Effectively using multifidelity optimization for wind turbine design

John Jasa[1], Pietro Bortolotti[1], Daniel Zalkind[1], and Garrett Barter[1]

[1]National Renewable Energy Laboratory, Boulder, CO

Correspondence: John Jasa (john.jasa@nrel.gov)

Abstract. Wind turbines are complex multidisciplinary systems that are challenging to design because of the tightly coupled interactions between different subsystems. Computational modeling attempts to resolve these couplings so we can efficiently explore new wind turbine systems early in the design process. Low-fidelity models are computationally efficient but make assumptions and simplifications that limit the accuracy of design studies, whereas high-fidelity models capture more of the

5 actual physics but with increased computational cost. This paper details the use of multifidelity methods for optimizing wind turbine designs by using information from both low- and high-fidelity models to find an optimal solution at reduced cost. Specifically, a trust-region approach is used with a novel corrective function built from a nonlinear surrogate model. We find that for a diverse set of design problems—with examples given in rotor blade geometry design, wind turbine controller design, and wind power plant layout optimization—the multifidelity method finds the optimal design using 38%–58% the

10 computational cost of the high-fidelity-only optimization. The success of the multifidelity method in disparate applications suggests that it could be more broadly applied to other wind energy or otherwise generic applications.

1 Introduction

Wind turbines are complex systems, where aerodynamic, structural, acoustic, controls, manufacturing, logistics, and technoe-conomic considerations are all design drivers. To design the optimal wind energy system, multidisciplinary design optimization

15 (MDO) approaches help capture the interconnected trade-offs among these disciplines while dramatically reducing the times and costs of design processes compared to sequential single-discipline design approaches. The past two decades have seen the development of a variety of MDO models for wind turbine design, such as Giguere and Selig (2000); Fuglsang and Madsen (1999); Ning et al. (2014); Ashuri et al. (2014); Fischer et al. (2014); Pourrajabian et al. (2016); Ning and Petch (2016); Bortolotti et al. (2016); Barlas et al. (2020), among many others.

20 Choosing the correct fidelity level of analyses used in the MDO process is a crucial decision for the designer, who must meet the need of reasonable accuracy with tractable computational costs. Although lower fidelity tools offer the possibility to explore a broad solution space and investigate uncommon design choices thanks to the lower computational costs, they often run the risk of oversimplifying the design problem, which could lead to solutions that in reality might underperform or violate unresolved constraints. In contrast, higher fidelity models are usually incompatible with numerical optimizations that

25 rely on hundreds or thousands of function evaluations. Higher computational costs can therefore be tolerated only when doing spot-checks and potential design changes are small.

During the design process, model fidelity usually increases from low to high as more configuration and sizing choices are finalized. Low-fidelity models are usually applied during the early, conceptual stages of design, when many different options, architectures, goals, constraints, etc., are being considered. Because designing these systems using only high-fidelity tools would lead to impossibly long development cycles caused by the computational expense, designers generally start with low-fidelity tools and increase simulation fidelity as the design cycle progresses. This approach usually works fairly well in practice, but it requires frequent interventions and the expertise of the designers. More importantly, it still runs the risk of leading to suboptimal design solutions when the designs are evaluated using more accurate models.

An alternative pathway consists of adopting formal multifidelity design optimization approaches. These methods capture realistic physical trends while reducing the computational cost compared to optimizing using only high-fidelity methods. For applications where single-fidelity design and model iterations work well, multifidelity approaches simply make the design process more efficient by combining information from the low- and high-fidelity models.

Multifidelity optimization methods have a long history across multiple fields, including applied mathematics (Kennedy and O'Hagan, 2000; Forrester et al., 2007; Peherstorfer et al., 2018), aerospace engineering (Robinson et al., 2008; March and Willcox, 2012), and wind energy. Specifically for wind energy, multifidelity methods have been used for aeroelastic blade design (Maki et al., 2012; McWilliam et al., 2017; Abdallah et al., 2019), wind plant layout optimization (Rahbari et al., 2014; Réthoré et al., 2014), and wake steering uncertainty quantification (Quick et al., 2019).

Delving deeper into the literature on multifidelity optimization within wind energy, Maki et al. (2012) use a series of nested and sequential optimizations along with metamodels to minimize cost of energy for a given turbine design considering multiple fidelities. Réthoré et al. (2014) introduce TOPFARM, a tool for multifidelity layout optimization of wind farms, and they demonstrate sequential optimizations at increasing levels of fidelity to show how optimal results from one model can speed up the design process for a higher fidelity model. McWilliam et al. (2017) used an approximation model management framework (AMMF) approach to perform multifidelity aerostructural optimization of a wind turbine blade. That work established AMMF as a reasonable tool to enable multifidelity blade design, but it showed that the additional complexity of the AMMF algorithm led to slower overall convergence than high-fidelity-only optimization. This shows that the efficacy of multifidelity methods to wind turbine design is both method- and model-dependent, and there is room for improvement for developing a method that works more generally without expert intervention.

Building on previous work, questions remain about how to best use multifidelity methods for different wind energy applications. For example, depending on the relative computational cost between the low- and high-fidelity models, different approaches might be more effective. Additionally, the optimization problem size directly impacts how effective certain methods are. We address these questions and examine how multifidelity methods can achieve high-performing designs using less computational cost.

In this paper, we present best practices for using multifidelity optimization methods for wind energy design applications. We do so by first detailing a trust-region-based multifidelity method with a novel correction function built on top of nonlinear surrogate models. We then formulate and solve three optimization problems: aerodynamic blade design for the IEA 15-MW reference wind turbine; a controls optimization using both linearized and nonlinear state-space models; and a wind power

plant layout problem. The problems studied here were selected as they examine different disciplines within the larger wind plant design problem and serve as meaningful representative cases to benchmark the multifidelity optimization method against commonly-used single-fidelity methods. By solving disparate optimization problems involving different simulation models,

65 we demonstrate how multifidelity methods can be effectively used for the design of complex wind energy systems.

Section 2 introduces the trust-region multifidelity optimization method we have implemented. The next three sections then compare the computational cost and design performance of the multifidelity optimization method as compared to both low- and high-fidelity optimization. Sections 3, 4, and 5 examine case studies concerning aerodynamic blade design, controller design, and power plant layout, respectively. Each case study section details the design problem, approach, and tools used, which differ

70 for each study. Lastly, Section 6 presents key findings and takeaways, including which problem types within wind energy are best suited for multifidelity optimization.

2 Multifidelity optimization methodology

2.1 Multifidelity methods

In an optimization problem, we seek the minimum of a function within a design space subject to arbitrary constraints. The

75 optimizer selects design variable values, evaluates computational models at that design point to obtain objective and constraint values, and then repeats until convergence is reached. In multifidelity optimization, multiple different types of computational models are queried and information from each model is combined to determine where to sample the design space next. A comprehensive survey of multifidelity methods is presented by Peherstorfer et al. (2018), where various approaches are categorized as one of *adaption*, *fusion*, or *filtering*, with guidelines for matching methods to application. In this work, we focus on the

80 multifidelity optimization method using the *adaptation* model management strategy and simplified physics-based low-fidelity models. We selected this model management strategy because it is non-intrusive and straightforward to implement in a general manner.

Here we loosely define *fidelity* as a qualitative measure of the accuracy of the underlying physical equations being modeled compared to the real world. Related to fidelity is the concept of *resolution*, or how finely discretized a domain or set of inputs

85 might be.

2.2 Trust-region optimization method

We use a trust-region approach to perform multifidelity optimization. This method is well-studied in the fields of applied mathematics, computational sciences, and aerospace engineering (Alexandrov et al., 1998, 2001; March and Willcox, 2012). It has also been used in other wind energy research, though those studies focused on different applications (Park and Law, 2015;

90 Yu et al., 2018) or used simplified corrective functions (McWilliam et al., 2017).

The trust-region method progression is shown in Figure 1 using the XDSM graphical data-flow format from Lambe and Martins (2012). The XDSM diagram format shows analysis and optimization processes as on-diagonal blocks in green and

blue respectively. Off-diagonal gray boxes show what data is passed between those process blocks and they are connected by gray lines to show data flow. Following this diagram for the trust-region method, the low- and high-fidelity models are first

95   called at a set of initial design points to establish the corrective function between the fidelities. Then, a subset of the design space is established where the corrected low-fidelity model is trusted. Local optimization within this region is then performed, and the high-fidelity model is queried at the locally optimal point. Based on the actual reduction in the objective value compared to the expected reduction, the trust region is either expanded or contracted. The local optimization is then repeated within this new trust region, and the process is repeated.
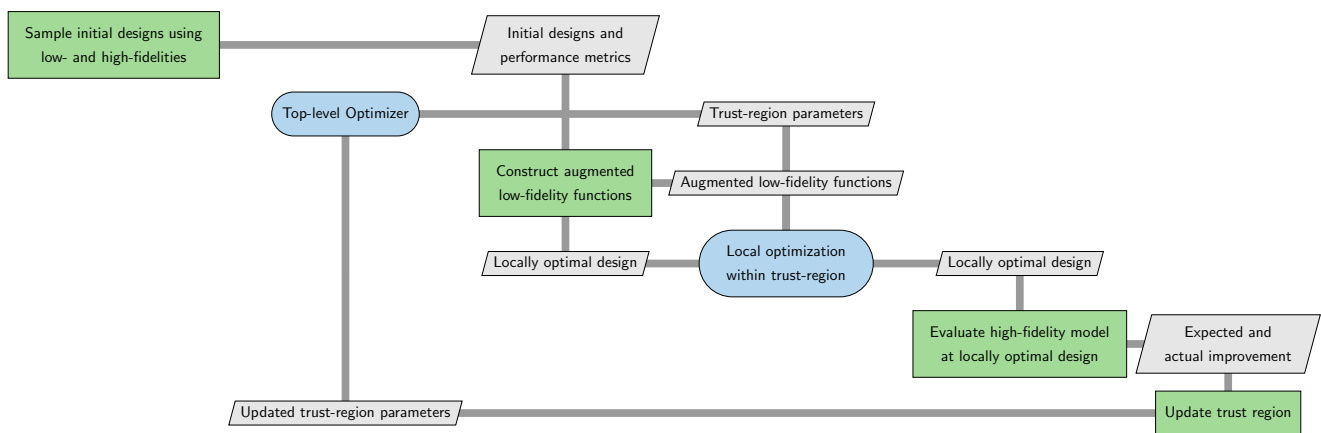
Figure 1. XDSM diagram of the trust-region method, specifically showing the top-level and inner-level optimizations.

100   A simple 1D example of how the trust region converges is shown in Figure 2, which highlights how the corrected low-fidelity model is used to approximate the high-fidelity model. The trust region for local optimization is shown with a gray bar.

2.3   Corrective function between low- and high-fidelity models

Within the trust-region method, we need to construct an approximation for the high-fidelity model using the low-fidelity model and a corrective function. This approximation is devised such that it is equal to the high-fidelity model at the points where

105   we have high-fidelity data. In previous work in wind energy (McWilliam et al., 2017), this corrective function was simply an additive and/or multiplicative factor, leading to a first-order linear correlation between the models. When dealing with nonlinear design spaces, it makes sense to use a more complex corrective function that can account for nonlinear differences between the fidelity levels because there is no assurance that the trust region is small enough to support a linear approximation of the high-fidelity space.

110   In this work, we use a nonlinear surrogate model to construct the corrective function as $f_{\text{corrected}} = f_{\text{low-fidelity}} + f_{\text{surrogate}}$, which allows us to capture arbitrary correlations between the models. This nonlinear surrogate formulation is especially useful when we do not have an *a priori* simple understanding of the correlation between the different fidelity levels, which is

(a) 0th iteration   (b) 2nd iteration   (c) 11th iteration

Figure 2. Trust-region optimization progression in the 1D example. In the 0th iteration, the initial trust region does not contain the entire design space. In the 2nd iteration, the trust region has grown as the local optimizer found the best answer at the bounds of the trust region. In the 11th iteration, the trust region narrows around the local minima.

common in physics-based modeling. Recent advances in surrogate modeling have increased the accuracy for a given model when using a fixed number of data points while simultaneously decreasing computational cost. One example is the Kriging

115 partial least squares (KPLS) method (Bouhlel et al., 2016), which is based on the Kriging method (Cressie, 1988). Typically, as problem dimensionality increases, the cost of training the surrogate model increases as well; however, KPLS has lower training costs than ordinary Kriging due to its internal dimension reduction, leading to a lower computational cost when training the model (Bouhlel et al., 2016). Additionally, the gain in surrogate accuracy is generally worth the increased cost compared to using a simple piecewise-linear fit.

120   Figure 3 shows the impact of the corrective function on the augmented low-fidelity model for a canonical 1D problem and relatively few data points. Figure 3a shows that using a simple piecewise linear fit achieves reasonable results, but it does not capture the high-fidelity function well. Sensibly, increasing the order of the corrective function might produce a better result, but Figure 3b shows that a piecewise cubic fit leads to a worse fit far from the data points. Last, Figure 3c shows how a KPLS-based corrective function does very well at capturing the high-fidelity model trends in between the data points.

125   Although the trends shown in these figures suggest that KPLS is the best corrective function, the performance and accuracy of these corrections is entirely problem-dependent. That said, these advanced surrogate modeling techniques generally capture multidimensional nonlinear correlations much better than more simplistic functions, especially when using a small number of high-fidelity data points. Additionally, if we wanted to obtain a better fit with the high-fidelity model, we could use gradient information at each data point to ensure that the corrected low-fidelity model has the same gradient values at those points.

130 For this work, however, we purposefully do not assume that we have any high-fidelity gradient information, which makes the methods presented here applicable to a wider range of real-world tools and applications. For the following case studies, we use a KPLS-based corrective function with the number of sampling points depending on the application. The KPLS implementation is taken from the open-source surrogate modeling toolbox based on the work of Bouhlel et al. (2019).

Figure 3. Corrective function formulation with (a) piecewise-linear reasonably approximating the higher fidelity model, (b) piecewise-cubic introducing unwanted oscillations in the approximation, and (c) KPLS-based corrective function matching the high-fidelity most accurately.

3 Case study: blade design optimization

135 This case study focuses on creating an aerodynamically optimal blade, a common and challenging problem in wind turbine design. Blades are commonly designed using relatively low-fidelity aerodynamic models, such as steady-state blade-element momentum theory (BEMT), which does not capture the effects that unsteady 3D flows have on blade performance. Using multifidelity optimization methods for blade design would allow for more accurate aerodynamic considerations earlier in the design cycle.

3.1 Model descriptions and tools used

140

The multifidelity optimization method is implemented in the Wind Energy with Integrated Servo-control (WEIS) framework (NREL, 2021c). WEIS is a new design tool that enables multifidelity wind turbine design by integrating the capabilities of multiple tools from the National Renewable Energy Laboratory (NREL). Of the numerous WEIS component models, the ones active in the first two case studies in this paper include the systems engineering framework Wind-Plant Integrated System Design & Engineering Model (WISDEM®) (NREL, 2021d), the aeroservoelastic solver OpenFAST (NREL, 2021b), the

145 auto-tuning Reference OpenSource Controller (ROSCO) (NREL, 2020), the wind solver TurbSim (Jonkman, 2009), as well as several pre- and post-processing routines. The primary goal of WEIS is to provide a framework for the controller codesign of floating wind turbines alongside turbine and platform geometry at multiple fidelity levels. In this paper, we do not include floating dynamics because incorporating that degree of complexity into the other case studies is the focus of future work. The

150 next subsections present more details on the models of WEIS adopted in this work and the formulation of the optimization problem.

Within WEIS, users have the option to individually activate WISDEM and OpenFAST, with additional customization available for all of the various sub-modules. In this way, numerous simulation pathways are available, creating a spectrum of fidelity options.

155 WISDEM is built using OpenMDAO (Gray et al., 2019), the open-source Python-based optimization framework developed at the National Aeronautics and Space Administration's Glenn Research Center. WISDEM models the wind turbine as an assembly of blocks, where each block models a specific component of the machine. The blocks are ordered following the load path—namely, from the blades toward the tower—and once the machine is sized, cost models are called to compute the levelized cost of energy. WISDEM computes only steady-state performance and loads and is therefore considered a lower

160 fidelity simulation tool.

Wind turbine aerodynamics in WISDEM are computed with the CCBlade module, which implements the formulation of the BEMT presented in Ning (2014), with hub and tip losses accounted for. The RotorSE module in WISDEM combines the CCBlade-computed aerodynamic loads with a 1D element beam solver, based on Frame3DD (Gavin, 2014), which accounts for centrifugal stiffening but otherwise assumes a rigid rotor with no aeroelastic iteration.

165 OpenFAST is a multiphysics, multifidelity tool for simulating the coupled dynamic response of wind turbines in the time domain. It is well represented in the literature and has undergone numerous validation studies. In this work, OpenFAST serves as the high-fidelity level of the multifidelity optimization approach.

The aerodynamics in OpenFAST are handled by the module AeroDyn15, whose theory is described in Moriarty and Hansen (2005). AeroDyn15 implements various permutations of the BEMT theory and, since recently, a free-wake vortex aerodynamic

170 model (Shaler et al., 2020). Among the unsteady effects, the airfoil aerodynamics include the Onera stall model. Full aeroelastic coupling is implemented in OpenFAST by combining the aerodynamic loads from AeroDyn with the blade structural dynamics simulated by ElastoDyn using Rayleigh-Ritz shape functions. The user can model the wind as a steady-state flow or via turbulent wind grids with the affiliated TurbSim (Jonkman, 2009) model. OpenFAST also includes two aerodynamic models of the tower—namely, the Powles and the Eames models—and couples the turbine elastic behavior with the rotor and tower

175 aerodynamics.

3.2 Optimization problem formulation

To study the efficacy of the trust region multifidelity method, we set up a simple blade design optimization case study using the IEA 15-MW reference wind turbine (Gaertner et al., 2020) as the baseline. This reference turbine has a rotor diameter of 242.2 m and a hub height of 150 m. The objective of the study is to maximize the electrical power of the generator at a given

180 wind speed, 9 m/s, by varying the blade twist and chord. Notably, the problem focuses only on rotor aerodynamic performance and does not consider any structural constraints or subsystem design constraints. For the low-fidelity model, we use the steady-state BEMT solver CCBlade, as described in Section 3.1. For the high-fidelity model, we use the unsteady BEMT solver within AeroDyn15 with the dynamic generator torque controller active. The inflow includes turbulence, and flapwise and edgewise blade flexibility is accounted for.

The blade and twist profiles along the blade are controlled by continuous spline interpolations. Each profile is independently parameterized using six control points, with the first two points fixed for both twist and chord and the outermost point fixed for chord. The twist control point design variables act as an adder on top of the original distribution, and the chord control point design variables act as a multiplier. This results in seven design variables, as shown in Table 1.

Table 1. Optimization problem formulation for the unconstrained power maximization case.

Category	Name	Quantity	Lower bound	Upper bound	Units
Objective	Power	1	–	–	MW
Design variables	Twist adder	4	-14.32	14.32	degrees
	Chord multiplier	3	0.5	1.50	–

To fairly evaluate the performance of the multifidelity method, we conducted three different blade chord and twist design optimizations:

1. Design optimization using only the low-fidelity model, CCBlade

2. Design optimization using only the high-fidelity model, AeroDyn15

3. Design optimization using the trust-region multifidelity method with the KPLS corrective function.

As a last step, we cross-checked the designs by computing the performance for each of the three blade shapes in both CCBlade and AeroDyn15.

3.3 Optimization results

The results of the single-fidelity and multifidelity optimizations are reported in Table 2. The optimization of the chord and twist in CCBlade achieves the highest power of the three designs when evaluated by CCBlade but the lowest of the three in AeroDyn15. On the contrary, the single-fidelity optimization in AeroDyn15 and the multifidelity optimization successfully identify the configuration generating the highest power in the high-fidelity model, with only small numerical differences in performance between them. This result is even more compelling when considering the computational cost of the three optimizations. The low-fidelity-only optimization completed in just under 9 central processing unit (CPU) hours with 360 calls to CCBlade. The high-fidelity-only optimization took nearly three times as long, with more than 1,511 function calls to Aero-Dyn15 due to the noisy gradients that are common in unsteady turbulent simulations. In contrast, the multifidelity optimization made only 63 calls to AeroDyn15 but over 2,500 calls to CCBlade, with a net CPU time of 10.61 hours, only 19% higher than the low-fidelity-only optimization.

Figure 4 shows the design solutions identified in the three optimizations. The single-fidelity optimization with CCBlade increases the power generation by simultaneously decreasing the twist and chord, effectively increasing the angles of attack along the blade span and narrowing the margin to stall. Both the single-fidelity optimization in AeroDyn15 and the multifidelity optimization choose instead the opposite route and increase the chord and twist, effectively reducing the angles of attack along

Table 2. Optimal results for aerodynamic blade power maximization.

Design	Lo-fi calls	Hi-fi calls	CPU time	Performance evaluation	
				Lo-fi power	Hi-fi power
	–	–	hrs	kW	kW
Low fidelity	360	–	8.91	9683.5	10733
High fidelity	–	1511	23.97	9646.1	11043
Multifidelity	2564	63	10.61	9412.5	11023

the blade span. The different design trends can be explained by the unsteadiness of the operational angles of attack at high fidelity caused by the turbulent wind. Such oscillations are less problematic with a higher twist and lower angles of attack, whereas when the blade operates at a lower twist and higher angles of attack, the turbulent wind frequently pushes the blade close to or into stall, increasing drag and decreasing power.
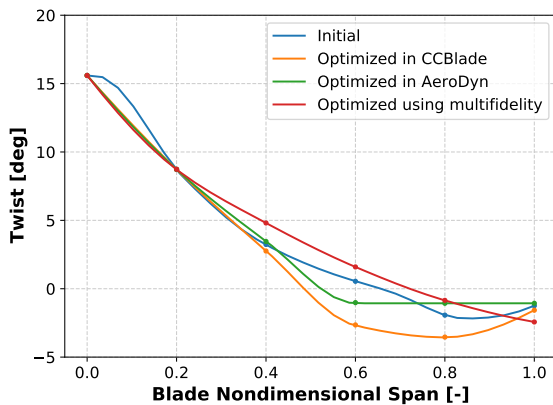
(a) Twist distributions along the blade span   (b) Chord distributions along the blade span

Figure 4. Baseline and optimized chord and twist distributions.

4 Case study: controls optimization

215

Optimal turbine control, or specifically determining how to vary the pitch and yaw of the turbine for optimal performance and longevity, is a complex and commonly-studied field. To demonstrate multifidelity optimization on a wind turbine control problem, we tune the control bandwidth, ω_{PC}, of the above-rated pitch controller to minimize tower fatigue loads with a constraint on the maximum generator speed. When the generator speed exceeds some limit, the supervisory controller triggers a
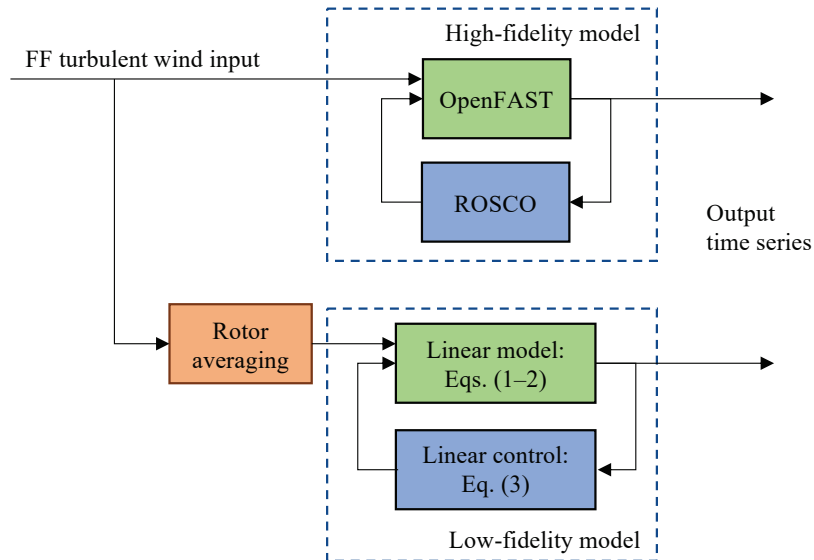
Figure 5. High- and low-fidelity control models. The high-fidelity model runs OpenFAST with the ROSCO controller and a full-field (FF) turbulent wind input. The low-fidelity models are described in (1)–(3) and use a rotor-averaged wind speed as the input. Both models output time series that can be processed to derive operational and load measures.

shutdown procedure, which reduces the net annual energy production (AEP). Tower loads drive the tower design and its capital expenditures. The pitch control bandwidth determines the proportional-integral (PI) gains of the blade pitch controller. Generally, lower bandwidths reduce tower loads but increase generator speed transients, so we expect the results of this optimization procedure to seek the lowest bandwidth such that the generator speed constraint is not violated.

4.1 Model descriptions and tools used

We simulate both a linearized and nonlinear version of the IEA 15-MW wind turbine with the University of Maine's VolturnUS semisubmersible (Allen et al., 2020) in extreme turbulence with a mean wind speed of 16 m/s. For the nonlinear simulation, we use OpenFAST with the ROSCO controller (NREL, 2020) and a full-field turbulent wind input generated using TurbSim. When this turbulent wind input is sampled by the blades, it results in 3P (per-revolution) oscillating loads on the tower. The nonlinear OpenFAST model is run for 800 simulation seconds, which requires approximately 3 minutes on a standard laptop computer, and represents the high-fidelity model for this case study.

To serve as the low-fidelity model, we simulate a linearized turbine and control model, which requires less than 3 seconds on a standard laptop computer. To create these low-fidelity models, we run OpenFAST in its linear mode, which creates linearized snapshots of the turbine at several azimuth positions for a fixed wind speed (Jonkman and Jonkman, 2016). These linear snapshots are averaged using the multiblade coordinate transform (Bir, 2010) to create a linear time-invariant system relative

235 to the turbine's operating points:

$$\dot{x} \quad = A(u_h)x + B(u_h)[u - u_{\text{op}}(u_h)] \tag{1}$$

$$y - y_{\text{op}}(u_h) \quad = C(u_h)x + D(u_h)[u - u_{\text{op}}(u_h)] \tag{2}$$

where u, x, and y are the inputs, states, and outputs of the linearized turbine, respectively. The input and output operating points, u_{op} and y_{op}, respectively, and the state-space matrices A, B, C, and D are determined during the OpenFAST linearization

240 process. When multiple wind speeds, u_h, are linearized, we construct a set of state-space systems, which can be interpolated based on the mean wind speed, u_h, so the system matrices and operating points are a function of u_h. In this study, we focus on above-rated control, and we linearize the turbine model at mean wind speeds of 14, 16, and 18 ms^{-1}.

For the pitch control input, which is part of u, we connect the output of a linearized ROSCO controller:

$$\theta_c = k_P(\omega_g - \omega_{\text{rat}}) + k_I \int (\omega_g - \omega_{\text{rat}}) + k_{\text{float}} \int \ddot{x}_{\text{IMU}}, \tag{3}$$

245 where k_P and k_I are the PI gains of the pitch controller, and $k_{\text{float}} = -9.49$ seconds is the floating feedback gain. The PI gains are a function of the bandwidth, ω_{PC}, and turbine parameters (Abbas et al., 2021). Generally, as the design variable ω_{PC} increases, the PI gains also increase.

The inputs to the controller are the generator speed, ω_g, and an acceleration measurement from the nacelle inertial measurement unit (IMU) in the nodding direction, $\ddot{x}_{\text{IMU}}$; these are in y. When the linear turbine and control models are connected, we

250 have a set of closed-loop linear turbine models that depend on the wind speed.

Instead of a full-field turbulent wind input, as in the high-fidelity model, the rotor average wind speed is used to simulate the linear model. The mean rotor average wind speed is used to determine the single closed-loop linear model from the set by linearly interpolating the state-space matrices and operating points. Then, we integrate the linear system over time, which results in a time series that is similar to the nonlinear model (Figure 6). Nonlinear aerodynamic and hydrodynamic

255 effects are not captured in the linear state-space model, but they are part of the operating points. In the linear simulations, a constant operating point is chosen for the whole 800-second simulation (with the first 200 seconds typically omitted as startup transients).

Both the linear and nonlinear turbine outputs can be processed to compute the generator speed maxima (constraint) and the damage equivalent loading (DELs) on the tower (objective), as shown in Figure 7. In general, trends, or changes, in the

260 linear and nonlinear models are in agreement and as expected: increasing the pitch control bandwidth increases tower DELs and platform motion while decreasing generator speed transients. The linear models do not capture the 3P harmonic loading on the tower, which accounts for most of the difference in the tower base fore-aft DELs between the two models. Finally, the magnitude of the optimization constraints (maximum generator speed and platform pitch angle) are more accurately sampled from the nonlinear simulations; therefore, these constraints are active only in the nonlinear simulations, which creates a good

265 stress test for the multifidelity optimization where some constraints are violated only in the high-fidelity simulation.

Figure 6. Comparison between the time histories generated from the nonlinear and linearized OpenFAST simulations. The linearized model generally has a smoother response than the nonlinear simulation.

4.2 Optimization problem formulation

The objective, design variables, and constraints for the controls optimization problems are shown in Table 3.

Table 3. Optimization problem formulation for the controls optimization case.

Category	Name	Lower bound	Upper bound	Units
Objective	Tower base moment DEL	–	–	–
Variables	ω_{PC}	0.1	0.4	–
Constraints	Generator speed	–	9	rpm
	Platform pitch	–	5	deg

Figure 7. The linearized and nonlinear functions of interest as we vary the pitch controller bandwidth, ω_{PC}.

4.3 Optimization results

As in the previous case study, we performed single-fidelity optimization using both the low- and high-fidelity models and
270 compared the results to the multifidelity trust-region method. Table 4 contains the optimal ω_{PC} values and the corresponding functions of interest from each optimization. Using the high-fidelity optimization to evaluate true performance, the low-fidelity-only optimization finds an infeasible solution that violates the generator speed constraint. Revisiting Figure 7, this is expected due to the linearized model not resolving the same magnitude or trends found in the nonlinear model. Although the optimal pitch control bandwidths in the high-fidelity and multifidelity optimizations differ, the actual difference in objective value is
275 relatively small, approximately 0.03%, and the constraints are satisfied in both cases.

Table 4. Optimal results for the controls problem.

Design	Lo-fi calls	Hi-fi calls	CPU time	ω_{PC}	Hi-fi DEL	Hi-fi max gen speed	Hi-fi max plat pitch
	–	–	hrs	–	–	rpm	degrees
Low fidelity	8	–	0.263	0.1000	105389	9.8246	4.4500
High fidelity	–	74	5.10	0.1935	111080	8.9379	4.4606
Multifidelity	755	23	1.96	0.1891	111114	8.9867	4.5088

Table 4 also shows that the multifidelity method finds an optimal answer using 62% less computational expense than the high-fidelity optimization. The one-time cost of linearizing the model across three wind speeds is included for both the low-fidelity and multifidelity computational cost columns. Specifically, this upfront cost requires 944 core-seconds, but then each

function call to the low-fidelity model is quite low. at 0.55 core-seconds. Each function call to the high-fidelity model requires
280 248 core-seconds.

5 Case study: wind power plant layout optimization

Wind power plant layout optimization is the practice of placing wind turbines within a plant to minimize the power production losses caused by wakes from upstream turbines. This is a well-studied and challenging optimization problem due to the inherent multimodality of the design space (Samorani, 2013; Baker et al., 2019; Khan and Rehman, 2013; Stanley and Ning, 2019).
285 Turbine-wake interactions require high-fidelity simulations, including large-eddy simulations, to correctly resolve the highly complex flows within a wind power plant (Fleming et al., 2013; Churchfield et al., 2016); however, the large computational expense of these simulations limits their use in design optimization problems, which has encouraged the development of wind power plant simulation tools that straddle multiple levels of fidelity (Sprague et al., 2020; Réthoré et al., 2014). In this case study, we optimize the layout of turbines using multiple different wake models and resolutions to represent different levels of
290 fidelity.

5.1 Model descriptions and tools used

To more easily study how wind turbine layout and controls affect plant performance using less computational cost, multiple analytic wake models have been developed, including the Jensen (Jensen, 1983), Gaussian (Bastankhah and Porté-Agel, 2014), and Gauss-Curl Hybrid (GCH) (King et al., 2020) models. Listed in order of increasing fidelity, these analytic models
295 capture simplified wake physics and have been verified against high-fidelity simulations and validated against experimental results (King et al., 2020).

In this paper, we use the Jensen and GCH as the low- and high-fidelity wake models, respectively. The Jensen wake model uses a simplistic velocity deficit to represent the wake, and this deficit is summed when wakes interact using the sum-of-squares method (Jensen, 1983). Additionally, the velocity deficit fans out linearly behind the turbine. The wakes from the Jensen model
300 for the initial plant used in this study are shown in Figure 8a. The GCH model modifies the Gaussian model (Bastankhah and Porté-Agel, 2014) by including analytic approximations from the curl model (Martínez-Tossas et al., 2019), which leads to a wake model that better resembles results from high-fidelity simulations. These more complex flow interactions are visible in Figure 8b, which also uses a sum-of-squares method for wake interaction.

These wake models are already integrated into FLOw Redirection and Induction in Steady State (FLORIS) (NREL, 2021a),
305 a controls-oriented wake modeling tool that performs wind power plant simulation and optimization. FLORIS is an open-source tool that provides a common application programming interface for multiple wake models, which allows us to easily investigate different levels of fidelity.

In addition to using different wake models, our low- and high-fidelity models for this problem use different wind roses and wind speed bin resolutions, leading to accuracy and computational differences caused by both fidelity and resolution.
310 The low-fidelity model samples six equally spaced wind directions (60° bins) and five wind speeds from 0–26 m/s, whereas

(a) Jensen wake model	(b) GCH wake model

Figure 8. Simplistic velocity deficits of the Jensen wake model (Jensen, 1983) (a) versus the more accurately resolved flow field of the GCH model (King et al., 2020).

the high-fidelity model samples 18 wind directions ($20°$ bins) and 14 wind speeds from 0–26 m/s. These relatively coarse discretizations were selected so the optimization studies could be easily run on a laptop workstation. Both models use a Weibull distribution for the wind speed frequencies.

5.2 Optimization problem formulation

315 For this study, we optimize the locations of seven wind turbines within an area of 360000 square meters. Additionally, we impose a two-rotor-diameter (2D or 262-unitm) spacing constraint between turbines to create a well-posed optimization problem. We aggregate these turbine-turbine spacing constraints using the Kreisselmeier–Steinhauser functional (Poon and Martins, 2007), which reduces the number of constraints from 21 to 1, producing a less complex optimization problem. This problem formulation leads to 14 design variables, one objective, and one constraint, as shown in Table 5. The wind turbine model is 320 based on the NREL 5-MW reference turbine (Jonkman et al., 2009) and is provided within FLORIS.

Table 5. Optimization problem formulation for the wind power plant layout AEP maximization case.

Category	Name	Quantity	Lower bound	Upper bound	Units
Objective	AEP	1	–	–	GWh
Variables	Turbine x-locations	7	0.0	600.0	m
	Turbine y-locations	7	0.0	600.0	m
Constraints	Turbine spacing	1	262 (2D)	–	m

5.3 Optimization results

As in the first two case studies, we performed single-fidelity and multifidelity optimizations for this plant layout case, with the high-fidelity AEP evaluated at the optimal design from each method shown in Table 6. Each call to the low- and high-fidelity

models took 0.212 second and 7.13 seconds, respectively, meaning that the high-fidelity model is 33.6 times as expensive as the
low-fidelity model to evaluate. Overall, we see that the multifidelity method takes 58% as many core-hours to find an optimal
answer as the high-fidelity method. The multifidelity method resulted in a better layout than the low-fidelity optimization,
however, this AEP value was less than that from the high-fidelity optimization. Examining the physical layouts from the
high-fidelity and multifidelity cases shown in Figure 9, the results do not appear drastically different, although only one wind
direction and speed from the wind rose is shown. The main difference between the two cases lies in the location of the central
turbine, which is farther north in the high-fidelity case. Note that in all cases, the turbine spacing constraint is not active at the
optimal design; thus, the trade-off between the computational savings and the optimality of the obtained design would vary
based on the number of wind turbine locations optimized.

Table 6. Optimization results for the wind power plant layout AEP maximization case.

Design	Lo-fi calls	Hi-fi calls	CPU time	Hi-fi AEP
	–	–	hrs	GWh
Low fidelity	6382	–	0.376	81.855
High fidelity	–	3975	7.87	82.367
Multifidelity	56939	618	4.58	81.972

(a) High-fidelity layout   (b) Multifidelity layout

Figure 9. The optimal layout found by the (a) high-fidelity-only optimization (82.368 GWh AEP) and (b) the multifidelity method
(81.972 GWh AEP).

This wind power plant layout problem presents an interesting case for the multifidelity method due to the highly nonlinear
design space as well as the number of design variables. The corrective function used to correlate the two fidelity levels needs
to be able to capture sharp changes in AEP with respect to changes in turbine location. By not using simple additive or
multiplicative factors, and instead using surrogate corrective functions, as detailed in Section 2.3, we are able to account
for the design space nonlinearities. As the number of design variables increases, however, the number of points needed to
correctly correlate the two fidelities also increases. This trend is not due to the type of corrective function used but is instead

due to the well-known "curse of dimensionality," which dictates that the cost of constructing an accurate representation of a
340 high-dimensional space increases greatly as the number of dimensions increases. These costs are problem-dependent, and this
power plant layout problem is known to be highly nonlinear and high-dimensional, which leads to a relatively large number of
training points to correctly correlate the low- and high-fidelity models.

6 Conclusions

We have shown that multifidelity optimization methods are effective for a variety of wind energy applications to decrease
345 the computational cost needed to find an optimal design. Optimizing using only a low-fidelity model might miss important
physical trends that the high-fidelity and multifidelity approaches will correctly capture. Across three distinct applications—
aerodynamic blade design, controls tuning, and wind power plant layout optimization—we have shown that obtaining an
optimal result requires 38% to 58% of the computational cost compared to high-fidelity optimization, as depicted in Figure 10.
Showing the percentage change for each problem normalizes the computational costs across differences in core-hours and
350 function calls, which may vary depending on system type, parallelilzation, and compiler choices.

Problem	Multifidelity cost relative to hi-fi	
Blade design		0.44
Controls		0.38
Plant layout		0.58

Figure 10. The normalized relative computational cost for multifidelity optimization in each application case compared to high-fidelity
optimization shows that the trust-region method finds an optimal answer in less time.

Although we used a traditional trust-region approach for multifidelity optimization, we offered a new corrective function
technique based on efficient KPLS surrogate models, and we demonstrated its efficacy across three case studies. In this way,
the methods and results presented in this paper should be useful to wind energy researchers who seek optimal designs when
using multiple levels of model fidelity.
355 There are some limitations to the types of design problems for which multifidelity methods are effective. Specifically, there
needs to be an established model fidelity hierarchy with one model known to be of higher fidelity than another. If the accuracy
of the models is unknown, then the trust-region method presented here is ill-posed. Each model used in the multifidelity method
must receive the same inputs and return the same outputs so the corrective function between fidelity levels can be constructed.
Finally, multifidelity methods are less beneficial when there is not a large difference between the computational expense of the
360 models. Many engineering design problems meet these requirements, but special care is needed to select appropriate levels of
model fidelity and to pose a reasonable optimization problem.
Future work could involve more complicated design problems, additional fidelity tiers, or different types of model disci-
plines. As the optimization problems increase in complexity, the best multifidelity strategy might differ, including which type

of corrective function to use or how many correlative design points to use. A series of model fidelities could also be considered, with nested trust regions to conduct the model fidelity management. Last, in this paper we examined multiple disciplines in wind energy systems engineering, but there are additional subsets of model disciplines that could benefit from design exploration through these multifidelity methods.

Acknowledgements. This work was authored by the National 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 U.S. Department of Energy Advanced Research Projects Agency-Energy. 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 U.S. Government purposes.

This 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.

References

Abbas, N., Zalkind, D., Pao, L., and Wright, A.: A Reference Open-Source Controller for Fixed and Floating Offshore Wind Turbines, Wind Energy Science Discussions, 2021, 1–33, https://doi.org/10.5194/wes-2021-19, https://wes.copernicus.org/preprints/wes-2021-19/, 2021.

Abdallah, I., Lataniotis, C., and Sudret, B.: Parametric hierarchical kriging for multi-fidelity aero-servo-elastic simulators—Application to extreme loads on wind turbines, Probabilistic Engineering Mechanics, 55, 67–77, 2019.

Alexandrov, N. M., Dennis, J., Lewis, R. M., and Torczon, V.: A trust-region framework for managing the use of approximation models in optimization, Structural optimization, 15, 16–23, 1998.

Alexandrov, N. M., Lewis, R. M., Gumbert, C. R., Green, L. L., and Newman, P. A.: Approximation and model management in aerodynamic optimization with variable-fidelity models, Journal of Aircraft, 38, 1093–1101, 2001.

Allen, C., Viselli, A., Dagher, H., Goupee, A., Gaertner, E., Abbas, N., Hall, M., and Barter, G.: Definition of the UMaine VolturnUS-S Reference Platform Developed for the IEA Wind 15-Megawatt Offshore Reference Wind Turbine, Tech. Rep. NREL/TP-76773, International Energy Agency, 2020.

Ashuri, T., Zaaijer, M. B., Martins, J. R. R. A., van Bussel, G. J. W., and van Kuik, G. A. M.: Multidisciplinary Design Optimization of Offshore Wind Turbines for Minimum Levelized Cost of Energy, Renewable Energy, 68, 893–905, https://doi.org/10.1016/j.renene.2014.02.045, 2014.

Baker, N. F., Stanley, A. P., Thomas, J. J., Ning, A., and Dykes, K.: Best practices for wake model and optimization algorithm selection in wind farm layout optimization, in: AIAA Scitech 2019 Forum, p. 0540, 2019.

Barlas, T., Ramos-García, N., Pirrung, G. R., and González Horcas, S.: Surrogate based aeroelastic design optimization of tip extensions on a modern 10MW wind turbine, Wind Energy Science Discussions, 2020, 1–20, https://doi.org/10.5194/wes-2020-108, https://wes.copernicus.org/preprints/wes-2020-108/, 2020.

Bastankhah, M. and Porté-Agel, F.: A new analytical model for wind-turbine wakes, Renewable Energy, 70, 116–123, 2014.

Bir, G. S.: User's Guide to MBC3: Multi-Blade Coordinate Transformation Code for 3-Bladed Wind Turbine, Tech. Rep. NREL/TP-500-44327, National Renewable Energy Laboratory, https://www.nrel.gov/docs/fy10osti/44327.pdf, 2010.

Bortolotti, P., Bottasso, C. L., and Croce, A.: Combined preliminary–detailed design of wind turbines, Wind Energy Science, 1, 71–88, https://doi.org/10.5194/wes-1-71-2016, https://wes.copernicus.org/articles/1/71/2016/, 2016.

Bouhlel, M. A., Bartoli, N., Otsmane, A., and Morlier, J.: Improving Kriging Surrogates of High-Dimensional Design Models by Partial Least Squares Dimension Reduction, Structural and Multidisciplinary Optimization, 53, 935–952, https://doi.org/10.1007/s00158-015-1395-9, 2016.

Bouhlel, M. A., Hwang, J. T., Bartoli, N., Lafage, R., Morlier, J., and Martins, J. R. R. A.: A Python surrogate modeling framework with derivatives, Advances in Engineering Software, p. 102662, https://doi.org/https://doi.org/10.1016/j.advengsoft.2019.03.005, 2019.

Churchfield, M., Wang, Q., Scholbrock, A., Herges, T., Mikkelsen, T., and Sjöholm, M.: Using high-fidelity computational fluid dynamics to help design a wind turbine wake measurement experiment, Journal of Physics: Conference Series, 753, 032 009, 2016.

Cressie, N.: Spatial prediction and ordinary kriging, Mathematical geology, 20, 405–421, 1988.

Fischer, G. R., Kipouros, T., and Savill, A. M.: Multi-objective optimisation of horizontal axis wind turbine structure and energy production using aerofoil and blade properties as design variables, Renewable Energy, 62, 506–515, 2014.

Fleming, P., Gebraad, P., van Wingerden, J.-W., Lee, S., Churchfield, M., Scholbrock, A., Michalakes, J., Johnson, K., and Moriarty, P.: SOWFA super-controller: A high-fidelity tool for evaluating wind plant control approaches, Tech. rep., National Renewable Energy Lab.(NREL), Golden, CO (United States), 2013.

Forrester, A. I., Sóbester, A., and Keane, A. J.: Multi-fidelity optimization via surrogate modelling, Proceedings of the royal society a: mathematical, physical and engineering sciences, 463, 3251–3269, 2007.

Fuglsang, P. and Madsen, H. A.: Optimization method for wind turbine rotors, Journal of Wind Engineering and Industrial Aerodynamics, 80, 191–206, 1999.

Gaertner, E., Rinker, J., Sethuraman, L., Zahle, F., Anderson, B., Barter, G., Abbas, N., Meng, F., Bortolotti, P., Skrzypinski, W., Scott, G., Feil, R., Bredmose, H., Dykes, K., Sheilds, M., Allen, C., and Viselli, A.: Definition of the IEA 15-Megawatt Offshore Reference Wind Turbine, Tech. Rep. NREL/TP-75698, International Energy Agency, https://www.nrel.gov/docs/fy20osti/75698.pdf, 2020.

Gavin, H. P.: Frame3DD. Static and Dynamic Structural Analysis of 2D and 3D Frames, http://frame3dd.sourceforge.net/, version 0.20140514+, 2014.

Giguere, P. and Selig, M.: Blade geometry optimization for the design of wind turbine rotors, in: 2000 ASME Wind Energy Symposium, p. 45, 2000.

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.

Jensen, N.: A note on wind generator interaction, no. 2411 in Risø-M, Risø National Laboratory, 1983.

Jonkman, B.: TurbSim User's Guide: Version 1.50, Tech. Rep. NREL/TP-500-46198, National Renewable Energy Laboratory, https://www.nrel.gov/docs/fy09osti/46198.pdf, 2009.

Jonkman, J., Butterfield, S., Musial, W., and Scott, G.: Definition of a 5-MW reference wind turbine for offshore system development, Tech. rep., National Renewable Energy Lab.(NREL), Golden, CO (United States), 2009.

Jonkman, J. M. and Jonkman, B. J.: FAST modularization framework for wind turbine simulation: full-system linearization, Journal of Physics: Conference Series, 753, 082 010, https://doi.org/10.1088/1742-6596/753/8/082010, 2016.

Kennedy, M. C. and O'Hagan, A.: Predicting the output from a complex computer code when fast approximations are available, Biometrika, 87, 1–13, 2000.

Khan, S. A. and Rehman, S.: Iterative non-deterministic algorithms in on-shore wind farm design: A brief survey, Renewable and Sustainable Energy Reviews, 19, 370–384, 2013.

King, J., Fleming, P., King, R., Martínez-Tossas, L. A., Bay, C. J., Mudafort, R., and Simley, E.: Controls-oriented model for secondary effects of wake steering, Wind Energy Science Discussions, pp. 1–22, 2020.

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.

Maki, K., Sbragio, R., and Vlahopoulos, N.: System design of a wind turbine using a multi-level optimization approach, Renewable Energy, 43, 101–110, 2012.

March, A. and Willcox, K.: Provably convergent multifidelity optimization algorithm not requiring high-fidelity derivatives, AIAA journal, 50, 1079–1089, 2012.

Martínez-Tossas, L. A., Annoni, J., Fleming, P. A., and Churchfield, M. J.: The aerodynamics of the curled wake: a simplified model in view of flow control, Wind Energy Science, 4, 127–138, 2019.

450   McWilliam, M. K., Zahle, F., Pavese, C., and Blasques, J. P.: Multi-fidelity optimization of horizontal axis wind turbines, in: 35th Wind Energy Symposium, p. 1846, 2017.

Moriarty, P. J. and Hansen, A.: AeroDyn Theory Manual, Tech. Rep. NREL/TP-500-36881, National Renewable Energy Laboratory, https://www.nrel.gov/docs/fy05osti/36881.pdf, 2005.

Ning, A. and Petch, D.: Integrated design of downwind land-based wind turbines using analytic gradients, Wind Energy, 19, 2137–2152,
455   https://doi.org/10.1002/we.1972, 2016.

Ning, S. A.: A simple solution method for the blade element momentum equations with guaranteed convergence, Wind Energy, 17, 1327–1345, https://doi.org/10.1002/we.1636, 2014.

Ning, S. A., Damiani, R., and Moriarty, P. J.: Objectives and constraints for wind turbine optimization, Journal of Solar Energy Engineering, 136, 2014.

460   NREL: ROSCO. Version 2.1.1, https://github.com/NREL/rosco, 2020.

NREL: FLORIS. Version 2.2.5, https://github.com/NREL/floris, 2021a.

NREL: OpenFAST, v2.5.0, https://github.com/OpenFAST/openfast, 2021b.

NREL: WEIS, v1.0, https://github.com/WISDEM/WEIS, 2021c.

NREL: WISDEM, v3.2.0, https://github.com/WISDEM/WISDEM, 2021d.

465   Park, J. and Law, K. H.: A Bayesian optimization approach for wind farm power maximization, in: Smart Sensor Phenomena, Technology, Networks, and Systems Integration 2015, vol. 9436, p. 943608, International Society for Optics and Photonics, 2015.

Peherstorfer, B., Willcox, K., and Gunzburger, M.: Survey of Multifidelity Methods in Uncertainty Propagation, Inference, and Optimization, SIAM Review, 60, 550–591, https://doi.org/10.1137/16M1082469, 2018.

Poon, N. M. K. and Martins, J. R. R. A.: An Adaptive Approach to Constraint Aggregation Using Adjoint Sensitivity Analysis, Structural
470   and Multidisciplinary Optimization, 34, 61–73, https://doi.org/10.1007/s00158-006-0061-7, 2007.

Pourrajabian, A., Afshar, P. A. N., Ahmadizadeh, M., and Wood, D.: Aero-structural design and optimization of a small wind turbine blade, Renewable Energy, 87, 837–848, 2016.

Quick, J., Hamlington, P. E., King, R., and Sprague, M. A.: Multifidelity uncertainty quantification with applications in wind turbine aero-dynamics, in: AIAA Scitech 2019 Forum, p. 0542, 2019.

475   Rahbari, O., Vafaeipour, M., Fazelpour, F., Feidt, M., and Rosen, M. A.: Towards realistic designs of wind farm layouts: Application of a novel placement selector approach, Energy Conversion and Management, 81, 242–254, 2014.

Réthoré, P.-E., Fuglsang, P., Larsen, G. C., Buhl, T., Larsen, T. J., and Madsen, H. A.: TOPFARM: Multi-fidelity optimization of wind farms, Wind Energy, 17, 1797–1816, 2014.

Robinson, T. D., Eldred, M. S., Willcox, K. E., and Haimes, R.: Surrogate-Based Optimization Using Multifidelity Models with Variable
480   Parameterization and Corrected Space Mapping, AIAA Journal, 46, 2814–2822, https://doi.org/10.2514/1.36043, 2008.

Samorani, M.: The wind farm layout optimization problem, in: Handbook of wind power systems, pp. 21–38, Springer, 2013.

Shaler, K., Branlard, E., and Platt, A.: OLAF User's Guide and Theory Manual, Tech. Rep. NREL/TP-5000-75959, National Renewable Energy Laboratory, https://www.nrel.gov/docs/fy20osti/75959.pdf, 2020.

Sprague, M. A., Ananthan, S., Vijayakumar, G., and Robinson, M.: ExaWind: A multifidelity modeling and simulation environment for wind
485   energy, Journal of Physics: Conference Series, 1452, 012 071, 2020.

Stanley, A. P. and Ning, A.: Massive simplification of the wind farm layout optimization problem, Wind Energy Science, 4, 663–676, 2019.

Yu, X., Zhang, W., Zang, H., and Yang, H.: Wind power interval forecasting based on confidence interval optimization, Energies, 11, 3336, 2018.