbf{x}' + \sigma_t^2 \\
 & - \sigma_t^2 \left(\int_{\mathbf{x}} p(\mathbf{x}) \mathbf{r}'_t(\mathbf{x}) d\mathbf{x} \right) \mathbf{R}_t^{-1} \left(\int_{\mathbf{x}'} \mathbf{r}_t(\mathbf{x}') p(\mathbf{x}') d\mathbf{x}' \right),
 \end{aligned} \tag{A18}$$

which exhibit a more complicated – but no less feature-rich – form than their single-fidelity integrations in Eq. (A6) and Eq. (A7).