Received: 12 August 2016 | Revised: 29 August 2017 | Accepted: 29 August 2017

DOI: 10.1002/we.2143

WILEY

RESEARCH ARTICLE

A comparison of methods for assessing power output in non-uniform onshore wind farms

Andrea Staid[1] | **Claire VerHulst**[2] | **Seth D. Guikema**[3]

[1]Discrete Math and Optimization, Sandia National Laboratories, Albuquerque, NM, USA
[2]Mechanical Engineering, Johns Hopkins University, Baltimore, MD, USA
[3]Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI, USA

Correspondence
Andrea Staid, Discrete Math and Optimization, Sandia National Laboratories, P.O. Box 5800, Mail Stop 1326, Albuquerque, NM, 87185-1326, USA.
Email: astaid@sandia.gov

Funding information
National Science Foundation, Grant/Award Number: OISE 1243482

Abstract

Wind resource assessments are used to estimate a wind farm's power production during the planning process. It is important that these estimates are accurate, as they can impact financing agreements, transmission planning, and environmental targets. Here, we analyze the challenges in wind power estimation for onshore farms. Turbine wake effects are a strong determinant of farm power production. With given input wind conditions, wake losses typically cause downstream turbines to produce significantly less power than upstream turbines. These losses have been modeled extensively and are well understood under certain conditions. Most notably, validation of different model types has favored offshore farms. Models that capture the dynamics of offshore wind conditions do not necessarily perform equally as well for onshore wind farms. We analyze the capabilities of several different methods for estimating wind farm power production in 2 onshore farms with non-uniform layouts. We compare the Jensen model to a number of statistical models, to meteorological downscaling techniques, and to using no model at all. We show that the complexities of some onshore farms result in wind conditions that are not accurately modeled by the Jensen wake decay techniques and that statistical methods have some strong advantages in practice.

KEYWORDS

onshore wind farms, statistical modeling, turbine wakes, wind power assessment

1 | INTRODUCTION

Wind energy is a fast-growing segment of the energy sector worldwide. As wind starts to play a larger role in our electricity systems, there is increasing interest in using it efficiently. Wind power is, for the most part, not a dispatchable form of energy; forecasts are not perfect, and an individual wind farm cannot be set to a desired production level unless the wind conditions are sufficient for that level of power. Managing a farm's variability and intermittency is still a big challenge, but recent research shows promising results in optimizing farm production.[1-4] Knowledge of the amount of power that a wind farm will produce in a given hour, day, or year will improve the way that the farm is operated and financed. Perfect knowledge of future wind power production is impossible, but recent research has made great strides in estimating these quantities. Power estimation and prediction is important for multiple applications within the timeline of a wind energy project. In the initial planning stages of a farm, one must develop accurate estimates of what the wind conditions will be once the farm is built. This typically involves the installation of meteorological towers (or met towers) for data collection. This information is used to determine the turbine layout and alignment within the farm that will best capture the available energy. Initial estimates of long-term wind conditions are used to assess the financial viability of a wind farm. Thus, the decision of farm design and resource assessment is critical, as the return on investment depends on the ability of the wind farm to generate enough energy to recoup costs and remain profitable.

Improved farm planning has spurred research on the interactions among turbines within a farm, as this can be used to optimize the placement, spacing, and number of turbines. Wakes from operational wind turbines result in lower wind speeds behind the blades, since some of the energy has been extracted from the ambient flow. This results in lower power production for other turbines sitting behind a leading turbine in its wake. Characterizing these wake effects is critical for choosing turbine layouts and developing accurate estimates of farm power production.

Wind speed is the greatest determinant of wind power, but turbine wake interactions complicate the simple translation of speed to power. Wake models can be used to calculate the winds seen throughout a farm given known input conditions, and the turbine power curve can then be

 Copyright © 2017 John Wiley & Sons, Ltd.

used to calculate the power produced by each turbine in the farm. This can also be done using statistical models or data-mining techniques if a farm has been operational long enough, but this usually is not the case in planning contexts.

Much recent research has focused on offshore wind installations, leading to improvements in techniques and methods for offshore applications. However, onshore wind farms will still play an important role in future energy systems, and several factors are likely to differ when analyzing a wind farm that is onshore instead of off. Offshore sites see steadier wind; onshore locations typically have higher turbulence with more mixing and less dramatic wake losses.[5] Onshore farms also have local geography to take into account. Terrain can strongly influence the wind flow in a region. In addition, offshore farms are often built with a uniform grid of turbines in evenly spaced rows and columns. This is not always the case onshore where there are significant constraints to turbine placement, such as roads, houses, hillsides, or protected lands. In theory, the models used to analyze wind farms can easily be modified to account for these factors. In practice, however, much of the research and model validation has recently been conducted on offshore (and occasionally very near-shore) farms.[5-8] The model results do not always match up with reality in onshore farms and therefore should be used carefully so as not to promote false assumptions in the industry.

This manuscript aims to take an analytical look at the performance of various techniques for analyzing wind farm power production in real-world settings. This presents challenges in terms of data availability and managing sources of uncertainty, as it is not often known which uncertainties are present in a real system. We choose to focus on simple, easy to use methods that could feasibly be used to assist decisions in a planning context. This is very different from other research settings in which many uncertainties are purposely controlled to discover the exact nature of the physical relationships. Here, we deal only with the data that a wind farm planner or operator would have available to them, and we highlight the limitations that this presents.

In contrast to much of the existing research into wakes and turbine interactions, the models presented here are applied to onshore farms that are not aligned in a uniform grid. The findings show that validation done in offshore settings does not always translate well when applied in these contexts. We focus on the performance of a simple wake model and compare it to other alternatives for estimating wind farm power production. This work presents a comparative analysis of several different techniques for estimating wind farm power production. It is not meant to be a comprehensive study of all available methods; there are many different wake decay models that could be used in addition to the Jensen wake model used here.[9] We chose to use the Jensen model because of its popularity, simplicity, and accuracy in offshore settings. There are a number of studies that compare different types of wake models, and the Jensen model has been found to perform well in a variety of settings, primarily offshore.[6,10] This research focuses on the level of accuracy achieved for onshore wind farms instead.

1.1 | Background

Estimates of wind farm power production are needed for many types of planning problems. Most power production estimates are used for short-term tasks such as load scheduling and dispatching, and these are usually conducted for an individual wind farm or a small group of farms within a region.[11,12] There is a sizable body of research on predicting power output for existing wind farms, and the methods are varied.[11,13] Generally, power forecasting models combine wind speed forecasts and farm-specific data, such as turbine characteristics, farm layout, and past generation data to generate power forecasts for that farm.

In practice, most of the work done on predictions for wind energy applications has been focused on generating accurate wind speed forecasts. As a result, there have been many advancements over the years in wind speed forecasting, both from a physical and statistical standpoint.[14-16] Many models focus on wind speed only, as this is obviously the most important factor in determining the amount of wind power that will be produced. Other models take a wind speed forecast as input and focus on improving the power estimate itself. See Giebel et al[17] for an overview. Even with a perfect wind speed forecast, translating from wind speed to wind power for an entire farm is not as straightforward as many would like.

Despite the growing focus on wind speed predictions, the subsequent conversion to wind power is often highly simplified. In fact, a lot of research does not distinguish between forecasts for wind speed and power, and this is generally acceptable for predictions regarding just one turbine.[18] The power curve of an individual turbine is well understood, and for the most part, an accurate wind speed is the only piece of information needed to then determine the power output. However, the process is much more complex when dealing with an entire farm. Turbine wakes strongly affect the wind at downstream turbines. Wake effects depend on a number of factors, including farm layout, terrain, wind direction, wind speed, turbulence intensity, and atmospheric stability.[19] The performance of an entire farm can differ greatly based on the wind conditions at a given time. With reliable wind speed forecasts and detailed farm data, power estimates can be fairly accurate.[20]

In practice, the wind speeds seen by each individual turbine in a large wind farm are rarely known exactly. Turbine wake interactions and complex terrain (for onshore farms) result in different speeds at each turbine throughout a farm. There has thus been a strong focus on advancing the modeling techniques to better understand and represent fluid flow within a wind farm. Even with an accurate forecast for wind speed, conversions from wind speed to power are typically done by heuristics in practice, such as simply adding up the individual power curves of each turbine and applying a penalty to account for wakes.[21] Although these estimates are far from perfect, they are still useful for many large-scale planning decisions. Long-term planning and integration studies need estimates of wind production for virtual wind farms, and these estimates use a lot of simplifying assumptions about actual farm behavior.[22,23] Often the turbines are aggregated, despite the acknowledgement that deviations from the farm power curve are common and can be caused by a number of factors.[22]

Farm-level power estimation is difficult to do in a large wind farm or when planning for the production of a new wind farm. A site with met towers can provide information on wind conditions, but the resulting power production depends on the turbine layout. Many tools exist for turbine layout

planning. The goal of these tools is to model the farm as it would exist in an operational setting. It is therefore important to capture the wake effects of the turbines in the farm and the terrain effects of surrounding geography. Many of these models rely on computational fluid dynamics to model fluid interactions among turbines, but fast and simple analyses using the Jensen model, for example, are often sufficiently accurate.

The Jensen model is a basic wake decay model that has become the standard baseline for research and industry.[9] It is based on a simple, geometric calculation of the wake expansion behind a turbine and can be used to calculate the velocity deficit relative to the freestream at any point behind a turbine. It has been used as a point of comparison in countless experiments (sometimes referred to as the Park or Jensen/Park model). See previous studies[6,7,10,24] for a selection of examples. There have been advancements in wake modeling that move beyond the Jensen model, especially as large eddy simulations of large turbine arrays became computationally feasible. However, the Jensen wake model is still used frequently with high degrees of accuracy.[6] There are some known weaknesses of the Jensen model due to its simplicity. The model focuses on individual wakes, and for very large farms, this can result in an underprediction of wake losses.[5] Even in onshore settings, the Jensen model has been shown to perform relatively well under limited conditions.[25] However, there may be other factors that complicate the application in large, onshore settings when wind conditions throughout the farm cannot be known. Still, the Jensen model serves as a strong benchmark measure, and it is used here to estimate the wake deficits, and therefore power production, within farms.

2 | DATA

We have wind farm production data from 2 different onshore farms in the United States. The details of the farms must remain anonymous, and they will be referred to throughout as Farms 1 and 2. Farm 1 is located in the south-central region of the country and is made up of 100 turbines with a rotor diameter of 77 m and a hub height of 80 m. The turbines have a cut-in wind speed of 3.5 m/s, a cut-out wind speed of 25 m/s, and a rated wind speed of 14 m/s. Farm 1 sits in a relatively flat geographical area in the midst of agricultural land. Farm 2 is located in the north-central region of the country and is made up of 140 turbines (of the same make and model as Farm 1.) It also sits in a flat geographical area in the midst of agricultural land. The turbine layout of Farms 1 and 2 can be seen in Figures 1 and 2, respectively. Both farms have 2 met towers installed, the locations of which are also shown in Figures 1 and 2. Included in these figures are the wind roses for each of the farms, showing that the predominant wind direction for Farm 1 comes from the south, and for Farm 2, the predominant winds are from the SSW. Thus, the met towers see freestream wind flow under most operating conditions. Under some wind directions, the met tower measurements will represent wind flow within turbine wakes, and we have no way of knowing the freestream conditions seen by the wind farms. This adds uncertainty to the assessment of power output, but these conditions are representative of those faced by wind farm operators and are important to still analyze to represent the real-world decision space. Turbine spacing varies significantly across the farm and with varying wind direction. The minimum turbine spacing is approximately 3 and 3.6 turbine diameters for Farms 1 and 2, respectively.

The available data from Farm 1 runs from November 2010 to October 2013 and Farm 2 from November 2010 through October 2014. Individual turbine data include 10-minute averaged power production, wind speed (measured at nacelle), availability, and curtailment. Availability and curtailment data are 0-1 metrics representing the portion of the 10-minute period that a turbine was either available or curtailed, respectively. The met tower data consists of 10-minute averaged wind speed, direction, temperature, and, for Farm 2 only, air pressure.

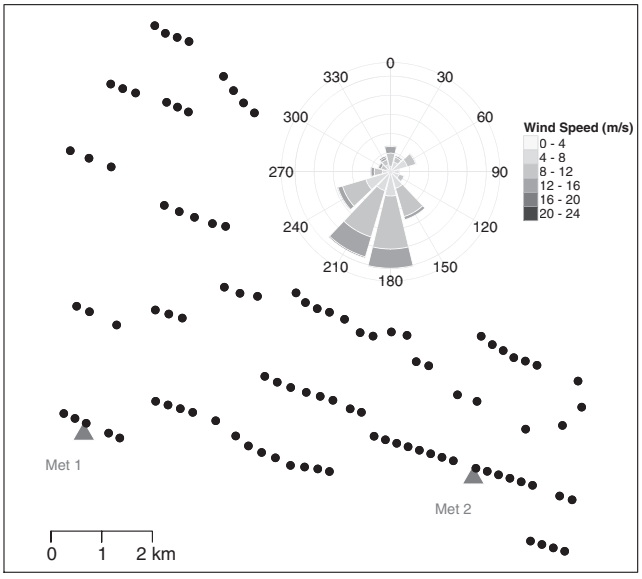

FIGURE 1 Turbine layout and meteorological tower locations for Farm 1

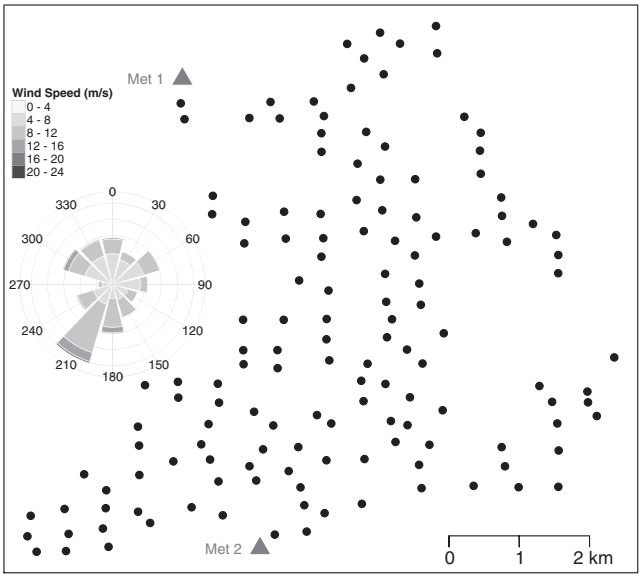

FIGURE 2 Turbine layout and meteorological tower locations for Farm 2

Our analysis is specific to the data used. The findings represent 2 wind farms, although they are fairly representative of other large onshore wind farms in the United States. Ideally, all results would be validated with additional farms, but it is often difficult to gain access to production data from wind farms and this is left as a future research area should more data become available.

The data for both farms were cleaned, filtered, and corrected to ensure accuracy. Details can be found in the Supporting Information that accompanies this manuscript. The exercise highlights the importance of ensuring data quality when working with real data, which is often subject to errors due to measurement, calibration, or recording.

3 | IMPLEMENTING THE JENSEN MODEL

The Jensen model takes 2 input parameters: the wake decay coefficient and the thrust coefficient of the wind turbines. Use of the model when analyzing a real wind farm, however, requires that certain assumptions be made. The Jensen model can be used to calculate the velocity deficits in the turbine wakes, and this calculation is based on steady-state wind. For one constant wind direction, the deficits are calculated as a snapshot in time. Thus, the Jensen model is best suited for large-scale resource assessment and is not the best option for evaluating real-time or rapidly changing wind power production. For the cases presented here, the data exists as 10-minute averages and, for some of the analysis, has been filtered to represent only conditions that are as close to steady state as possible by using only the datapoints where both met towers in the farm agree in the measured wind direction. The Jensen model should work well under these limitations. Although other filters could be applied to further restrict the relevant data, we chose not to in order to ensure that the results are still representative of the actual farm behavior. A discussion of the relevance of wake models for evaluating both farms can be found in the Supporting Information.

The Jensen model determines the effect that a turbine wake will have on a downstream turbine. The wake velocity deficit at a downstream turbine relative to the velocity seen by an upstream turbine is determined by

$$\delta V = \frac{1-\sqrt{1-C_T}}{(1+2kx)^2}\left(\frac{A_o}{A}\right),$$

where C_T is the turbine thrust coefficient, k is the wake decay coefficient, and x is the downstream distance at which the deficit is measured. For a turbine that sees only a partial wake from an upstream turbine, a correction is applied based on the fractional area of overlap, A_o, and the swept rotor area, A. For a large wind farm with multiple turbines, the deficit seen by any one turbine is a combination of all of the upstream wakes that interact with that turbine. This results in the following formulation that gives the overall velocity deficit as a result of the superposition of multiple upstream wakes as given by Katic et al.[26]

$$\frac{u}{u_0} = 1 - \sqrt{\sum_{j\in J}\left(\delta V_j\right)^2},$$

where J is the set of upstream turbines with wakes impacting a given downstream turbine. The choice of the parameter k determines the rate at which the wake decays behind a turbine, and it is dependent on the ambient turbulence, turbine-induced turbulence, and atmospheric stability. Atmospheric stability has a strong influence on wake behavior, and it should determine the choice of k.[8] Namely, stable conditions result in a slower

wake recovery and therefore larger velocity deficits at downstream turbines. There is consensus on a value of 0.075 for onshore applications and lower values (ie, 0.04 or 0.05) for offshore applications to reflect the lower turbulence generally found offshore or near-shore.[27] The thrust coefficient, C_T, comes from a turbine-specific curve that is a function of incoming wind velocity. A typical curve has the highest C_T value at low wind speeds (ie, $C_T = 0.9$ for winds 5 m/s) and then drops off as wind speed increases (ie, to a value of 0.2 above 20 m/s.)

We conducted a sensitivity analysis on the choice of the 2 parameters in the model, and the details are discussed in the Supporting Information. With set values of k and C_T, the velocity deficits can be evaluated at each turbine in a farm for a given wind direction and compared to the actual data averaged over that same wind direction. The actuals were determined by applying the inverse turbine power curve to the power values measured at each turbine, and this process introduces some error as the wind speed cannot be known exactly at rated power values. The difference between the actual data and the Jensen deficit calculations for a wind direction of 180° are shown in Figures 3 and 4 for Farms 1 and 2, respectively. The Jensen models used here were run with $k = 0.075$ and $C_T = 8/9$. The actuals were calculated using a subset of the farm data that most closely resembles the conditions used in the Jensen model. A subset of farm data was used containing only those datapoints with wind directions that fall between 180°±5° for *both* met towers at both farms, as there is sufficient data from this direction for both farms. Since the 2 towers are spaced fairly far apart in both farms, agreement in the wind direction indicates a relatively steady flow. Even with the actual data reduced to the "best" cases, there is a large discrepancy between the Jensen wind speeds and the actual data for both farms. The leading turbines in the Jensen model all have a normalized velocity equal to 1, by definition. The turbines further back in the farm see lower velocities as the wakes from upstream turbines show their influence. In the actual farm data, however, the results deviate significantly from what is expected. Instead of seeing a trend toward lower

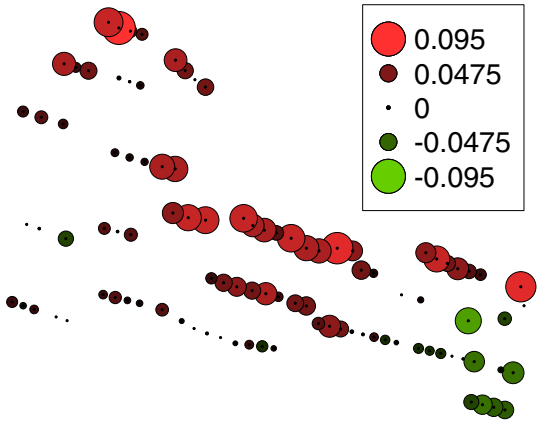

FIGURE 3 The average difference between the actual data and the Jensen model estimates for wind velocity deficit relative to leading turbines for Farm 1. The wind direction is from the south (180°), and the wind speed deficit is relative to the leading turbines on the south side of the farm. The plot shows the actual data minus the Jensen estimates, so positive values are underestimates of the actual wind speeds whereas negative values are overestimates [Colour figure can be viewed at wileyonlinelibrary.com]

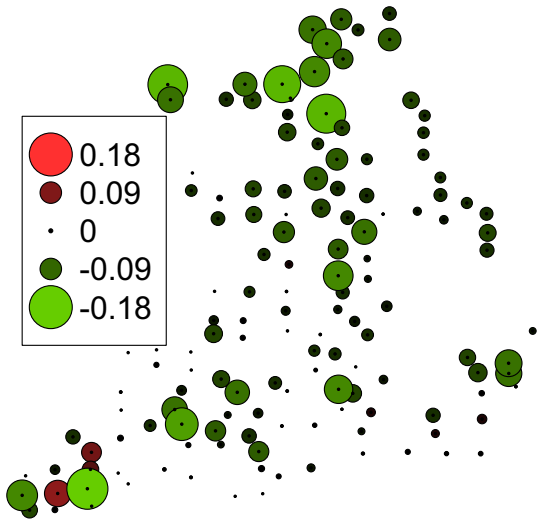

FIGURE 4 The average difference between the actual data and the Jensen model estimates for wind velocity deficit relative to leading turbines for Farm 2. The wind direction is from the south (180°), and the wind speed deficit is relative to the leading turbines on the south side of the farm. The plot shows the actual data minus the Jensen estimates, so positive values are underestimates of the actual wind speeds whereas negative values are overestimates [Colour figure can be viewed at wileyonlinelibrary.com]

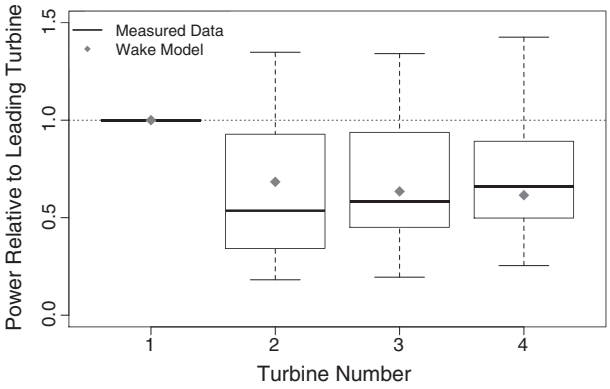

FIGURE 5 Comparison of wind velocity along a line of turbines. Shown here are the wind speeds (relative to Turbine 1) for Turbines 1-4 in Farm 1, which are closely spaced in a straight line, when the wind is aligned with the line of turbines ±2°. The Jensen model shows the expected drop in wind speeds behind each turbine, and the actual farm data shows similar behavior on average, even if there is a large spread in the data. The standard errors for the relative power of turbines 2, 3, and 4 are 0.071, 0.084, and 0.084, respectively

wind speeds at the back of the farm in Farm 1, some of the rear turbines have much higher speeds than expected. Although most of the turbines have fairly small deficit errors, Farm 1 tends to have higher actuals than the Jensen model predicts, while Farm 2 tends to have lower actuals.

There is not enough data to determine the exact cause of the mismatch between the Jensen model predictions and the actual farm data. High turbulence or atmospheric instability are likely at play, although the extent of their influence is not known. The farms are quite large, spanning more than 6 km in some directions. Two measurement points (ie, the 2 met towers) may simply not be enough to provide information on the conditions present throughout the entire farm. The surrounding areas of both farms are flat and homogeneous by many standards, but small deviations could result in different conditions far from the met towers that would not show up in the available data. As given, the Jensen model does not capture the dynamics of these 2 farms and results in significant errors for production estimation using the given data.

This mismatch can be partially explained by the complex wake structure and propagation that is likely present in the farms. For the case of Farm 1, another study has measured these flow structures using Doppler radar, and the wake structures present are driven by short-term wind transients and terrain features.[28] This behavior has been shown to result in high momentum channels throughout the farm, and this could explain the discrepancies observed. Turbines exposed to these higher-speed channels further back in the farm produce more power than otherwise would be expected, and in some cases, they produce more than the leading turbines in the farm.[28] This phenomena is observed in the actual data from Farm 1.

The Jensen model can also be used to look at some standard benchmarking performance measures for wind farms. In a limited context, we can evaluate whether the Jensen model is capturing the localized effects expected in a wind farm. In closely-spaced lines of turbines, the wind speed typically drops off significantly after the leading turbine and may sustain smaller decreases in the downstream turbines.[5] This behavior is also predicted by the Jensen model and is expected to be seen in any line of turbines subjected to oncoming winds parallel to the line. Figure 5 shows the comparison of the expected behavior of the Jensen model and the actual farm data. There are not many datapoints for this particular wind direction (105°), but there are enough to show that the wind farm does behave as expected in a limited setting with simple turbine interactions.

4 | ALTERNATIVE METHODS FOR POWER PREDICTION

The Jensen model fails to capture the true dynamics of either Farm 1 or Farm 2 regarding wind speed deficits and farm power production. It cannot be used as an accurate tool for either resource assessment or power prediction for the case studies presented here. We compare the accuracy of alternatives to wake models for both resource assessment and power prediction. Statistical models have been used for short-term wind power forecasting as well as wind farm production characterization (ie, such as developing a farm-level power curve[20]). For resource assessment, the National Renewable Energy Laboratory (NREL) has a database of potential wind farm sites and virtual production data, which will also be used to compare power production estimates from the 2 farms.

4.1 | Statistical models

Statistical models have no need for explicit information about the physical details of the farm, atmosphere, surface roughness, or turbulence. Unlike wake-based models, any relationships are intrinsically captured in the existing data when training the model. Statistical models have limitations; they are dependent on the data available for training and cannot always effectively extrapolate beyond the training data. Some models are heavily influenced by outliers, leading to erroneous estimates.

Generalized linear models, generalized additive models, random forests, and multivariate adaptive regression splines were the models initially tested on these wind farms for their predictive accuracy. Since the nonparametric models (random forest and multivariate adaptive regression

TABLE 1 Summary statistics for the statistical model input data

		Minimum	Median	Mean	Maximum
Farm 1	Wind speed (m/s)	0.8	9.0	9.0	21.1
	Wind direction (°)	0.0	191.1	187.2	360.0
	Temperature (°C)	−14.4	19.1	18.3	43.1
	Minimum turbine spacing (rotor diam.)	2.9	18.5	19.8	121.4
Farm 2	Wind speed (m/s)	1.7	7.8	8.1	21.4
	Wind direction (°)	0.0	196.1	185.3	360.0
	Temperature (°C)	−18.7	7	8.7	35.6
	Minimum turbine spacing (rotor diam.)	3.3	13.0	15.5	139.2

TABLE 2 Farm-level training and testing paradigms tested

Model	Train	Test
$Farm1_{self}$	70% of Farm1 data	30% Farm1 data
$Farm1_{other}$	All Farm2 data	All Farm1 data
$Farm2_{self}$	70% of Farm2 data	30% Farm2 data
$Farm2_{other}$	All Farm1 data	All Farm2 data

splines) resulted in much lower errors than the parametric models tested (generalized linear models and generalized additive models), we further tested nonparametric models, comparing their performance using out-of-sample test data under different training conditions, including varying the number of trees in the case of the random forest models. While the 2 model types both performed well, the random forest model (implemented in R using the package RandomForest[29,30]) consistently outperformed the others with lower mean absolute and root mean squared errors. For each farm, several versions of a random forest model were created, both on the turbine level and the farm level. For the turbine level, a separate model was created for each individual turbine in the wind farm either by training the model using only data from the met towers or with met tower data and turbine layout data, denoted T_{met} and T_{layout}, respectively, in Tables 3 and 4. The met tower data consisted of wind speed, direction, and temperature. The layout data for each turbine contains the streamwise distance and angle relative to any upstream turbines within a symmetric 30° cone pointing upstream. These values change with wind direction, and the resulting information is used in the statistical model as a proxy for possible wake interactions at a given turbine. Models incorporating seasonal and diurnal effects were also tested, but the resulting errors were generally higher, with changes in mean absolute error ranging from a 7% increase to a 0.4% decrease, and they are not presented here. Summary statistics for the data used in the statistical models is shown in Table 1. All upstream turbine distances were included, but only the most critical distance parameter, the minimum, is shown here.

The farm-level models were implemented with 2 different training methods for each farm, laid out in Table 2. We used each statistical model to predict power production for both farms, meaning that we trained a model on Farm 1, and used it to predict both Farms 1 and 2 power production, and vice versa. This was done as a test of the ability of these models to generalize across wind farms, especially wind farms that are fairly similar in terms of terrain and geography, as these two are. Generalizable statistical models are promising for planning contexts when sufficient measured data are not available or too expensive to obtain. For the models trained using the same farm data, a holdout analysis was also performed to assess the consistency of the results. The training sample is drawn randomly in each case, using 70% of the available data to predict the remaining 30% of the observations. For the models trained on the opposite (other) farm, the entire dataset was used. The models trained on one farm and used to predict the other lack the intrinsic farm-specific information usually captured. Specific relationships between wind direction differences and farm power could imply turbulence in one farm that affects power production. That may not be the case in the other farm, and farm-level power curves have been found to be specific to an individual farm.[20] Nevertheless, if such a model can provide a reasonable degree of accuracy, this technique can go a long way toward assessing wind production for farms that may not exist yet.

4.2 | NREL wind integration national dataset

The NREL provides an open-source database of potential wind sites that can be used for resource assessment, planning, and integration studies. The dataset consists of 112 471 onshore grid cells that contain up to 8 virtual turbines each. Each turbine is assumed to be 2 MW with 100-m hub height. The power curve used in each grid cell depends on the assessed wind class and is made up of a number of representative manufacturer power curves for similar turbines.[21] Grid cell power production is estimated using the appropriate power curve and wind speeds from meteorological reanalysis data. Instead of explicitly modeling wake effects, an empirical correction is applied to reduce the wind speed to account for the presence of wakes. The power curve is then applied based on these corrected wind speeds. The correction is

$$C_{wake} = 1 - \frac{1}{20}\left(\frac{n_{turbines} - 1}{7}\right),$$

where $n_{turbines}$ is the number of turbines in the grid cell, and with a maximum of 8, the wake reduction is at most 5%. C_{wake} is multiplied by the reanalysis wind speeds, resulting in corrected grid cell wind speed for all turbines.[21]

The dataset includes multiple grid cells within the areas of both Farms 1 and 2. For the most accurate comparison between the actual data and the NREL dataset, we chose the individual grid cells located closest to the 2 met towers in each of the farms. The 2 sites in Farm 1 closest to met towers 1 and 2 contain 6 and 8 turbines, respectively, and use the power curve for class 1 winds. The 2 sites in Farm 2 closest to met towers 1 and 2 contain 8 turbines each and use the power curve for class 3 winds. The power data from these grid cells are normalized based on total grid capacity and then used as the predictions for comparison to the actual farm data. Predictions are used individually from each met location and as an average of the two grid cells' production in each farm, incorporating wind conditions from locations spread across the area of the farms.

4.3 | Aggregated power curves

If wakes are ignored with steady wind conditions, it can be assumed that wind speeds are consistent across the farm. Therefore, a farm-wide power prediction can be made using no model at all, but by aggregating individual power curves across the farm using met tower wind data. This wind speed was then translated into power using the turbine power curve, and this power value is multiplied by the number of turbines to get the farm-level power. This was done using both met towers 1 and 2 data separately and then using the mean value predicted from both. This aggregated method is the simplest to implement; only the number of turbines is relevant. In theory, this method overestimates power production since wake effects are ignored.

5 | PREDICTION ERRORS

Using each of the methods previously mentioned, we normalized all predictions based on the installed farm capacity so that errors can be compared across farms and across methods. Comparison results for Farms 1 and 2 are shown in Tables 3 and 4, respectively. The tables show the mean absolute error and root mean squared error for statistical models, Jensen wake models, aggregated predictions (ie, no model), and the NREL dataset predictions. All predictions have been normalized by total capacity, so the errors are in the units of normalized capacity as well. A 20-replication holdout analysis was performed for some of the statistical models. For these, errors reported in Tables 3 and 4 represent the mean value across all holdouts; the number in parenthesis is the standard error across holdouts. There is little variation across holdouts, showing that the models are robust in their predictive accuracy. For the statistical models, T refers to models created for each turbine individually and then aggregated; F refers to models trained on the farm as a whole. Subscripts *met* and *layout* refer to models trained using met data only and both met and farm layout data, respectively. Subscripts *self* and *other* refer to models that were trained on Farm 1 to predict its own power and trained on Farm 2 to predict Farm 1's power, respectively. Mean values are the errors when the predictions from both met towers 1 and 2 were averaged, and the average used as the prediction.

The statistical models have the lowest errors overall. There is no advantage to be gained by including information about the layout of the farm; using the met data alone outperforms models that incorporate data on upstream turbines. Most surprisingly, errors stay low when predicting for the *opposite* farm, which the models were not trained on. Statistical models perform well estimating power production on an independent farm with no farm-specific data. The Jensen models have already been shown to disagree with the actual data, and this is reflected in high error values. However, there is some improvement gained by averaging predictions from 2 different Jensen models using the data from each met tower separately. There is often poor agreement between the 2 met towers, and the errors tend to cancel out, resulting in a lower combined error. The combined errors are still higher than the statistical model errors for Farm 2, but the mean Jensen predictions in Farm 1 fall below the level of the farm-level statistical model errors when predicting the opposite farm.

It is surprising that the errors when predicting performance on the opposite farm are as low as they are. The reasons for this are likely due to the complexities of modeling fluid flow within a wind farm. With only 2 met towers for a large farm, it is impossible to know the incoming wind for

TABLE 3 Normalized Farm 1 power prediction errors

	Statistical models				Jensen model			Aggregated turbines			NREL
	T_{met}	T_{layout}	F_{self}	F_{other}	Met1	Met2	Mean	Met1	Met2	Mean	Mean
MAE	0.050 (0.0008)	0.053 (0.0008)	0.050 (0.0006)	0.082	0.095	0.101	0.079	0.095	0.095	0.080	0.156
RMSE	0.074 (0.0021)	0.077 (0.0021)	0.076 (0.0015)	0.110	0.140	0.144	0.114	0.139	0.139	0.113	0.217

TABLE 4 Normalized Farm 2 power prediction errors

	Statistical models				Jensen model			Aggregated turbines			NREL
	T_{met}	T_{layout}	F_{self}	F_{other}	Met1	Met2	Mean	Met1	Met2	Mean	Mean
MAE	0.051 (0.0006)	0.051 (0.0006)	0.052 (0.0007)	0.087	0.127	0.111	0.099	0.116	0.116	0.085	0.227
RMSE	0.072 (0.0010)	0.070 (0.0009)	0.073 (0.0012)	0.114	0.160	0.147	0.128	0.148	0.148	0.112	0.295

Abbreviations: MAE, mean absolute error; NREL, National Renewable Energy Laboratory; RMSE, root mean squared error.

all of the leading-edge turbines. Statistical models can intrinsically learn wind farm behavior from historical data, and they are then able to apply the same relationships to other, complex farms in similar terrain. It would be interesting to study whether this performance advantage still holds as the farms being compared become less similar in terms of number of turbines and geographical landscape. Unfortunately, we do not have sufficient data available to test this, and it is left as an interesting avenue for future research.

The third method, aggregating individual turbine power curves (thus ignoring any wake effects), results in errors comparable to the Jensen model. Again, the errors are reduced by averaging the predictions from each met tower, and these combined error measurements are similar to those obtained by the farm-level statistical model trained on the opposite farm. The NREL dataset predictions had the highest errors out of all the models tested and do not provide a strong estimate for farm power production at these locations. The individual site errors from the NREL data were similar to that of the mean value, and they are not included in the tables.

The high NREL errors are largely due to the added layer of uncertainty present in this data. The wind speeds come from reanalysis data, and they are downscaled from large-scale weather models. They might be accurate for capturing the dynamics of large weather patterns, but it is difficult to model weather at small spatial scales.[31] Wind in particular is difficult to model locally; complex terrain complicates this further.[32] To support this point, correlations between the temperature values in the NREL dataset and the measured temperature at the met towers ranges from 0.97 to 0.98. Correlations for wind speeds, on the other hand, range from 0.66 to 0.69. This discrepancy in local wind speeds translates into discrepancies in power estimates. We included this comparison to demonstrate the need for local wind condition data when making planning decisions, as large-scale estimates can lead to large local errors.

6 | DISCUSSION

Statistical models achieve the lowest errors among the methods tested for estimating wind farm power production; errors are low even when the models are not trained on the same farm that they are predicting, and this finding provides a promising result for wind farm planning moving forward. The 2 farms tested are similar in that they are both situated in flat, agricultural areas, but they differ in location, turbine layout, and wind conditions. Existing literature suggests that turbine layout is critical to characterizing farm power production because of wake effects. The Jensen model explicitly models turbine wakes, but the complexities of real farms are not captured with commonly available data. There are a number of issues present when comparing wake models to actual measurements, including measurement errors, accounting for wake transport time, natural fluctuations of speed and direction, and time averaging.[5,7] In addition, the actual wind conditions within an onshore wind farm with even minor terrain variation are more complex than the steady input flow used in the Jensen model. Potential sources of error are plentiful, and we can speculate as to the sources. A thorough investigation into the causes of these errors is left as an avenue for future research. Met tower data issues are one suspect (calibration and location, if the met towers are not in freestream flow), but the natural variation of wind throughout such a large spatial area could also result in a significant mismatch between what the Jensen model predicts and the actual data.

There are multiple sources of uncertainty in the analyses presented here, and further research would be needed to elucidate the exact nature of the uncertainty and its effect on the results. However, our careful analysis leads us to believe that the largest impact is due to the lack of knowledge of wind conditions within such large wind farms. Two met towers are not sufficient to determine the nature of the wind flow seen by each turbine within the farm, and this is a critical parameter for the Jensen wake model. The results presented here are designed to represent data availability in real-world wind farm contexts, where the uncertainties often cannot be reconciled as they would in a controlled lab (or wind tunnel) setting. The Jensen model is known to perform well if the wind conditions are known, but we show that the data used for decision making is not sufficient to determine upstream wind conditions in large spatial areas. In these settings, other methods are preferable.

Statistical models do not explicitly model turbine wakes, but they have the ability to intrinsically capture relationships among turbines that are present in the data. If the wind farms have unique but consistent behavior patterns that may be driven by wake effects, statistical models can capture these relationships. It is especially surprising that statistical models maintain accuracy when predicting power production for another farm, since any unique wake behavior should not carry over. Still, the results have important implications. It is likely that the geographical similarities are enough to result in accurate cross-farm statistical models even with different layouts. In addition, for rough estimates of farm power production, one can argue that no model is needed at all. Aggregating power curves of a number of turbines is as accurate, or better than, the Jensen model. No farm model nor data on turbine positions is needed. This method can be used on any farm, even one that is not yet built. Reasonable estimates can be obtained with just a met tower installed.

Wake models are heavily tested on offshore wind farms with uniform turbine alignment. Techniques are continually improving and accurately capture the wake behavior in a number of case studies. Any user of these models, however, should be careful when applying them to onshore sites. The 2 case studies presented here demonstrate that onshore farms do not always match wake-model behavior.

Luckily, alternatives do exist, and statistical models can be used for power prediction if appropriate training data is available. For resource assessment when no training data is available, it may be advantageous to use a simple summation to aggregate the turbine power curves. While farm layouts should be chosen carefully, real-world effects can result in drastically different behavior than what a model may predict, and the models should be used to inform but not determine a farm's design. Statistical techniques are simple to implement, and we believe that they can play an important role in wind farm planning going forward. With additional data and farms for testing and validation, wind developers can begin to expand the regimes in which statistical models can be used to improve planning decisions for a wider spectrum of farm types and locations.

The results presented here also bring up practical questions about the value of wind forecasting. Accurate wind forecasts are highly sought after and are used to make operating decisions regarding wind integration in power systems. It has been shown, however, that even with a perfect forecast (ie, the wind conditions are known exactly at the met towers in real time), predictions of power produced still contain errors. Even with a perfect forecast, translating wind speed to power is nontrivial.

ACKNOWLEDGEMENTS

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc, for the US Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. This work was conducted with the support of the WindINSPIRE Project, National Science Foundation grant OISE 1243482, and most of the research was done while the authors were at Johns Hopkins University. The authors would like to thank our data provider, who wishes to remain anonymous, for allowing us access to this rich data set.

ORCID

Andrea Staid http://orcid.org/0000-0001-6125-5814

REFERENCES

1. DeCesaro J, Porter K, Milligan M. Wind energy and power system operations: a review of wind integration studies to date. *Electricity J*. 2009;22(10):34-43.
2. Mirzaei M, Göçmen T, Giebel G, Sørensen PE, Poulsen NK. Turbine control strategies for wind farm power optimization. American Control Conference (ACC), 2015 Chicago, IL: IEEE; 2015:1709-1714.
3. Muyeen SM, Takahashi R, Murata T, Tamura J. A variable speed wind turbine control strategy to meet wind farm grid code requirements. *IEEE Trans Power Syst*. 2010;25(1):331-340.
4. Hansen AD, Sørensen P, Iov F, Blaabjerg F. Centralised power control of wind farm with doubly fed induction generators. *Renewable Energy*. 2006;31(7):935-951.
5. Barthelmie RJ, Hansen K, Frandsen ST, et al. Modelling and measuring flow and wind turbine wakes in large wind farms offshore. *Wind Energy*. 2009;12(5):431-444.
6. Peña A, Réthoré PE, Laan MP. On the application of the Jensen wake model using a turbulence-dependent wake decay coefficient: the Sexbierum case. *Wind Energy*. 2015.
7. Gaumond M, Réthoré PE, Ott S, Peña A, Bechmann A, Hansen KS. Evaluation of the wind direction uncertainty and its impact on wake modeling at the Horns Rev offshore wind farm. *Wind Energy*. 2014;17(8):1169-1178.
8. Hansen KS, Barthelmie RJ, Jensen LE, Sommer A. The impact of turbulence intensity and atmospheric stability on power deficits due to wind turbine wakes at Horns Rev wind farm. *Wind Energy*. 2012;15(1):183-196.
9. Jensen NO. *A Note on Wind Generator Interaction*: Risø National Laboratory; 1983.
10. Beaucage P, Brower M, Robinson N, Alonge C. Overview of six commercial and research wake models for large offshore wind farms. Proceedings of the European Wind Energy Associate (EWEA). 2012.
11. Landberg L, Giebel G, Nielsen HA, Nielsen T, Madsen H. Short-term prediction—an overview. *Wind Energy*. 2003;6(3):273-280.
12. Pease J. Challenges with the integration of large scale wind by a regional utility, Bonneville Power Administration; 2006. https://www.bpa.gov/Power/PGC/wind/Challenges_of_large_scale_wind_integration_paper.pdf.
13. Costa A, Crespo A, Navarro J, Lizcano G, Madsen H, Feitosa E. A review on the young history of the wind power short-term prediction. *Renewable and Sustainable Energy Reviews*. 2008;12(6):1725-1744.
14. Kavasseri RG, Seetharaman K. Day-ahead wind speed forecasting using f-ARIMA models. *Renewable Energy*. 2009;34(5):1388-1393.
15. Sloughter JM, Gneiting T, Raftery AE. Probabilistic wind speed forecasting using ensembles and Bayesian model averaging. *J Am Stat Assoc*. 2010;105(489):25-35.
16. Staid A, Pinson PD, Guikema S. Probabilistic maximum-value wind prediction for offshore environments. *Wind Energy*. 2014.
17. Giebel G, Brownsword R, Kariniotakis G, Denhard M, Draxl C. The state-of-the-art in short-term prediction of wind power: a literature overview. *ANEMOS. plus*. 2011.
18. Barthelmie RJ, Murray F, Pryor SC. The economic benefit of short-term forecasting for wind energy in the UK electricity market. *Energy Policy*. 2008;36(5):1687-1696.
19. Meyers J, Meneveau C. Optimal turbine spacing in fully developed wind farm boundary layers. *Wind Energy*. 2012;15(2):305-317.
20. Wan YH, Ela E, Orwig K. Development of an equivalent wind plant power curve. Tech.Rep. NREL/CP-550-48146, NREL, Golden, CO; 2010.
21. King J, Clifton A, Hodge BM, RePPAE LLC. *Validation of Power Output for the Wind Toolkit*: National Renewable Energy Laboratory; 2014.
22. Potter CW, Gil HA, McCaa J. Wind power data for grid integration studies. In: Power Engineering Society General Meeting, 2007 Tampa, FL. IEEE. IEEE; 2007:1-6.
23. Potter CW, Lew D, McCaa J, Cheng S, Eichelberger S, Grimit E. Creating the dataset for the western wind and solar integration study (USA). *Wind Engineering*. 2008;32(4):325-338.
24. Stevens RJAM, Gayme DF, Meneveau C. Coupled wake boundary layer model of wind-farms. *Journal of Renewable and Sustainable Energy*. 2015;7(2):023115.

25. Jeon S, Kim B, Huh J. Comparison and verification of wake models in an onshore wind farm considering single wake condition of the 2 MW wind turbine. *Energy*. 2015;93:1769-1777.
26. Katić I, Hjstrup J, Jensen NO. A simple model for cluster efficiency. In: European Wind Energy Association Conference and Exhibition; Rome, Italy: 1986:407-410.
27. Jacobsen HS. Wasp wake effect model. http://www.wasp.dk/wasp#details_wakeeffectmodel; 2015.
28. Hirth BD, Schroeder JL, Gunter WS, Guynes JG. Coupling Doppler radar-derived wind maps with operational turbine data to document wind farm complex flows. *Wind Energy*. 2015;18(3):529-540.
29. Breiman L. Random forests. *Machine Learning*. 2001;45(1):5-32.
30. Liaw A, Wiener M. Classification and regression by randomForest. *R News*. 2002;2(3):18-22. http://CRAN.R-project.org/doc/Rnews/.
31. Dee D, Fasullo J, Shea D, Walsh J. The climate data guide: atmospheric reanalysis: overview & comparison tables. *National Center for Atmospheric Research*. 2015. https://climatedataguide.ucar.edu/climate-data/atmospheric-reanalysis-overview-comparison-tables.
32. Schwartz MN, George RL. On the use of reanalysis data for wind resource assessment. In: 11th Applied Climatology Conference. National Renewable Energy Laboratory; Golden, CO: 1998.

SUPPORTING INFORMATION

Additional Supporting Information may be found online in the supporting information tab for this article.

How to cite this article: Staid A, VerHulst C, Guikema SD. A comparison of methods for assessing power output in non-uniform onshore wind farms. *Wind Energy*. 2018;21:42–52. https://doi.org/10.1002/we.2143