Model-based closed-loop wind farm control for power maximization using Bayesian Optimization: a large eddy simulation study

Bart M. Doekemeijer[†], Daan C. van der Hoek[†] and Jan-Willem van Wingerden[†]

Abstract—Modern wind farm control (WFC) methods in the literature typically rely on a surrogate model of the farm dynamics that is computationally inexpensive to enable real-time computations. As it is very difficult to model all the relevant wind farm dynamics accurately, a closed-loop approach is a prerequisite for reliable WFC. As one of the few in its field, this paper showcases a closed-loop wind farm control solution, which leverages a steady-state surrogate model and Bayesian Optimization to maximize the wind-farm-wide power production. The estimated quantities are the time-averaged ambient wind direction, wind speed and turbulence intensity. This solution is evaluated for a wind farm with nine 10 MW wind turbines in large-eddy simulation, showing a time-averaged power gain of 4.4%. This is the first WFC algorithm that is tested for wind turbines of such scale in high fidelity.

I. INTRODUCTION

As wind turbines extract energy from the wind, they shed a slower, more turbulent pocket of air behind them, called the “wake”. The number of wind turbines placed close to one another continues to increase for financial reasons, and the resulting wake effects play a prominent role in the performance of such wind farms [1]. Specifically, a wind turbine operating in another turbine’s wake extracts less energy from the air flow (e.g., Barthelmie et al. [2] estimated the power losses due to wake interaction to be between 10% and 23% for a number of operational offshore wind farms) and typically experiences more structural degradation due to the increased turbulence in the air [3]. Clearly, wake interaction has a negative impact on the financial feasibility of energy from wind.

The field of wind farm control attempts to address this issue by the coordinated control of (a subset of) the wind turbines inside a farm. One solution proposed in the literature is by derating upstream turbines (purposely extracting less energy than possible) to reduce wake formation downstream. This methodology has led to mixed results in the literature [1]. A more promising method is the steering of the wake by purposely misaligning the rotor plane with the inflow. While this leads to a loss in power capture of the misaligned turbine, significant gains can be achieved downstream, leading to improvements of up to 15% in power capture situationally (e.g., [4], [5]).

Typically, wind farm control algorithms proposed in the literature rely on a simplified surrogate model of the wind

turbines and wake interactions [1]. While model-free methods exist, their long convergence times and the relatively high-frequent changes in ambient conditions significantly complicate such methods. Model-free methods are not further explored in this paper. The success of model-based methods heavily relies on the accuracy of the surrogate model used. Due to the complicated dynamics of the air flow and the turbine structure on a large range of spatial and temporal scales, the development of an accurate and computationally efficient surrogate model is a nontrivial research topic. At the time of writing, no model exists which is consistently accurate over the various operating conditions that a wind farm experiences annually [1]. This problem therefore requires a closed-loop approach, in which the surrogate model is calibrated in real time. The framework is depicted in Fig. 1.

The literature on closed-loop model-based wind farm control is scarce, and there is only a handful of solutions that have been tested in high-fidelity simulation or through experiments (e.g., [6]). Furthermore, existing literature is limited to wind farms with 5 MW turbines or smaller. Hence, this paper may provide some insight into the effect of wind farm control on larger-scale wind turbines.

In this work, a novel closed-loop model-based wind farm control solution will be synthesized and demonstrated on a 3 by 3 wind farm through high-fidelity simulations. The contributions of this work specifically are:

- With respect to previous work [6], a more sophisticated optimization methodology is used for model adaptation and control setpoint optimization, leveraging Gaussian Process (GP) techniques. This should demand fewer function evaluations to approach the global optimum.¹
- Furthermore, the evaluated wind farm consists of nine future-scale 10 MW turbines developed by the Technical University of Denmark (DTU) [7]. This is the first paper in its field to showcase wind farm control for such large-scale turbines in a high-fidelity environment. This will provide insight into the scalability of wind farm control algorithms with respect to turbine size.
- Finally, compared to previous work [6], the high-fidelity simulations now involve higher spatial and temporal resolutions and an actuator line approach to approximate the turbine rotors, leading to more reliable simulation results.

[†]Bart Doekemeijer, Daan van der Hoek and Jan-Willem van Wingerden are with the Delft Center for Systems and Control (DCSC), Faculty of Mechanical, Maritime and Materials Engineering (3mE), Delft University of Technology, The Netherlands {b.m.doekemeijer; d.c.vanderhoek; j.w.vanwingerden}@tudelft.nl

¹Furthermore, GP techniques provide insight into the probability distribution of the optimal arguments of the to-be-optimized function, which is planned to be used to increase robustness of the controller in future work.

Fig. 1. Closed-loop model-based wind farm control framework

The structure of the paper is as follows. The surrogate model used for controller synthesis is described in Section II. The high-fidelity model used for simulation is described in Section III. Then, the closed-loop controller is synthesized in Section IV. Simulation results are shown and analyzed in Section V. Finally, the paper is concluded in Section VI.

II. SURROGATE MODEL

As mentioned previously, a simplified yet accurate surrogate model of the wind farm is an invaluable component of the closed-loop model-based wind farm control solution. The mathematical model used in this work is largely similar to that described in Annoni et al. [8], referred to interchangeably as the “FLOW Redirection and Induction in Steady-state” (FLORIS) model in this paper. Specifically, FLORIS consists of four components:

- 1) a wake model for a single turbine capturing the effects of turbine derating and yaw misalignment,
- 2) a wake summation model that combines the effect of multiple wakes,
- 3) a submodel that captures the turbine-induced turbulence in the flow, and
- 4) a submodel that determines the power captured by a turbine under a certain inflow.

Recently, FLORIS has evolved into a modular software package containing various submodel definitions from the literature [9], [10].

Firstly, in this work, the single wake submodel is largely derived from Bastankhah et al. [11], displayed in Fig. 2. In this figure, U_∞ is the freestream wind speed, γ is the turbine misalignment (yaw) angle, and δ_f is the lateral wake deflection angle due to the rotation of the rotor and due to the yaw misalignment. The model consists of two parts: a near-wake region called the “potential core” with length x_0 in which the wind speed is uniform, and a far-wake region in which the wind speed deficit is shaped as a two-dimensional Gaussian with standard deviations σ_y and σ_z . The mathematical derivation of this model is beyond the scope of this work. The reader is referred to the article by Bastankhah et al. [11] for more information.

Fig. 2. A schematic top-view drawing of the single wake submodel, taken and modified from [11].

Secondly, the wake summation model of Katic et al. [12] is followed to combine wakes. Thirdly, turbine-induced turbulence is modeled according to Crespo and Hernandez [13]. Fourthly, the submodel that captures the turbine power production under a certain inflow relies on an abstraction from an aero-elastic turbine simulation code from the National Renewable Energy Laboratory (NREL) called OpenFAST [14]. OpenFAST was used to generate a database of power and thrust coefficients for various inflow wind speeds, yaw misalignment angles and turbine derating settings for the DTU 10 MW wind turbine. Within FLORIS, the wind speed on the rotor plane is spatially averaged and then used for interpolation in this three-dimensional database.

Finally, these submodels come with their set of tuning parameters that are used to calculate, e.g., δ_f , σ_y and σ_z . In this work, the model parameters are taken from Table I of Doekemeijer et al. [6]. In that publication, the model parameters were tuned for the NREL 5 MW turbine and should theoretically be independent of turbine size.

III. HIGH-FIDELITY MODEL

The synthesized wind farm control solution in this work will be validated through high-fidelity simulation using the Simulator fOr Wind Farm Applications (SOWFA), a large-eddy simulation model for wind farms from NREL [15], [16].

In this work, the turbine rotors are parametrized using the actuator line model [17]. SOWFA then solves the three-

dimensional, unfiltered, unsteady Navier-Stokes equations over a discretized domain, including Coriolis effects and geostrophic forcing terms. SOWFA has been used on multiple occasions for surrogate model calibration, model validation, and wind farm controller verification [1].

IV. CLOSED-LOOP CONTROLLER SYNTHESIS

The surrogate model described in Section II is adapted to the actual conditions inside the wind farm periodically, after which the turbine control setpoints are optimized for a prescribed objective. Each step is described next.

A. Real-time model adaptation

As the optimal control settings are highly sensitive to the ambient conditions, most importantly the wind direction, it is important to estimate the actual conditions inside the farm. This is done in two steps.

Firstly, the wind direction is estimated following the approach described in Bertele et al. [18]. Specifically, they use blade load sensors to reconstruct the inflow profile at the wind turbine. For example, for an inflow wind speed of 8 m/s, the wind direction can be estimated with a standard deviation of 6° . For simplicity, the algorithm is not implemented, but rather an estimate for the wind direction is assumed to be available under the specified uncertainty.

Secondly, the wind speed and turbulence intensity are estimated through a model-inversion approach, in which the error between the power predicted by FLORIS and the time-averaged measured power is minimized. Such power measurements are already available in modern wind turbines. The to-be-minimized cost function is:

$$\Xi_{\text{opt}} = \arg \min_{\Xi} \sum_i^{N_t} (P_{\text{SOWFA}}^i - P_{\text{FLORIS}}^i(\Xi)), \quad (1)$$

with N_t the number of turbines, $\Xi = [I_\infty, U_\infty]$ a vector with the to-be-estimated variables, and I_∞ the freestream turbulence intensity.

B. Control setpoint optimization

In this work, the aim is to maximize the collective power production of the wind farm. After model adaptation, the surrogate model is assumed to accurately capture the current conditions inside the farm. A robust optimization approach is then followed based on the work from Rott et al. [19], in which the yaw angles are optimized for a probability distribution of wind directions, rather than one deterministic wind direction. The to-be-maximized cost function is:

$$\vec{\gamma}_{\text{opt}} = \arg \max_{\vec{\gamma}} \int_{-\pi}^{\pi} \rho(\phi) \sum_i^{N_t} P_{\text{FLORIS}}^i(\phi, \gamma_i) d\phi, \quad (2)$$

with $\vec{\gamma} = [\gamma_1 \ \gamma_2 \ \dots \ \gamma_{N_t}]$, and ρ a probability distribution of the wind direction ϕ . Further, in the optimization, γ_i is bounded to $[-30, 30]$ degrees to manage the increase in structural loads compared to greedy operation. To mitigate computational cost, ρ is discretized at 5 points, as exemplified in Fig. 3. A more refined discretization of ρ would result

Fig. 3. An example of the probability distribution of the wind direction. The red circles show a possible discretization of the distribution.

in a proportional increase in computational cost and a more accurate posterior probability distribution of the farm's power production. However, this is not guaranteed to improve the algorithm due to model mismatches.

C. Optimization methodology

The optimization methodology adopted in this paper for both model adaptation and control setpoint optimization consists of Bayesian Optimization (BO) leveraging a Gaussian Process (GP) as a surrogate model of the cost function [20]. Using a GP to model the model adaptation cost function provides a straightforward extension of the robust control setting approach, as it is able to estimate the wind direction in a probabilistic manner. For simplicity, the focus in this section is on the estimation part, but $\mathbf{x}$ can correspond to either the to-be-estimated ambient conditions or the to-be-optimized yaw angles.

1) *Gaussian processes*: The cost function given in Eq. 1 (resp. Eq. 2 for the control optimization) can be modelled as a GP by assuming that the noisy observations of a cost function $\mathbf{J} = [J^1 \ J^2 \ \dots \ J^n]$ belong to a multivariate Gaussian distribution:

$$\begin{bmatrix} \mathbf{J} \\ \mathbf{J}^* \end{bmatrix} \sim \mathcal{N} \left(\mathbf{0}, \begin{bmatrix} K(\mathbf{x}, \mathbf{x}) + \sigma_n^2 \mathbf{I} & K(\mathbf{x}, \mathbf{x}^*) \\ K(\mathbf{x}^*, \mathbf{x}) & K(\mathbf{x}^*, \mathbf{x}^*) \end{bmatrix} \right), \quad (3)$$

where $\mathbf{x} = [x^1 \ x^2 \ \dots \ x^n]$ is a matrix containing sets of ambient conditions (resp. yaw angles) that we previously evaluated in FLORIS. Furthermore, $\mathbf{J}^*$ indicates the estimation error at the trial inputs $\mathbf{x}^*$ (i.e., sets of possible ambient conditions), K refers to the covariance function and the variance of the measurement noise is given by σ_n^2 .

In this paper, the squared exponential covariance function, which is useful for smooth functions in general, is used to quantify the similarity between two input vectors $\mathbf{x}$ and $\mathbf{x}^*$, i.e.:

$$K(\mathbf{x}, \mathbf{x}^*) = \lambda_f^2 \cdot \exp \left(-\frac{1}{2} (\mathbf{x} - \mathbf{x}^*)^T \boldsymbol{\Lambda}_x^{-1} (\mathbf{x} - \mathbf{x}^*) \right). \quad (4)$$

The variables λ_f , $\boldsymbol{\Lambda}_x$ and σ_n are the so-called hyperparameters and they are updated during each iteration of the optimization algorithm [20].

Once we have tested a new set of ambient conditions in FLORIS, a prediction can be made of the cost function at the trial points $\mathbf{x}^*$ by conditioning the prior distribution from

Eq. 3 to obtain the posterior distribution according to:

$$\boldsymbol{\mu}^* = K(\mathbf{x}^*, \mathbf{x}) (K(\mathbf{x}, \mathbf{x}) + \sigma_n^2 \mathbf{I})^{-1} \mathbf{J}, \quad (5)$$

$$\boldsymbol{\Sigma}^* = K(\mathbf{x}^*, \mathbf{x}^*) - K(\mathbf{x}^*, \mathbf{x}) (K(\mathbf{x}, \mathbf{x}) + \sigma_n^2 \mathbf{I})^{-1} K(\mathbf{x}, \mathbf{x}^*), \quad (6)$$

where $\boldsymbol{\mu}^*$ and $\boldsymbol{\Sigma}^*$ are the mean and variance of the cost function at the trial points $\mathbf{x}^*$, respectively.

2) *Bayesian optimization*: Using the GP framework, a statistical model of the cost function is obtained, which in turn can be used to find the optimally estimated ambient conditions. The next ambient conditions to be tested in the estimation algorithm is found by optimizing an acquisition function. In this paper, the Upper Confidence Bound (UCB) is selected as acquisition function and is optimized accordingly [21]:

$$\mathbf{x}^{n+1} = \arg \max_{\mathbf{x}} (\boldsymbol{\mu}^* + \kappa \sqrt{\boldsymbol{\Sigma}^*}), \quad (7)$$

where κ is a user-defined variable and is set to $\kappa = 2$ in order to allow for some exploration of the cost function. Using alternative acquisition functions might lead to improved results or faster convergence, but such a comparison is outside the scope of this paper.

The BO algorithm can be run for a specified number of iterations or until a stopping criterion is reached. In this paper, n depicts the current iteration number and was set to a maximum of $N = 50$ and $N = 200$ in order to converge to the optimal settings for the model adaptation and control setpoint optimization, respectively. The optimal input variables were those that resulted in the optimal value for each cost function.

V. SIMULATION RESULTS

Theoretical proofs of wind farm control algorithms are extremely difficult due to the uncertainties and hard-to-model temporal and spatial dynamics involved. Rather, high-fidelity simulations are performed in pursuit of empirical validation.

A. Simulation setup

The proposed wind farm control solution is tested in a high-fidelity simulation for a farm with nine DTU 10 MW turbines. The simulation settings are displayed in Table I. First, a simulation in which the turbines are operated under

TABLE I
SIMULATION SETTINGS FOR SOWFA

Variable	Value
Timestep	0.20 s
Rotor model	Actuator line model [17]
Cell size (near rotor)	2.5 m $\times$ 2.5 m $\times$ 2.5 m
Cell size (outer region)	10.0 m $\times$ 10.0 m $\times$ 10.0 m
Blade epsilon, ϵ	5.0 m
Turbine layout	3 $\times$ 3 turbines at 5D $\times$ 3D spacing
Freestream wind speed, U_∞	8.0 m/s
Freestream turbulence intensity, I_∞	5.0 %
Freestream wind direction, ϕ	0.0 rad

traditional “greedy” control is performed, where all turbines

are aligned with the wind. Secondly, the closed-loop controller is employed to periodically optimize the yaw angles of the turbines inside the wind farm. A top view of one wind farm simulation showing the flow at turbine hub height is displayed in Fig. 5.

B. Simulation results and analysis

First, the effect of the synthesized wind farm control solution on the farm-wide power production is investigated. A timeseries of the collective power capture is shown in Fig. 4, and time-averaged quantities for various time windows are shown in Table II.

Fig. 4. Comparison of timeseries for the power capture of the wind farm for greedy operation and for optimized operation

TABLE II
COMPARISON OF GREEDY AND THE CLOSED-LOOP CONTROLLER CASES

Time window	$P_{\text{farm}}(\gamma = \tilde{\gamma}_{\text{greedy}})$	$P_{\text{farm}}(\gamma = \tilde{\gamma}_{\text{opt, CL}}^{\text{rob}})$
0-600 s	26.1 MW	26.1 MW
600-900 s	21.5 MW	21.2 MW
900-1200 s	21.4 MW	22.1 MW
1200-1500 s	21.3 MW	22.3 MW
1500-1800 s	21.4 MW	25.0 MW
0-2000 s	22.9 MW	23.9 MW

During the first 600 s, the power production is identical between the greedy and the controller case, as both simulations operate under greedy control. At $t = 600$ s, the control setpoints are optimized, resulting in a yaw misalignment in the order of -20° for the upstream turbines, and of -5° to -10° for the second row of turbines. This wake deflection initially leads to a loss in power of 0.3 MW at the yawed turbines at $t = 600 - 900$ s (Table II), but as the reduced wakes propagate downstream, leads to a wind-farm-wide gain of 0.7 MW at $t = 900 - 1200$ s. A second and third control setpoint optimization occur at $t = 1200$ s and $t = 1800$ s, respectively, in which the optimization leads the turbines to be yawed in the other direction. Table III shows the optimized control setpoints. The horizontal flowfield at hub height is shown in Fig. 5. This sign change in yaw angles can be explained by inspecting the real-time model adaption results, shown in Table IV.

In this table, it can be seen that the estimation is generally consistent and accurate, only slightly deviating from the true $I_\infty = 5.0\%$ and $U_\infty = 8.0$ m/s. The change in the yaw angles is due to the change in sign of the estimated wind direction, and the robust optimization approach appears to be insufficient

TABLE III

CONTROL SETPOINTS OF THE CLOSED-LOOP CONTROLLER FOR TURBINES 1-6. TURBINES 7-9 FOLLOW GREEDY CONTROL; $\gamma_{7-9} = 0^\circ$.

Time window	Yaw angles of turbines 1 to 6, $\gamma_{1-6}^{\text{opt, CL}} (^\circ)$
0-600 s	[0.0, 0.0, 0.0, 0.0, 0.0, 0.0] (greedy)
600-1200 s	[-17.2, -18.1, -17.9, -9.2, -4.4, -8.9]
1200-1800 s	[18.2, 17.1, 18.1, 8.1, 7.1, 10.6]
1800-2000 s	[19.0, 18.5, 19.2, 16.1, 17.0, 16.1]

Fig. 5. Top-view at hub-height in m/s at $t = 1750$ s

to deal with such variations. However, the total wind-farm-wide power capture increases with 4.4% over the 2000 s of simulation, and situationally increases with 16.8% for the 1500 – 1800 s timeslot.

The Bayesian optimization algorithm used in this paper takes $7 \cdot 10^1$ s on a single Intel i7-6600U core for a single estimation of the ambient conditions, in comparison to $2 \cdot 10^2$ s for the optimization algorithm in previous work [6], at a negligible loss of performance. A stronger CPU would further reduce the computational cost, as the problem can be parallelized. However, it is uncertain how fast this estimation has to be, as the optimal frequency of yaw setpoint updates is still an open question in the literature.

Furthermore, note that the optimized yaw angles shown in Table III are relatively conservative, never exceeding a misalignment of 20 degrees. This can be explained by comparing the impact of yaw misalignment on a turbine's power production as modeled in SOWFA, and as modeled in FLORIS. See Fig. 6. Here, it is seen that FLORIS overestimates the loss

Fig. 6. Normalized power as a function of the yaw misalignment for a single turbine in freestream flow

TABLE IV

ESTIMATION PERFORMANCE OF CLOSED-LOOP CONTROLLER

Time window	Estimated ambient conditions
300-600 s	$\Xi_{\text{opt}} = [6.5\% \quad 8.0 \text{ m/s}], \phi = 2.3^\circ$
900-1200 s	$\Xi_{\text{opt}} = [7.0\% \quad 8.1 \text{ m/s}], \phi = -0.7^\circ$
1500-1800 s	$\Xi_{\text{opt}} = [4.0\% \quad 7.9 \text{ m/s}], \phi = -2.4^\circ$

in power production due to a yaw misalignment. Therefore, the yaw angles at $t = 600$ s may be lower than what is truly optimal. While good results are achieved, these observations stress the importance of appropriate surrogate model tuning prior to controller synthesis [6].

In comparison to a near-identical simulation with 5MW turbines [6], several things can be noted. Firstly, the optimal yaw angles for the upstream turbines are close to 20° for both simulations. Secondly, the initial loss in power production due to a yaw misalignment is compensated by an increase at downstream turbines (under the assumption of negligible changes in ϕ).

C. Discussion

While the results appear promising, there are a number of important limitations to this work. Firstly, the inflow in SOWFA is quasi-static, with one mean wind direction, speed, and turbulence intensity. In reality, these conditions change continuously. Secondly, the measurements and yaw settings are idealized. Thirdly, the low-frequent changes in inflow also gives rise to the question on how often to go through the model adaptation and control setpoint optimization cycle – whether it is every 20 seconds, every minute, or every 10 minutes. This remains an open question in the literature.

In addition, in this work, a comparison was made between a wind farm control (WFC) algorithm and traditional wind farm operation, rather than comparing this algorithm with other WFC algorithms in the literature. This is particularly difficult, as such algorithms must be compared under identical situations (e.g., wind turbine, wind farm layout, test environment). This is out of the scope of this paper.

Also, wind turbines are currently not designed to purposely operate under a yaw misalignment. The effect of the WFC algorithm proposed in this work on the fatigue loads of the turbine structure will be investigated in future work.

Furthermore, an advantage of the GP framework is that it can be extended to obtain an approximation of the confidence bounds [22] on the estimated wind direction. This can directly be translated into the control setpoint optimization, as ρ in Eq. 2, to further promote robustness. This will be investigated in future work.

VI. CONCLUSIONS

This paper is the first of its field to demonstrate a closed-loop model-based wind farm control solution on wind farm with future-scale, 10 MW wind turbines. This control solution leverages Gaussian Process techniques for computational efficiency, and a robust optimization technique to deal with uncertainties in the estimated wind

direction. Without a priori tuning of the surrogate wind farm model, a time-averaged power increase of 4.4% was noted in high-fidelity simulation when compared to traditional, greedy operation. While results are positive, the mismatch between the simulation model and the surrogate model leads one to suspect that even higher gains can be achieved.

ACKNOWLEDGMENT

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 727477.

SOFTWARE AVAILABILITY

All software presented in this work is open-source. SOWFA is actively maintained by NREL [16]. This repository also includes the communication interface developed jointly by NREL and TU Delft that exchanges information to a MATLAB/Python-based WFC algorithm. The surrogate model is actively developed by TU Delft [9].

REFERENCES

- [1] S. Boersma, B. M. Doekemeijer, P. M. O. Gebraad, P. A. Fleming, J. Annoni, A. K. Scholbrock, J. A. Frederik, and J. W. van Wingerden, "A tutorial on control-oriented modeling and control of wind farms," *American Control Conference (ACC)*, pp. 1–18, 2017.
- [2] R. J. Barthelmie, S. C. Pryor, S. T. Frandsen, K. S. Hansen, J. G. Schepers, K. Rados, W. Schlez, A. Neubert, L. E. Jensen, and S. Neckelmann, "Quantifying the impact of wind turbine wakes on power output at offshore wind farms," *Journal of Atmospheric and Oceanic Technology*, vol. 27, no. 8, pp. 1302–1317, 2010.
- [3] S. Kanev and F. Savenije, "Active wake control: Loads trends," Energieonderzoek Centrum Nederland (ECN), ECN-E-15-004, 2015.
- [4] F. Campagnolo, V. Petrovi, C. L. Bottasso, and A. Croce, "Wind tunnel testing of wake control strategies," *Proceedings of the American Control Conference (ACC)*, pp. 513–518, 2016.
- [5] P. A. Fleming, J. Annoni, A. K. Scholbrock, E. Quon, S. Dana, S. Schreck, S. Raach, F. Haizmann, and D. Schlipf, "Full-scale field test of wake steering," *Journal of Physics: Conference Series*, vol. 854, no. 1, 2017.
- [6] B. M. Doekemeijer, P. A. Fleming, and J. W. van Wingerden, "A tutorial on the synthesis and validation of a closed-loop wind farm controller using a steady-state surrogate model," *American Control Conference (ACC)*, 2019.
- [7] C. Bak, F. Zahle, R. Bitsche, T. Kim, A. Yde, L. C. Henriksen, M. H. Hansen, J. P. A. Blasques, M. Gaunaa, and A. Natarajan, "The DTU 10-MW reference wind turbine," in *Danish Wind Power Research 2013*, 2013.
- [8] J. Annoni, P. Fleming, A. Scholbrock, J. Roadman, S. Dana, C. Adcock, F. Porte-Agel, S. Raach, F. Haizmann, and D. Schlipf, "Analysis of control-oriented wake modeling tools using lidar field results," *Wind Energy Science Discussions*, vol. 2018, pp. 1–17, 2018.
- [9] B. M. Doekemeijer and R. Storm, *FLORIS repository*, https://github.com/TUdelft-DataDrivenControl/FLORISSE_M, 2018.
- [10] R. M. Mudafort, J. R. Annoni, P. A. Fleming, and K. Fleming, *WISDEM/FLORIS Github repository*, <https://github.com/WISDEM/FLORIS>, 2018.
- [11] M. Bastankhah and F. Porté-Agel, "Experimental and theoretical study of wind turbine wakes in yawed conditions," *Journal of Fluid Mechanics*, vol. 806, pp. 506–541, 2016.
- [12] I. Katic, J. Højstrup, and N. Jensen, "A simple model for cluster efficiency," in *EWEC'86. Proceedings. Vol. 1*, W. Palz and E. Sesto, Eds. A. Raguzzi, 1987, pp. 407–410.
- [13] A. Crespo and J. Hernandez, "Turbulence characteristics in wind-turbine wakes," *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 61, no. 1, pp. 71–85, 1996.
- [14] National Renewable Energy Laboratory (NREL), *OpenFAST Github repository*, <https://github.com/OpenFAST/openfast>, 2018.
- [15] M. J. Churchfield, S. Lee, J. Michalakes, and P. J. Moriarty, "A numerical study of the effects of atmospheric and wake turbulence on wind turbine dynamics," *Journal of Turbulence*, vol. 13, no. 14, pp. 1–32, 2012.
- [16] National Renewable Energy Laboratory (NREL), *SOWFA Github repository*, <https://github.com/NREL/SOWFA>, 2018.
- [17] J. N. Sørensen and W. Z. Shen, "Numerical modeling of wind turbine wakes," *Journal of Fluid Mechanics*, vol. 124, no. 2, pp. 393–399, 2002.
- [18] M. Bertele, C. L. Bottasso, S. Cacciola, F. Daher Adegas, and S. Delport, "Wind inflow observation from load harmonics," *Wind Energy Science*, vol. 2, no. 2, pp. 615–640, 2017.
- [19] A. Rott, B. M. Doekemeijer, J. K. Seifert, J. W. van Wingerden, and M. Kühn, "Robust active wake control in consideration of wind direction variability and uncertainty," *Wind Energy Science*, vol. 3, no. 2, pp. 869–882, 2018.
- [20] C. E. Rasmussen and C. K. Williams, *Gaussian process for machine learning*. MIT press, 2006.
- [21] N. Srinivas, A. Krause, S. M. Kakade, and M. Seeger, "Gaussian process optimization in the bandit setting: No regret and experimental design," in *International Conference on Machine Learning*, 2010.
- [22] H. Bijl, T. B. Schon, J. W. van Wingerden, and M. Verhaegen, "A sequential Monte Carlo approach to Thompson sampling for Bayesian optimization," 2016.