

## Estimation of the Ambient Wind Field from Wind Turbine Measurements Using Gaussian Process Regression

Van Der Hoek, Daan; Sinner, Michael; Simley, Eric; Pao, Lucy; Van Wingerden, Jan Willem

**DOI**

[10.23919/ACC50511.2021.9483088](https://doi.org/10.23919/ACC50511.2021.9483088)

**Publication date**

2021

**Document Version**

Final published version

**Published in**

Proceedings of the American Control Conference, ACC 2021

**Citation (APA)**

Van Der Hoek, D., Sinner, M., Simley, E., Pao, L., & Van Wingerden, J. W. (2021). Estimation of the Ambient Wind Field from Wind Turbine Measurements Using Gaussian Process Regression. In *Proceedings of the American Control Conference, ACC 2021* (pp. 558-563). IEEE.  
<https://doi.org/10.23919/ACC50511.2021.9483088>

**Important note**

To cite this publication, please use the final published version (if applicable).  
Please check the document version above.

**Copyright**

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

**Takedown policy**

Please contact us and provide details if you believe this document breaches copyrights.  
We will remove access to the work immediately and investigate your claim.

# Estimation of the Ambient Wind Field From Wind Turbine Measurements Using Gaussian Process Regression

Daan van der Hoek<sup>1</sup>, Michael Sinner<sup>2,3</sup>, Eric Simley<sup>3</sup>, Lucy Pao<sup>2</sup>, and Jan-Willem van Wingerden<sup>1</sup>

**Abstract**—In the search for a lower leveled cost of wind energy, one approach is to increase the accuracy of wind turbine measurements such as wind speed and wind direction. The sensors available on wind turbines are susceptible to local turbulence and measurement bias, which can result in suboptimal turbine performance. As an alternative, recent research has considered using the sensor measurements in a coordinated manner. With such a cooperative approach, the local wind conditions can be estimated more accurately and reliably without the need for additional measurement equipment. In this paper, a novel wind field estimation approach is presented that estimates the local wind conditions based on turbine measurements using Gaussian processes. We show that the estimation framework is able to improve the accuracy of the wind direction estimate both in an offline and online manner, as well as identify possible biases in the sensors and reduce unnecessary wind turbine yaw activity.

## I. INTRODUCTION

As the drive to produce renewable electricity at lower and lower cost continues, engineers have expanded from controlling single wind turbines independently to controlling the wind farm, or “wind plant”, as a whole. Simulation studies [1], experiments [2], and, more recently, field campaigns [3, 4, 5], have shown that, by coordinating the efforts of turbines within a farm, the energy produced across the farm may be increased compared to the traditional “greedy” case, in which each turbine aims to maximize its own power output. Similarly, cooperating turbines may be more flexible in terms of following power reference signals for grid balancing purposes [6, 7, 8].

Such coordinated control techniques require more in-depth information about the state of the wind field impacting the farm and the interaction between turbine wakes than that required for individual turbine control [9]. Recent efforts have been made to model these complex interactions both at high fidelity [10] and using simplified models that may be used for real-time control and optimization [11, 12].

In order for these models to be useful for online control, accurate measurements of the wind field must be made available. An option is to use extra sensors, such as lidars and meteorological towers, to sample the wind field throughout the farm. On the other hand, the wind turbines themselves are already highly instrumented and capable of providing

point measurements of the wind speed and wind direction, among other quantities, at the turbine locations. These signals have long been used at the turbine locally for independent turbine-level control. A recent branch of research [13, 14] has considered reconstructing relevant characteristics of the wind field from the noisy turbine measurements. Such methods are particularly useful for estimating the underlying wind direction from turbine measurements, because the wind direction measurement plays a critical role in turbine yaw control and may therefore be important for wake steering control [15]. Annoni et al. used a nonparametric, distributed optimization technique for estimating the wind directions at the turbine locations [13], whereas Sinner et al. used a Kalman filtering approach to estimate the coefficients of a parametric wind field model [14], and both were tested on measurement data obtained from an operational wind farm. Comparing the two works, the implementation of Sinner et al. was computationally much simpler than that of Annoni et al., but did not explicitly identify turbines with faulty sensors. Another key difference between the two techniques was the use of a parametric model: the parametric model of Sinner et al. [14] can be easily used to produce wind field estimates at any location within the farm, whereas producing estimates away from the turbine locations is not straightforward using the nonparametric method of Annoni et al. [13].

In this work, we consider applying Gaussian process (GP) regression [16] for estimation of the wind direction within a wind farm. This approach strikes a balance between the qualities of parametric and nonparametric methods described above: the GP regression model, while not normally considered a parametric model, can produce an output estimate for a continuum of input locations. We also demonstrate a method of using the GP to identify faulty or biased turbines. To our knowledge, this is the first use of a GP to fully reconstruct the characteristics of a wind field online based on wind turbine measurements. Related works have used GPs to correct the errors in a simplified physics-based model for power prediction [17] and for short-term forecasting [18, 19, 20, 21, among others]. The latter techniques are generally interested in providing a single estimate of the wind speed or power over the entire farm for the purpose of forecasting the available farm power, whereas in this study we are interested in resolving the wind direction down to the turbine level.

The rest of this paper is organized as follows. The GP method is presented in Section II, which also provides an overview of the simulation technique used. Section III then describes the use of the methods of Section II and evaluates

<sup>1</sup>Delft Center for Systems and Control (DCSC), Faculty of Mechanical, Maritime and Materials Engineering (3mE), Delft University of Technology, The Netherlands {d.c.vanderhoek, j.w.vanwingerden}@tudelft.nl

<sup>2</sup>University of Colorado Boulder, USA {michael.sinner, pao}@colorado.edu

<sup>3</sup>National Renewable Energy Laboratory, Boulder, USA {michael.sinner, eric.simley}@nrel.gov

estimation performance. Finally, conclusions are given in Section IV.

## II. METHODOLOGY

This section presents the methodology that is required for the simulation and subsequent estimation of the wind field. A brief introduction to GP regression is given, including the application to wind turbine measurement data. A simple wind field model is introduced that will be used to generate data for the estimation framework. Lastly, a standard yaw control procedure is described. This will be used to evaluate the performance of the estimation framework.

### A. Gaussian process regression

A Gaussian process [16] is a nonparametric model consisting of a mean and covariance function. It assumes that a set of noisy function observations  $\mathbf{y} = [y^1 \ y^2 \ \dots \ y^n]$ , along with function evaluations  $\mathbf{y}^*$ , belong to a zero-mean multivariate Gaussian distribution:

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{y}^* \end{bmatrix} \sim \mathcal{N} \left( \mathbf{0}, \begin{bmatrix} K(\mathbf{x}, \mathbf{x}) + \sigma_n^2 \mathbf{I} & K(\mathbf{x}, \mathbf{x}^*) \\ K(\mathbf{x}^*, \mathbf{x}) & K(\mathbf{x}^*, \mathbf{x}^*) \end{bmatrix} \right), \quad (1)$$

where  $\mathbf{x} = [x^1 \ x^2 \ \dots \ x^n]$  is a matrix containing the measurement locations,  $\mathbf{x}^*$  is a matrix containing test locations, and the kernel matrix  $K$  and parameter  $\sigma_n$  are described below. In this case, we want to construct a spatio-temporal model of the wind field, which results in a vector  $x^i$  consisting of the turbine coordinates and timestamp belonging to a turbine measurement  $y^i$ .

The correlation between function values is determined by a covariance function based on the respective inputs and a set of hyperparameters. To ensure a certain level of smoothness between the function estimates, the squared exponential covariance function is used to quantify the similarity between two input vectors  $\mathbf{x}$  and  $\mathbf{x}^*$ , i.e.:

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

The covariance matrix is constructed using a set of hyperparameters  $\lambda_f^2$ ,  $\mathbf{\Lambda}_x = \{\lambda_{x_1^2}, \dots, \lambda_{x_n^2}\}$ , and  $\sigma_n^2$ , which denote the signal variance, characteristic length-scales and noise variance, respectively. The hyperparameters can be determined by optimizing the marginal likelihood of the data.

After conditioning the hyperparameters on the measurement data, we are able to infer the function value of the Gaussian process for any trial location  $\mathbf{x}^*$  with the following equation:

$$\boldsymbol{\mu}^* = K(\mathbf{x}^*, \mathbf{x}) (K(\mathbf{x}, \mathbf{x}) + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y}, \quad (3a)$$

$$\boldsymbol{\Sigma}^* = K(\mathbf{x}^*, \mathbf{x}^*) - K(\mathbf{x}^*, \mathbf{x}) (K(\mathbf{x}, \mathbf{x}) + \sigma_n^2 \mathbf{I})^{-1} K(\mathbf{x}, \mathbf{x}^*), \quad (3b)$$

where  $\boldsymbol{\mu}^*$  and  $\boldsymbol{\Sigma}^*$  are the mean and variance of the cost function at trial locations  $\mathbf{x}^*$ , respectively. This not only allows for estimation of the wind field at turbine locations, but at any location inside the farm, specified by  $\mathbf{x}^*$ .

### B. Wind field model

To test the effectiveness of the GP regressor in estimating the true wind directions within a wind farm, we require a representative wind field model. We follow a similar procedure to that presented by Simley et al. [15], and let the point wind direction, as measured by the wind vane on turbine  $T$  at time  $t$ , be broken down as

$$\phi^{(T)}(t) = \phi_1^{(T)}(t) + \phi_t^{(T)}(t), \quad (4)$$

where  $\phi_t$  is the turbulent high-frequency component and  $\phi_1$  is the low-frequency component, with spectra defined by (3b) & (3c) from Simley et al. [15], respectively. We assume that the high-frequency component is uncorrelated between the turbines, and randomly generate  $\phi_t^{(T)}$  for each turbine by assigning a uniformly random phase to the Fourier coefficients of the high-frequency component before computing the inverse Fourier transform.

On the other hand, we assume that the low-frequency component moves downstream at the mean wind speed  $U$  without changing. To do so, we assume a spatially and temporally constant underlying mean wind direction  $\bar{\phi}$ , which governs the advection of the flow field, for the duration of the simulation. The low-frequency component at turbine  $T$  is then

$$\phi_1^{(T)}(t) = \phi_1^{(0)}(t + d^{(T)}/U) \quad (5)$$

where  $\phi_1^{(0)}$  is the low-frequency wind direction component at some reference point in or near the farm, and  $d^{(T)}$  is turbine  $T$ 's distance downstream (using wind direction  $\bar{\phi}$ ) from the reference point. For turbines upstream of the reference point,  $d^{(T)} < 0$ . We therefore generate a single low-frequency wind direction signal  $\phi_1^{(0)}$  according to Simley et al. [15, (3c)], and use (5) to compute the low-frequency wind direction at the turbine locations. Note that this method of simulation means that the low-frequency wind direction components are completely correlated between turbines, which is a simplification. However, the combination of the correlated low-frequency component and uncorrelated high-frequency component (4) ensures that the coherence between the wind direction signals at different turbines is decreasing as a function of frequency.

### C. Turbine yaw control

Wind direction measurements are used predominantly for turbine yaw control. If the wind vane, located on the turbine nacelle, detects a significant and persistent nonzero value, the turbine is considered “misaligned”, and the yaw motor engages to turn the turbine into the wind and drive the misalignment to zero.

Letting  $\gamma^{(T)}$  denote the yaw angle of turbine  $T$  and the low-frequency wind direction component  $\phi_1^{(T)}$  be the “true” wind direction that we would like the yaw controller to track, we denote by

$$\tilde{\gamma}^{(T)} := \phi_1^{(T)} - \gamma^{(T)} \quad (6)$$

the turbine misalignment. Note that we are using a clockwise-positive coordinate system. The power lost due to




Fig. 1. Layout of a generic 12-turbine wind farm used for the case study. The turbines are spaced with a distance of 800 m in the horizontal direction and 600 m in the vertical direction.

yaw misalignment is modelled using a cosine-squared loss for simplicity [22].

The yaw controller we implement for this study can be described by a state machine with three modes of operation: yawing counterclockwise at a fixed yaw rate of  $\dot{\gamma} = 0.3$  deg/s; yawing clockwise at the same rate; and off (turbine not currently yawing). Based on an estimate of the wind direction  $\hat{\phi}_1^{(T)}$ , when the absolute yaw misalignment reaches a threshold value  $\tilde{\gamma}_{\text{threshold}} = 8$  deg, the turbine begins yawing to drive the misalignment to zero and stops yawing once zero is crossed. Typically,  $\hat{\phi}_1^{(T)}$  is taken to be a filtered version of the wind direction measured by the wind vane, which is presented as the baseline estimate in Section III-A. In Section III-B, the estimate will be replaced with the output of the GP model. The choice of yaw rate  $\dot{\gamma}$  and threshold value  $\tilde{\gamma}_{\text{threshold}}$  are based on the work of Bossanyi and Simley et al. [23, 15].

### III. CASE STUDY

In order to evaluate the performance of the GP wind field estimation framework, a case study is performed on the wind farm model shown in Fig. 1. The conditions in the wind farm are simulated using the wind field model presented in Section II-B for a constant wind speed of  $U = 8$  m/s. A distinction is made between online and offline estimation of the wind direction. This is due to the effect of adding new measurements to the data set used for GP regression, which results in less smooth function approximation over time.

Although we use a constant wind speed in this work to simplify the simulation procedure, the GP does not use wind speed information explicitly and is not limited to constant wind speeds. However, the GP hyperparameters will likely need to be relearned in situations where the wind speed, as well as other atmospheric conditions such as turbulence intensity, wind direction variability, and shear, change.

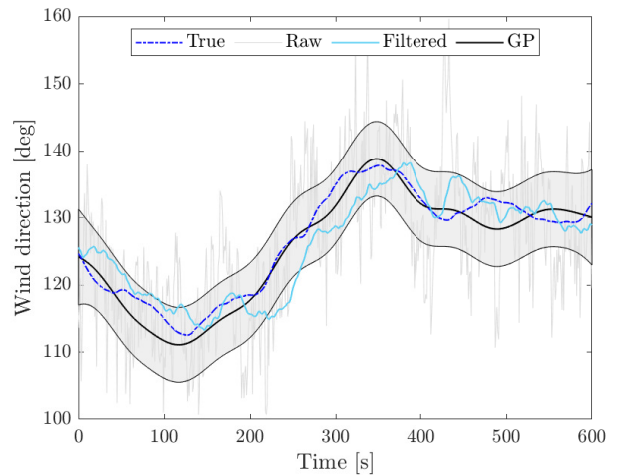


Fig. 2. Comparison of wind direction signals for a historical set of wind direction measurements of a single turbine ( $T_5$ ). The shaded area surrounding the GP estimate signifies the 95% confidence bounds of the function estimate.

#### A. Offline estimation

The GP estimation framework is initially applied to a historical set of noisy wind direction measurements for comparison with the low-frequency wind direction signals and the turbine measurements. Furthermore, a first-order low-pass filter is applied to the individual wind turbine measurements to serve as a baseline wind direction estimate over time. A time constant of  $\tau = 35$  s is used to construct the wind direction filter. This is consistent with the tuning used by Simley et al. to emulate the behavior of operational turbines [15].

The GP model for the entire wind farm is constructed by providing it with a set of noisy wind direction measurements  $\phi^{(T)}(t)$  from each turbine, along with the respective turbine coordinates and timestamps. After optimizing the hyperparameters for the given set of measurements, the GP model is able to provide the wind direction at any desired location and time according to (3). In order to adequately capture the low-frequency content of the wind direction in the GP model, ten minutes worth of wind turbine measurements with a time interval of  $\Delta t = 10$  s were used as input data. A visual comparison of the different wind direction signals is provided in Fig. 2 for a single turbine. It can be observed that in this instance the original low-frequency signal is closely approximated by the GP model mean and falls completely within the 95% confidence bounds of the function estimate. In contrast, the filtered signal is affected more by the turbulence on the raw measurement signal, resulting in a larger error with the low-frequency signal.

The previous comparison alone does not offer sufficient evidence that the GP model is always able to provide a better estimate of the true conditions for any turbine at all times. In order to test if this is the case, we performed a set of Monte Carlo simulations over the whole wind rose and using different realizations of the generated wind direction profile. The accuracy of the different wind direction signals

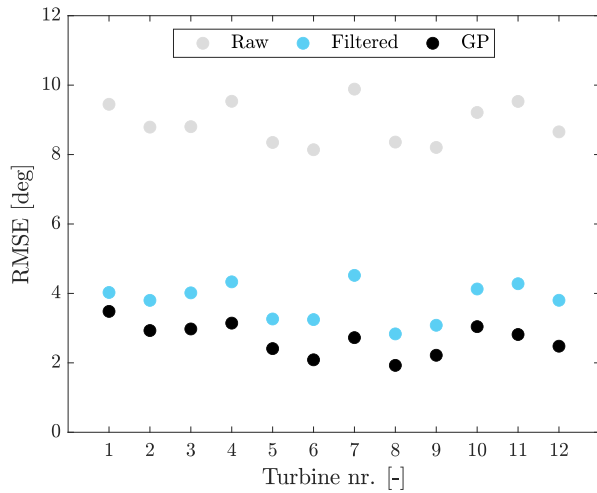


Fig. 3. Root-mean-square error between the different wind direction estimates and the simulated low-frequency signal for all 12 turbines based on a Monte Carlo simulation with 500 different realizations.

is expressed in terms of the root-mean-square error (RMSE) of the respective signals and the low-frequency signal. The results of the Monte Carlo simulation are presented in Fig. 3, and clearly show that the GP model outperforms standard first-order filtering of the raw measurements.

Until now, we have assumed that the turbine measurement signals are only influenced by local turbulence. In reality, these signals can also be affected by additional noise or bias due to a mechanical failure or inaccurate calibration. Because the Gaussian process model can estimate the wind direction at any location inside the wind farm, a simple leave-one-out cross-validation can be performed to determine whether a turbine measurement is biased.

The cross-validation for a single turbine is performed by removing the measurements from that particular turbine from the training data set and then estimating the wind direction at the turbine location using the measurements from surrounding turbines. This process is repeated for each of the turbines in the wind farm, after which the RMSE between the estimated wind direction signals and the filtered turbine measurements are evaluated. Based on the results from Fig. 2, a small error is expected between the estimated and filtered wind direction in the unbiased case. However, when a bias is present in one of the turbines, this should result in a clear distinction in the RMSE values of the different turbines. Ideally, the cross-validation procedure should be redone once a measurement bias has been identified, in order to prevent contamination of the other turbine results.

Fig. 4 shows the cross-validation results after a bias of  $\phi_{\text{bias}} = 7$  deg has been added to the measurements of turbine  $T_8$ . A clear peak is visible for this turbine, indicating a possible bias in the measurement data. When a potential bias has been identified, the measurement data of the corresponding turbine is removed from the GP model and the wind direction is estimated from nearby turbine measurements. An example of this is given in Fig. 5, where




Fig. 4. Results of the leave-one-out cross-validation expressed in the RMSE between the GP model estimate and the first-order filtered wind direction measurement.

the GP model is compared to the true low-frequency signal and the contaminated measurement signal. Notice that the confidence bounds of the GP model have increased compared to the estimate from Fig. 2, as a result from removing the measurements of that particular turbine from the GP model. Based on this figure we can conclude that it is possible to obtain a reasonable estimate of the local wind conditions, even when a turbine is affected by additional sensor noise or bias, if the biased turbine is identified and its measurements are excluded from the GP model.

### B. Online estimation

The previous section showed that the Gaussian process model is able to improve the accuracy of the wind direction




Fig. 5. Comparison of wind direction signals of a single turbine ( $T_8$ ) for a historical set of wind direction measurements, when the turbine measurements are affected by a bias. The GP model estimates the wind direction based on measurement data from surrounding turbines, and the shaded area surrounding the GP estimate signifies the 95% confidence bounds of the function estimate.




Fig. 6. Comparison of the original wind direction profile with the filtered measurement signal and the online GP model estimate. The resulting yaw orientations based on the latter two signals are shown using dashed lines.

measurement based on historical measurements. When the GP estimate is used as input to the yaw controller from Section II-C, some minor modifications to the framework are required. The initial procedure still consists of conditioning the hyperparameters to the last 10 minutes' worth of measurements. After this, the wind direction is inferred at time  $t$  for all wind turbine locations. At time  $t + \Delta t$ , with  $\Delta t = 10$ s, newly obtained measurements are added to the measurement set  $\mathbf{y}$  from (1), and the measurements of the oldest time instance are omitted. Once the data set has been updated, the mean and variance of the wind direction at time  $t + \Delta t$  are computed using (3). The online GP model is found to be particularly sensitive to new measurements at the edge of the input data, i.e., the current time step. This often results in a wind direction estimate that is less smooth over time when compared to the offline model. Additional care should be taken when the mean wind direction, as well as other atmospheric conditions, have changed significantly from the time of the original training data set. When this is the case, the hyperparameters should be optimized once more based on the current dataset.

An example of the online estimation procedure is given in Fig. 6. Here, the online GP estimate for a single turbine is shown next to the original low-frequency signal and the filtered wind direction measurement. Furthermore, both the GP model and the filtered measurement have been applied to the yaw controller described in Section II-C, resulting in two realizations of the yaw orientation over time. In this particular instance, it can be observed that the yaw orientation following from the GP model follows the original wind direction signal more closely than the yaw orientation from the filtered signal.

In order to find out whether the GP model is able to systematically improve the performance of the yaw controller, multiple realizations of the generated wind field along with different mean wind directions are applied to the online estimation framework. Figure 7 presents the average power

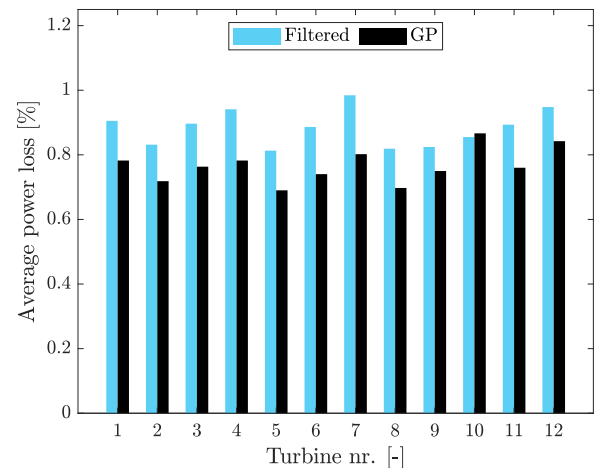


Fig. 7. Average power loss due to yaw misalignments for each of the turbines based on the first-order-filtered wind direction measurement and the GP model. The results are obtained from 1000 realizations of the wind direction model.

loss due to yaw misalignments as computed by the cosine-squared rule [22]. This figure shows that the GP model is able to reduce yaw misalignments, thereby decreasing the power that is lost due to misalignment with the incoming wind field. Additionally, the effect of the GP model on the yaw activity relative to the filtered measurement signal is shown in Fig. 8. From this we can conclude that, using the online GP model estimate in place of the raw wind direction measurement for turbine yaw control, we may be able to achieve a slight increase in power due to a smaller yaw misalignment while achieving a significant decrease in the amount of yaw activity.

#### IV. CONCLUSION

A Gaussian process (GP) regression model was introduced to model the ambient wind conditions inside a wind farm

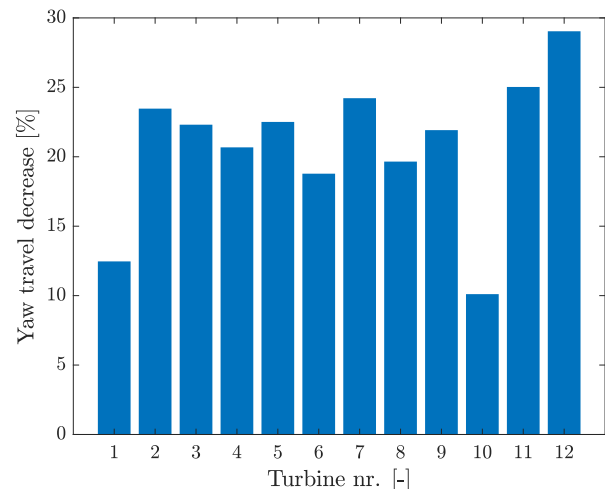


Fig. 8. The average decrease in distance traveled by yawing with the GP model compared to when the filtered measurement signal is used as input for the yaw controller. The results are obtained from 1000 realizations of the wind direction model.

based on wind turbine measurements. By conditioning the GP model on a historical set of noisy wind direction measurements, an accurate estimate of the wind direction over time at each turbine can be obtained. Using such an offline model of the wind direction, potential biases or sources of additional noise on a turbine level can be identified.

Estimating the wind direction in an online fashion leads to a slight deterioration in the accuracy of the estimate; however, the GP model is still able to achieve a smaller yaw misalignment compared to the baseline yaw controller. The performance is expected to increase even further when a bias in the wind direction measurements has been identified prior to the online estimation. A positive side effect of using the GP model as input to the yaw controller is a significant decrease in the amount of yaw activity.

In this work, only the wind direction was estimated with the GP model, but the current framework can easily be expanded to model wind speed as well. Furthermore, when additional measurements inside a wind farm are available from a meteorological mast or lidar, for example, these can be added to the input data of the GP model while taking into account the possibly higher accuracy of the additional measurement devices. Another possible application of the presented framework can be found when coupling it to dynamic wake models in the form of an observer. The observed wake model may subsequently be used to optimize the power of a wind farm using control methods such as wake steering.

#### ACKNOWLEDGMENT

This work is part of the research programme “Robust closed-loop wake steering for large densely space wind farms” with project number 17512, which is (partly) financed by the Dutch Research Council (NWO).

This work was authored in part 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 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 U.S. Government purposes.

#### REFERENCES

- [1] P. Fleming, P. M. Gebraad, S. Lee, J.-W. van Wingerden, K. Johnson, M. Churchfield, J. Michalakes, P. Spalart, and P. Moriarty, “Simulation comparison of wake mitigation control strategies for a two-turbine case,” *Wind Energy*, vol. 18, no. 12, pp. 2135–2143, 2015.
- [2] F. Campagnolo, V. Petrović, C. L. Bottasso, and A. Croce, “Wind tunnel testing of wake control strategies,” in *Proc. American Control Conf.*, 2016, pp. 513–518.

- [3] P. Fleming, J. Annoni, J. J. Shah, L. Wang, S. Ananthan, Z. Zhang, K. Hutchings, P. Wang, W. Chen, and L. Chen, “Field test of wake steering at an offshore wind farm,” *Wind Energy Science*, vol. 2, no. 1, pp. 229–239, 2017.
- [4] M. F. Howland, S. Lele, and J. Dabiri, “Wind farm power optimization through wake steering,” *Proc. National Academy of Sciences USA*, vol. 116, pp. 14495–14500, 2019.
- [5] B. M. Doekemeijer, S. Kern, S. Maturu, S. Kanev, B. Salbert, J. Schreiber, F. Campagnolo, C. L. Bottasso, S. Schuler, F. Wilts, T. Neumann, G. Potenza, F. Calabretta, F. Fioretti, and J.-W. van Wingerden, “Field experiment for open-loop yaw-based wake steering at a commercial onshore wind farm in Italy,” *Wind Energy Science*, vol. 6, no. 1, pp. 159–176, 2021.
- [6] P. Fleming, J. Aho, P. Gebraad, L. Pao, and Y. Zhang, “Computational fluid dynamics simulation study of active power control in wind plants,” in *Proc. American Control Conf.*, 2016, pp. 1413–1420.
- [7] J.-W. van Wingerden, L. Pao, J. Aho, and P. Fleming, “Active power control of waked wind farms,” *IFAC-PapersOnLine (Proc. IFAC World Congress)*, vol. 50, no. 1, pp. 4484–4491, 2017.
- [8] M. Vali, V. Petrović, L. Y. Pao, and M. Kühn, “Model predictive active power control for optimal structural load equalization in waked wind farms,” *IEEE Trans. Control Systems Technology*, pp. 1–15, 2021.
- [9] B. Doekemeijer and J.-W. van Wingerden, “Observability of the ambient conditions in model-based estimation for wind farm control: A focus on static models,” *Wind Energy*, vol. 23, no. 9, pp. 1777–1791, 2020.
- [10] P. Fleming, P. Gebraad, J. W. van Wingerden, S. Lee, M. Churchfield, A. Scholbrock, J. Michalakes, K. Johnson, and P. Moriarty, “SOWFA super-controller: A high-fidelity tool for evaluating wind plant control approaches,” National Renewable Energy Laboratory, Golden, CO, Tech. Rep. NREL/CP-5000-57175, 2013.
- [11] S. Boersma, B. Doekemeijer, M. Vali, J. Meyers, and J.-W. van Wingerden, “A control-oriented dynamic wind farm model: WFSim,” *Wind Energy Science*, vol. 3, no. 1, pp. 75–95, 2018.
- [12] 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*, vol. 3, no. 2, pp. 819–831, 2018.
- [13] J. Annoni, C. Bay, K. Johnson, E. Dall’Anese, E. Quon, T. Kemper, and P. Fleming, “Wind direction estimation using SCADA data with consensus-based optimization,” *Wind Energy Science*, vol. 4, no. 2, pp. 355–368, 2019.
- [14] M. Sinner, L. Y. Pao, and J. King, “Estimation of large-scale wind field characteristics using supervisory control and data acquisition measurements,” in *Proc. American Control Conf.*, 2020, pp. 2357–2362.
- [15] E. Simley, P. Fleming, and J. King, “Design and analysis of a wake steering controller with wind direction variability,” *Wind Energy Science*, vol. 5, no. 2, pp. 451–468, 2020.
- [16] C. E. Rasmussen and C. K. Williams, *Gaussian Processes for Machine Learning*. MIT press, 2006.
- [17] L. E. Andersson and L. Imsland, “Real-time optimization of wind farms using modifier adaptation and machine learning,” *Wind Energy Science*, vol. 5, no. 3, pp. 885–896, 2020.
- [18] J. Hu and J. Wang, “Short-term wind speed prediction using empirical wavelet transform and Gaussian process regression,” *Energy*, vol. 93, pp. 1456–1466, 2015.
- [19] C. Zhang, H. Wei, X. Zhao, T. Liu, and K. Zhang, “A Gaussian process regression based hybrid approach for short-term wind speed prediction,” *Energy Conversion and Management*, vol. 126, pp. 1084–1092, 2016.
- [20] J. Yan, K. Li, E. Bai, J. Deng, and A. M. Foley, “Hybrid probabilistic wind power forecasting using temporally local Gaussian process,” *IEEE Trans. Sustainable Energy*, vol. 7, no. 1, pp. 87–95, 2016.
- [21] A. Aziz Ezzat, “Turbine-specific short-term wind speed forecasting considering within-farm wind field dependencies and fluctuations,” *Applied Energy*, vol. 269, no. 115034, 2020.
- [22] T. Friss Pedersen, S. Gjerding, P. Ingham, P. Envoldsen, J. K. Hansen, and H. K. Jørgensen, “Wind turbine power performance verification in complex terrain and wind farms,” Risø National Laboratory, Roskilde, Denmark, Tech. Rep. 1330(EN), 2002.
- [23] E. Bossanyi, “Combining induction control and wake steering for wind farm energy and fatigue loads optimisation,” in *J. Physics: Conf. Series (Torque)*, vol. 1037, no. 032011, 2018.