Highlighting the Impact of Yaw Control by Parsing Atmospheric Conditions Based on Total Variation

Preprint

Nicholas Hamilton

National Renewable Energy Laboratory

Presented at NAWEA/WindTech 2019
Amherst, Massachusetts
October 14–16, 2019

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Conference Paper
NREL/CP-5000-74409
March 2020

Highlighting the Impact of Yaw Control by Parsing Atmospheric Conditions Based on Total Variation

Preprint

Nicholas Hamilton

National Renewable Energy Laboratory

Suggested Citation
Hamilton, Nicholas. 2020. *Highlighting the Impact of Yaw Control by Parsing Atmospheric Conditions Based on Total Variation: Preprint*. Golden, CO: National Renewable Energy Laboratory. NREL/CP-5000-74409. https://www.nrel.gov/docs/fy20osti/74409.pdf.

Conference Paper
NREL/CP-5000-74409
March 2020

National Renewable Energy Laboratory
15013 Denver West Parkway
Golden, CO 80401
303-275-3000 • www.nrel.gov

NOTICE

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 the U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Wind and Water Technologies Office. The views expressed herein 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.

Cover Photos by Dennis Schroeder: (clockwise, left to right) NREL 51934, NREL 45897, NREL 42160, NREL 45891, NREL 48097, NREL 46526.

NREL prints on paper that contains recycled content.

Highlighting the impact of yaw control by parsing atmospheric conditions based on total variation

Nicholas Hamilton

National Renewable Energy Laboratory, Golden, Colorado, USA

E-mail: nicholas.hamilton@nrel.gov

Abstract. Identification of atmospheric conditions within a multivariate atmospheric dataset is a necessary step in the validation of wind plant control strategies. Most often, operating conditions are characterized in terms of aggregated observations and assume that the atmosphere is 'quasi-steady'. Aggregation of observations without regard to covariance between time series discounts the dynamical nature of the atmosphere and is not su ciently representative of wind plant operating conditions. Identification and characterization of continuous time periods with atmospheric conditions that have a high value for analysis or simulation sets the stage for more advanced model validation and the development of real-time control and operation strategies. Controlling observational data for statistical stationarity highlights significant enhancements to the power production of waked turbines under wake steering wind plant control. Results in the current study emphasize the scope and intended range of wake models used for wind plant control and suggest that either models be defined to account for the transient nature of the atmosphere, or that their validation and application be geared to stationary atmospheric conditions.

1. Background

Wind plant control research has gained a great deal of attention in recent years as the potential to decrease the levelized cost of electricity through mitigation of wake losses has come into focus. In a general sense, wind plant control describes any operational strategy that coordinates control actions of individual constituent wind turbines in order to augment the performance of the wind plant as a whole. Plant performance in this sense is typically measured in terms of power production, but may also be quantified by bankability or expected return on investment, as business metrics may also incorporate operations and maintenance activities. One popular strategy for wind plant control is wake steering through the intentional introduction of a yaw offset with respect to the ambient wind direction, thereby increasing annual energy production through mitigating wake losses [20, 25].

Wake steering as a control strategy has been supported through theoretical development, wind tunnel experimentation [6, 14, 19, 22, 23, 24], and high-fidelity simulation [18, 1, 10]. In addition to demonstrating the merits of wake steering, experimental and numerical research have contributed to the development of computationally economical wake modeling platforms that can be used for myriad engineering processes in the wind energy industry [21, 16, 5, 11]. Control-oriented analytical wake models are now capable of modeling wake steering strategies to in turn support the design and analysis of wind plant controllers and estimate the benefits in terms of wind plant performance. Many wake velocity, turbulence, and deflection, models

are available for research, control design, and optimization in NREL's FLow Redirection and Inductions in Steady State platform (FLORIS) [21].

A utility-scale demonstration of the merits of wake steering for energy production is under way, executed by the National Renewable Energy Laboratory (NREL) at a wind plant in Colorado. A subset of turbines in a commercial wind farm was designated as a test and measurement site and now hosts an array of sensing equipment including a profiling lidar, meteorological (met) mast, and two sodars. Wake steering control has been prototyped to limit wake interactions that lead to power losses based on the predictions made by the FLORIS. The continued development and future validation of predicted energy gains introduced by wake steering wind plant control will require delicate consideration and quality control of field observations in order to control for the modeled atmospheric conditions.

Simultaneous observation of multiple thermodynamic and kinematic quantities reported in field observations are necessary to more completely characterize the dynamical state of the atmosphere. Directly considering multiple disparate data channels simultaneously represents a challenge in that each quantity has different engineering units and that variation within each channel may occur over a distinct scale. Atmospheric conditions are frequently characterized by considering wind speed, wind direction, and turbulence intensity or thermal stratification, each of which have different units, ranges, and statistical distributions. Additionally, direct comparison of statistical quantities (measures of central tendency, variability, or higher statistical moments) discount the inherent coupling between quantities of interest that underpin atmospheric physics.

A mathematically rigorous metric for the total variation of a sample of time series data is necessary to identify and characterize specific dynamical events and atmospheric conditions while taking into account covariance between atmospheric quantities. The total variation is used to identify and remove transient conditions from a set of dynamic atmospheric data in order to assess the benefits of yaw control for wake steering for a utility-scale wind plant. Stationary atmospheric conditions in particular are of interest in developing and validating engineering models that can be used to prototype wind plant control strategies.

2. Data and Quality Control

A small cluster of wind turbines from a commercial wind plant were made available for validation of wake steering control. The arrangement of turbines shown in Figure 1 is unaffected by the rest of the wind plant given the prevailing wind directions from the northwest during the winter and the southeast during the summer. Additionally, the wind plant abuts an escarpment to the south, providing additional variability of inflow conditions; northerly winds issue from flat terrain, southerly winds reflect complex terrain. Five turbines are included in the wind plant control field campaign. Given the prevailing wind directions, conditions in which the central turbine (T3) is waked by an upstream turbine (either T2 or T4) are fairly common, making wake losses a significant consideration and wake steering an attractive control strategy. To gauge the changes introduced to the energy production of the wind plant through wake steering, two machines (T1 and T5) are retained as reference turbines. Reference turbines remain uncontrolled and unwaked during conditions where control are enabled. Wind directions that lead to full waking of T3 (324° from T2 or 134° from T4) are shown in the figure.

Locations of the field instrumentation are also indicated in Figure 1. Profiles of wind speed and direction (denoted as u and θ respectively) were collected by a Leosphere Windcube v2 profiling lidar, estimated every second but averaged to 1-min intervals. The Windcube v2 samples line-of-sight velocities centered every 20 meters from 40 m up to 180 m in four cardinal directions with a fixed elevation angle of 28° from vertical, followed by a fifth beam oriented vertically. A time series of turbulence intensity (TI) is derived from a sonic anemometer on the top of the met mast. Wind speed and turbulence intensity estimates from the sodars are not used in the following analysis. Interested readers are referred to the recent work by Fleming,

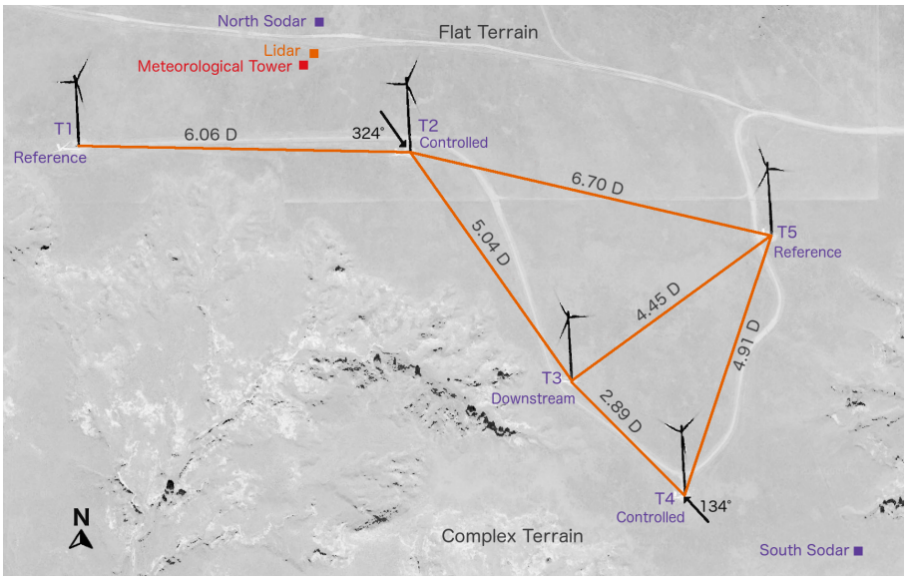

Figure 1: Subset of wind plant for controls testing. Wake steering implemented on Turbine 2 (T2) and Turbine 4 (T4) to mitigate wake losses by Turbine 3 (T3). Turbine 1 (T1) and Turbine 5 (T5) act as reference turbines. Image reproduced from Fleming, et al.[9]

et al. [9] for additional information on the experimental campaign and the additional data collected.

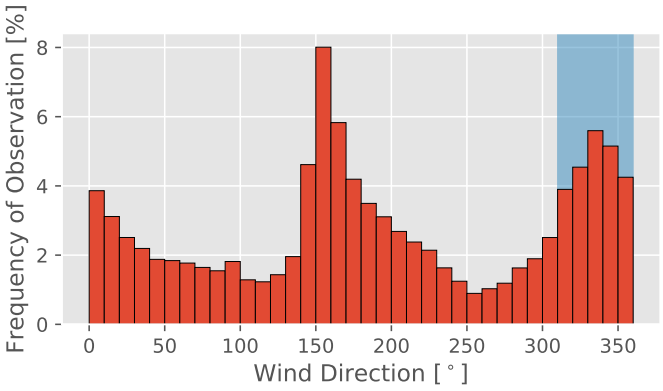

Figure 2: Distribution of observations by wind direction. The shaded region indicates direction sector for which yaw control may be enabled.

Data considered during the current analysis correspond to times between May 2, 2018 and June 4, 2019, only during periods when all instruments and turbines are reporting nominally. The full distribution of wind direction observations is shown in Figure 2 in the red histogram. The frequencies of observations are divided into wind direction sectors of 10°, and show the bimodality of the atmospheric conditions. Data reported in the following figures include observations made throughout the year and show both winter behavior (flow from the northwest) and summer (flows from south). In the results explored here, the data considered is further

limited only to conditions in which the wake from T2 may impinge on T3 without intervention. The blue band highlights the region of interest, winds coming from between $310° \leq \theta \leq 360°$. Throughout the data collection period, the yaw control was toggled on and off every hour in order to sample a wide range of atmospheric conditions over which to test the wake steering strategy.

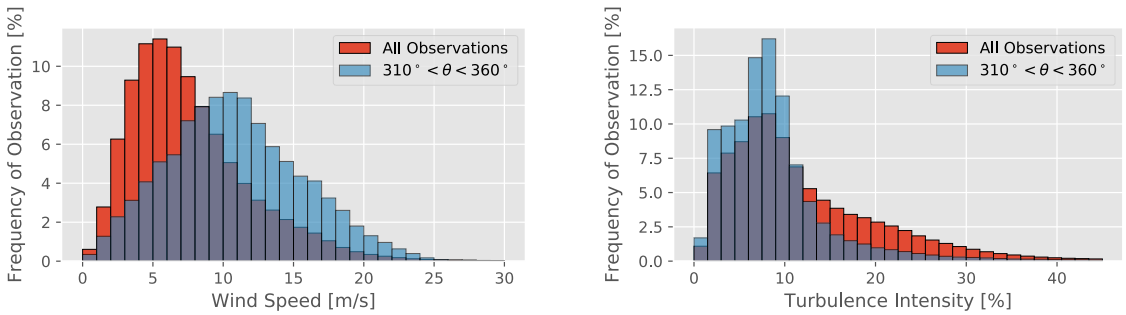

(a) Wind speed distribution

(b) Turbulence intensity distribution

Figure 3: Histograms of wind speed and TI. Data in red indicate the full range of observations, data in blue indicate observations from the wind direction sector of interest.

Wind speed and turbulence intensity are also considered to characterize the atmospheric conditions. Figure 3 compares histograms of wind speed and turbulence intensity over the full range of observations during the field campaign (red) and from the wind direction section for which controls are implemented ($310° \leq \theta \leq 360°$, blue). The distribution of wind speed indicates that the observations from the control sector have a higher mean value (11.4 m/s) and do not conform to a Weibull distribution as closely as the full range of observations. Observations in the control sector also delineate lower turbulence intensity as compared to the full range, which is expected given that the terrain to the north is much simpler than from the south. Wind speed and TI roses are shown in Figure 4, representing the distribution of observations sorting by wind direction and either wind speed or turbulence intensity. Wind roses, or other multidimensional histograms, hint at the inherent coupling between atmospheric variables for any given site, but still discount any dynamic behavior observed in reality.

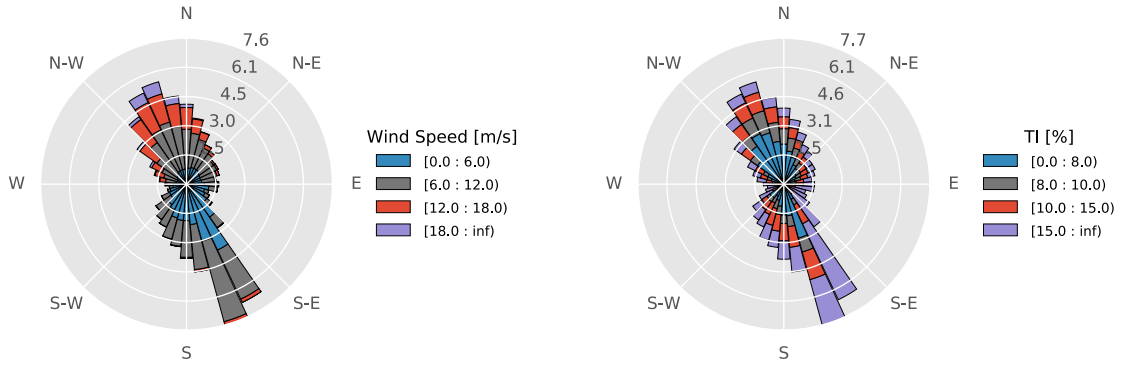

(a) Wind rose

(b) TI rose

Figure 4: Wind and TI roses from observational data.

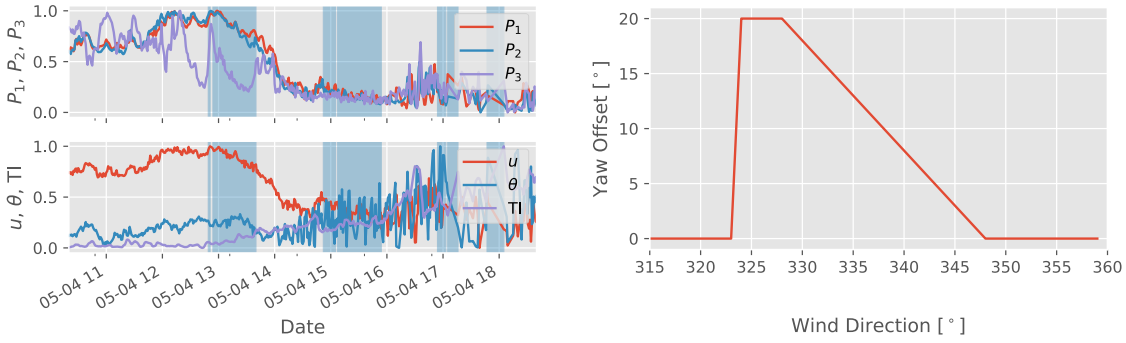

(a) Normalized power and atmospheric variables.

(b) Prescribed yaw offset T2 during control periods.

Figure 5: Sample power and atmospheric data. Shaded regions in (a) are subject to the yaw offset in (b).

To evenly sample atmospheric conditions for which benefits of yaw control could be assessed, controls were enabled in alternating hours. In addition to the toggling windows in which controls were permitted, the particular atmospheric conditions also had to be parsed to validate that the turbines were operating in conditions that would be favorable for control. Figure 5a shows a sample from the full data record. All data channels have ben normalized to a default interval of $[0, 1]$ for ease of comparison. Shaded regions in the figure indicate times in which the controls were enabled and active, that is times when yaw offset prescription was permitted and the bulk wind direction was favorable for testing wake steering to mitigate wake losses imposed on T3 by T2. Observations for which the controls are enabled (*controlled* conditions) were those in which T2 was given a yaw offset according to Figure 5b; when the control strategy was not enabled (*baseline* conditions), the turbines were operated according to their default behavior. The aggregate change in power production for T2 and T3 is estimated using T1 as an unaltered reference signal.

3. Theoretical Development

Aggregate statistical representation as in histograms shown in Figures 2 and 3 and in the rose figures shown in Figure 4 begins to describe the interdependence of the atmospheric variables considered in the current analysis, but cannot account for the dynamic nature of the atmosphere. A histogram, as a consequence of its composition, only denotes how frequently a given condition is observed without regard to what condition may precede or follow. For example, a given condition denoted by wind speed, wind direction, and turbulence intensity (u, θ, and TI) may represent the ensemble average of an inherently transient or dynamic event. Atmospheric conditions within any given sample may undergo a significant dynamical or transient change, but averaged values of the variables of interest can still fall within the stated bounds of a single bin within the full condition space.

Quantifying the variability of a set of data must include the correlation between data channels, or risk discounting any information regarding the relationship between variables. Any metric that combines the variability of each channel independently without accounting for covariance between the channels is incomplete and will not be su cient to fully quantify or characterize the state of a given system. Therefore, a method that accounts for not only the variation within each channel, but also the variation between channels is necessary to quantify the distribution of data across multiple channels into a single metric.

Below, each data block, $\mathbf{D}$, is a selected time period and corresponds to an array of size

of $[m, n]$. In this case, m is the length of the data block in time (30 min) and n is three, corresponding to the number of variables u, θ, and TI, (wind speed, direction, and turbulence intensity).

$$\mathbf{D} = [u(t),\ \theta(t),\ TI(t)] \tag{1}$$

TI is estimated by comparing the standard deviation of wind speed to the mean during a moving 10-minute window centered at each time stamp, t. Mentioned above, each channel represented in $\mathcal{D}$ has been normalized to an interval of $[0, 1]$. Normalization of the data makes comparison of data characterized by different ranges and engineering units mathematically meaningful.

The total variation, $\mathcal{V}$, of a system is a unitless metric to quantify statistical spread of a set of interdependent variables that is able to account for autocorrelation within each channel as well as covariance between channels. The covariance matrix $\mathbf{C}$ of the data block $\mathbf{D}$ is calculated for a subset taken from the full data, representing a continuous period of a specified duration,

$$\mathbf{C} = \mathbf{D}^T\mathbf{D} \tag{2}$$

$$= \begin{bmatrix} \overset{2}{u} & u\ \theta & u\ \mathrm{TI} \\ \theta\ u & \overset{2}{\theta} & \theta\ \mathrm{TI} \\ \mathrm{TI}\ u & \mathrm{TI}\ \theta & \overset{2}{\mathrm{TI}} \end{bmatrix} \tag{3}$$

In Eq. (3), $\mathbf{C}$ is a square matrix of size $n \times n$ representing the covariance between any pair of data channels. The principal components of the covariance matrix are derived through an eigenvalue decomposition,

$$\mathbf{Cv} = \quad \mathbf{v} \tag{4}$$

Eigenvectors are denoted as $\mathbf{v}$ and the eigenvalues as . By definition, eigenvectors are a set of orthonormal vectors that most e ciently span the space of the covariance matrix. Principal components are eigenvectors weighted by their respective eigenvalues, $\mathcal{P} = \mathbf{v}$. Total variation, $\mathcal{V}$, is the vector summation of all principal components. Because the principal components are orthogonal, the total variation can be equivalently expressed as the L_2-norm of the eigenvalues.

$$\mathcal{V} = ||\sum \mathcal{P}|| = \sum \tag{5}$$

Thus $\mathcal{V}$ is a single metric that can measure the spread of a multivariate data block. Any block in which the channels exhibit a large transient change or large variance results in a large value of $\mathcal{V}$, as does any block wherein multiple channels change in a coordinated manner. In the current work, large values of $\mathcal{V}$ indicate that the atmospheric conditions, described by time series of wind speed, wind direction, and turbulence intensity, change more during a given period. Small values of $\mathcal{V}$ indicate that the atmospheric conditions within a continuous 30 min block of data remain relatively stationary.

Effects of the wake steering implementation on the control and downstream turbine are highlighted by comparing to the production of the reference turbine. In energy ratio analysis, data are first limited to include only periods in which all turbines of interest were operating under normal conditions. All data including the power of T3 are binned by wind direction into sectors of 2° and split according to whether the wake steering controller was toggled on or off. For each bin, a ratio of energy is computed comparing the production of the subject turbine to that of the reference. When aggregate power production is of interest, production by all of the subject turbines is summed and the reference is increased by a factor equal to the number of subjects considered.

$$\rho = \frac{\Sigma_{i=1}^{N} P_{i|\mathrm{test}}}{N P_{\mathrm{ref}}} \tag{6}$$

where $i \in 1...N$ indexes each turbine in the aggregate energy ratio.

The energy ratio defined in Eq. (6) describes the energy produced by wind turbine or group of test turbines normalized by the respective reference values. When considering more than a single test turbine, a reference value for each test signal should be considered such that nominal unwaked power production corresponds to $\rho = 1.0$.

4. Results

The total variation, $\mathcal{V}$, is calculated for all continuous sets of observations of 30 minutes when the average wind direction issues from the NW ($310° \leq \theta \leq 360°$, denoted by the blue range in Figure 2). Figure 6 compares $\mathcal{V}$ for the observations considered in estimating the value of the implemented yaw-offset control. The full set of observations is shown in the red histogram (left), and split into the baseline and control conditions (center and right, respectively).

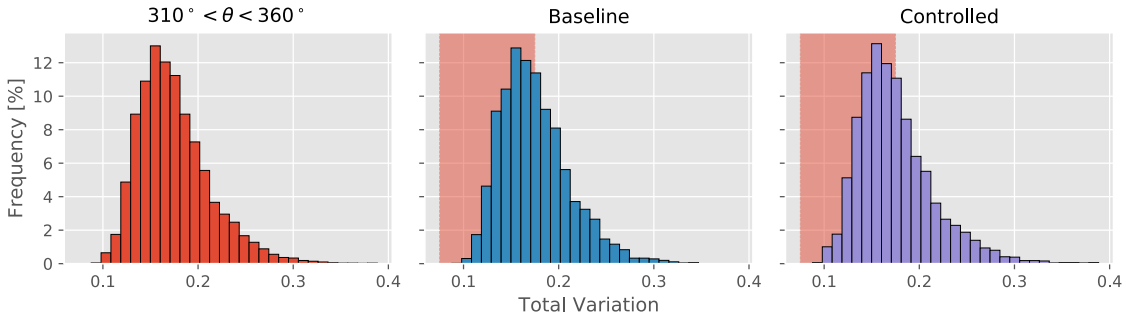

Figure 6: Distributions of $\mathcal{V}$ for all observations, periods during which yaw control is implemented, and periods in which the control is suppressed.

Shaded regions of the histograms of total variation during baseline and controlled conditions are those where the total value is less than the average of the full set of observations, $\mathcal{V} < \overline{\mathcal{V}}$. This threshold is somewhat arbitrary but does serve to separate the observations into observations of stationary and transient conditions. Establishing a threshold for $\mathcal{V}$ at the mean value reduces the number of observations considered in the evaluation of the benefits of the control action, which has a potentially negative impact on the statistical convergence of the sampled data. However, given that observations with low values of $\mathcal{V}$ are those with less variability, reducing the sample size does not appreciably increase the standard error of the sample.

Figure 7a compares the energy ratio for all observations to that calculated for stationary conditions. Figure 7a considers only T3 as the test turbine. Considering only T3 focuses the energy ratio on the turbine subject to wake losses by itself, without considering the aggregate power production of T2, which is itself the subject of the yaw control. Blue lines in the figure demonstrate the power production of T3 under baseline conditions, i.e. times when the yaw controller is not enabled and active. Purple lines show the complementary performance of T3 under the influence of the prescribed control action. Figure 7a shows that the minimum power produced by T3 occurs when the wind direction is 323°, when the wake of T2 impinges directly onto the rotor of T3.

To highlight the benefits of wake steering control, Figure 7b compares the energy ratio considering only times where $\mathcal{V}$ is below its mean value of 0.15. Quality controlling the sample for stationarity in the atmospheric conditions is done identically for the baseline conditions and the controlled conditions, i.e. transient events are filtered out of both samples. The most immediate effect of filtering based on $\mathcal{V}$ is to drive apart trends in the power production by T3 under baseline conditions and during yaw control. Changes seen in the energy ratio based on filtering out transient data suggests that simply binning atmospheric data based on average values of

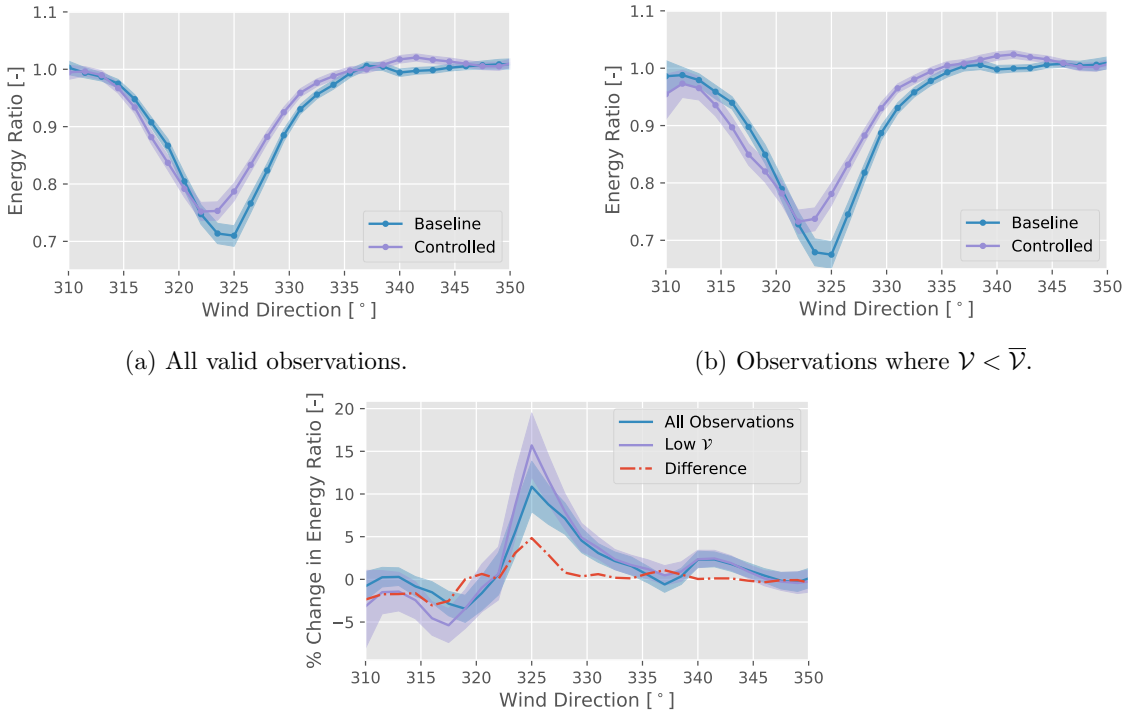

(a) All valid observations.

(b) Observations where $\mathcal{V} < \overline{\mathcal{V}}$.

(c) Change in energy ratio for T3.

Figure 7: Comparison of energy ratio for T3 alone considering all observations and only those corresponding to stationary conditions.

kinematic and dynamic qualities smears out distinct physical behaviors. Isolating observations during stationary periods highlights the underlying response of the wind turbines to the effects of yaw control.

Differences between the energy ratio considering the full range of valid observations and just those for which $\mathcal{V}$ remains below its average value are highlighted in Figure 7c, where the trend lines now indicate the percent change in energy ratio from one set of baseline conditions to the respective controlled conditions. Most notably, filtering observations based on $\mathcal{V}$ (purple line) shows a peak increase in the energy ratio of about 16%. This results suggests that T3 is able to produce 16% more energy under the influence of yaw control than under the baseline scenario given a wind direction of 325°. In contrast, the greatest change in the energy ratio for the full range of valid observations is approximately 10%, and occurs at around 325°. Together these trends suggest that an additional quality control check on the data to remove transient periods reveals an additional 6% change to the energy ratio that is obscured by dynamics in the atmospheric data.

The energy ratio can be similarly calculated for the combined power production of turbines T2 and T3, thus reconciling the benefits in power production for T3 against losses incurred by T2 by operating at an off-nominal set point. Suggested in eq. (6), the formulation for energy ratio can accept power production signals from multiple test turbines, but the reference needs to be multiplied by a factor equal to the number of test turbines in question. Figures 8a and 8b demonstrate the combined energy ratio for T2 and T3 for the full range of valid observations and the subset for which $\mathcal{V} < \overline{\mathcal{V}}$. Figure 8a indicates that only a nominal increase in the combined power production of T2 and T3 should be expected for $320° \leq \theta \leq 330°$, where wake interaction

is expected. Benefits of wake steering are also evident for wind directions $330° \leq \theta \leq 350°$, where the current results indicate that the combined power losses are lower during control periods than under baseline periods.

While the energy ratio for T2 and T3 in Figure 8b does show significant improvement where direct waking is avoided ($320° \leq \theta \leq 330°$), it also highlights a reduction in power losses for $330° \leq \theta \leq 350°$. Filtering observations to exclude transient data emphasizes the benefits for power production seen by T3, but it also underpins the off-nominal operation of the turbine for which the yaw offset is prescribed. Without controlling for transients in the atmospheric data, both of these effects are smeared out in the energy ratio. Comparing the percent change in energy ratio between baseline and controlled operation with and without the threshold of total variation still shows a significant benefit in terms of aggregate power production. Figure 8c shows that an additional 2% change to the energy ratio is illuminated by careful pretreatment of the observations.

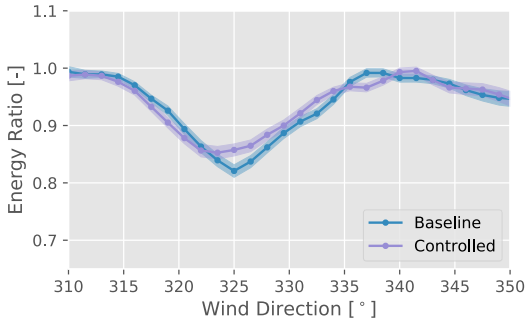

(a) All valid observations.

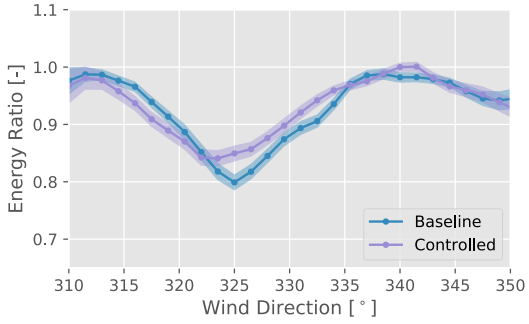

(b) Observations where $\mathcal{V} < \overline{\mathcal{V}}$.

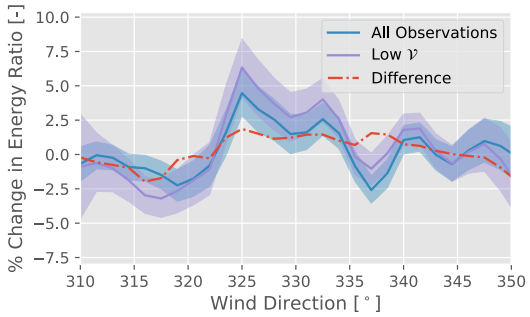

(c) Change in energy ratio for T2 and T3.

Figure 8: Comparison of energy ratio for T2 and T3 considering all observations and only those corresponding to stationary conditions.

5. Conclusions

Wind plant control by way of wake steering is a promising method for mitigating wake losses. Deflecting wakes by prescribing an offset in yaw with respect to the prevailing wind direction yields gains in power for downstream devices, at the cost of operating a turbine at an off-nominal set point. The current work examines changes in aggregate and individual wind turbine power production by comparing the energy ratio of selected turbines under baseline and controlled conditions. Changes in energy production are further highlighted by controlling the considered

samples for transients in the atmospheric data. Isolating periods in which the total variation of the atmosphere is less than or equal to the average variability increases the reported production of T3 from 10% to 16%.

Significant changes in the expected energy gains from wind plant control by removing transients in the considered data will impact the future design and optimization of wind turbine wake and control models. Many of the wake models currently included in FLORIS are assumed to take into account the stochastic nature of the atmosphere and are intended to represent average behavior given typical variability. However, the variability covered does not account for the types of dynamic events seen in observations of the atmosphere, e.g. wind speed ramps or changes in bulk flow direction. The current results highlight the need to consider the dynamic nature of the atmosphere in the evaluation of wind plant control strategies. Such consideration should come through explicit accounting for transient atmospheric events or through the definition of wake models.

The direct detection and classification of events or periods of interest within atmospheric data sets is vital to developing our understanding of wind plant response to new control strategies. The basis of the total variation method demonstrated here is easily extended to arbitrarily many data channels simultaneously and operates over any time scale or resolution within the limits of the data. Applying the total variation method with multiple atmospheric measurement systems across a wind plant will be an effective practice for identifying and accounting for spatial heterogeneity.

Regularizing a data block before calculation of the covariance matrix or the total variation can tune the method to pick out any event type of interest. For example, adding an objective function block, $\mathbf{f}$ to the consideration is a direct means of regularizing the data against transient events,

$$\mathbf{f} = [f_u(t),\ f_\theta(t),\ f_{\mathrm{TI}}(t)] \tag{7}$$

The difference between objective functions and their respective data is considered to be a regularized data block, and is noted with a caret,

$$\hat{\mathbf{D}} = \mathbf{D} \quad \mathbf{f} \tag{8}$$

The additional benefit from defining objective functions is that one can tune the analysis to show covariance specifically with respect to a desired form about which the data are regularized. In the current work, only stationary conditions are considered for quantification of the benefits from wake steering. Stationary conditions are those in which minimal simultaneous variation occurs in all data channels without additional regularization, i.e. setting the function block to $\mathbf{f} = 0$. Wind speed ramps are of particular interest in terms of balancing grid loads and have a significant impact on wind plant response to control schema. Identifying wind ramps in blocks of atmospheric data can be accomplished by assigning the objective function $f_u = c_0 t + c_1$, where c_0 then represents the rate of increase in the observed wind speed. The total variation method for identifying transient events in atmospheric time series data is demonstrated in forthcoming work by Hamilton [12].

Acknowledgments

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 the U.S. Department of Energy O ce of Energy E ciency and Renewable Energy Wind Energy Technologies O ce. 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] M. Abkar and F. Porté-Agel. Influence of atmospheric stability on wind-turbine wakes: A large-eddy simulation study. *Physics of fluids*, 27(3):035104, 2015.

[2] R. Barthelmie, P. Crippa, H. Wang, C. Smith, R. Krishnamurthy, A. Choukulkar, R. Calhoun, D. Valyou, P. Marzocca, D. Matthiesen, et al. 3d wind and turbulence characteristics of the atmospheric boundary layer. *Bulletin of the American Meteorological Society*, 95(5):743–756, 2014.

[3] J. M. S. Bartl, F. V. Mühle, and L. R. Sætran. Wind tunnel study on power output and yaw moments for two yaw-controlled model wind turbines. 2018.

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

[5] M. Bastankhah and F. Porté-Agel. Experimental and theoretical study of wind turbine wakes in yawed conditions. *Journal of Fluid Mechanics*, 806:506–541, 2016.

[6] M. Bastankhah and F. Porté-Agel. Wind farm power optimization via yaw angle control: A wind tunnel study. *Journal of Renewable and Sustainable Energy*, 11(2):023301, 2019.

[7] A. Bossavy, R. Girard, and G. Kariniotakis. Forecasting ramps of wind power production with numerical weather prediction ensembles. *Wind Energy*, 16(1):51–63, 2013.

[8] 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*, 18(12):2135–2143, 2015.

[9] P. Fleming, J. King, K. Dykes, E. Simley, J. Roadman, A. Scholbrock, P. Murphy, J. K. Lundquist, P. Moriarty, K. Fleming, J. van Dam, C. Bay, R. Mudafort, H. Lopez, J. Skopek, M. Scott, B. Ryan, C. Guernsey, and D. Brake. Initial results from a field campaign of wake steering applied at a commercial wind farm – part 1. *Wind Energy Science*, 4(2):273–285, 2019.

[10] P. Gebraad, F. Teeuwisse, J. Van Wingerden, P. A. Fleming, S. Ruben, J. Marden, and L. Pao. Wind plant power optimization through yaw control using a parametric model for wake effects—a cfd simulation study. *Wind Energy*, 19(1):95–114, 2016.

[11] P. M. Gebraad, F. Teeuwisse, J.-W. van Wingerden, P. A. Fleming, S. D. Ruben, J. R. Marden, and L. Y. Pao. A data-driven model for wind plant power optimization by yaw control. In *2014 American Control Conference*, pages 3128–3134. IEEE, 2014.

[12] N. Hamilton. Total variation of atmospheric data: covariance minimization about objective functions to detect conditions of interest. *Atmospheric Measurement Techniques*, 2019 (under review).

[13] Á. Hannesdóttir, M. Kelly, and N. Dimitrov. Extreme wind fluctuations: joint statistics, extreme turbulence, and impact on wind turbine loads. *Wind Energy Science*, 4(2):325–342, 2019.

[14] M. F. Howland, J. Bossuyt, L. A. Martínez-Tossas, J. Meyers, and C. Meneveau. Wake structure in actuator disk models of wind turbines in yaw under uniform inflow conditions. *Journal of Renewable and Sustainable Energy*, 8(4):043301, 2016.

[15] M. F. Howland, S. K. Lele, and J. O. Dabiri. Wind farm power optimization through wake steering. *Proceedings of the National Academy of Sciences*, 2019.

[16] T. Ishihara and G.-W. Qian. A new gaussian-based analytical wake model for wind turbines considering ambient turbulence intensities and thrust coe cient effects. *Journal of Wind Engineering and Industrial Aerodynamics*, 177:275–292, 2018.

[17] T. J. Larsen, H. A. Madsen, G. C. Larsen, and K. S. Hansen. Validation of the dynamic wake meander model for loads and power production in the egmond aan zee wind farm. *Wind Energy*, 16(4):605–624, 2013.

[18] L. A. Martínez-Tossas, J. Annoni, P. A. Fleming, and M. J. Churchfield. The aerodynamics of the curled wake: a simplified model in view of flow control. *Wind Energy Science (Online)*, 4(NREL/JA-5000-73451), 2019.

[19] D. Medici and P. Alfredsson. Measurements on a wind turbine wake: 3d effects and bluff body vortex shedding. *Wind Energy: An International Journal for Progress and Applications in Wind Power Conversion Technology*, 9(3):219–236, 2006.

[20] D. Medici and J. Dahlberg. Potential improvement of wind turbine array e ciency by active wake control (awc). 2003.

[21] NREL. FLORIS. Version 1.0.0, 2019.

[22] J. Park, S.-D. Kwon, and K. H. Law. A data-driven approach for cooperative wind farm control. In *2016 American Control Conference (ACC)*, pages 525–530. IEEE, 2016.

[23] J. Schottler, F. Mühle, J. Bartl, J. Peinke, M. S. Adaramola, L. Sætran, and M. Hölling. Comparative study on the wake deflection behind yawed wind turbine models. In *Journal of Physics: Conference Series*, volume 854, page 012032. IOP Publishing, 2017.

[24] L. Vollmer, G. Steinfeld, D. Heinemann, and M. Kühn. Estimating the wake deflection downstream of a wind turbine in different atmospheric stabilities: an les study. *Wind Energy Science*, 1(2):129–141, 2016.

[25] J. Wagenaar, L. Machielse, and J. Schepers. Controlling wind in ecn's scaled wind farm. *Proc. Europe Premier Wind Energy Event*, pages 685–694, 2012.