WIND ENERGY

WIND ENERGY
Wind Energ. 2017; **20**:1501–1513
Published online 21 March 2017 in Wiley Online Library (wileyonlinelibrary.com). DOI: 10.1002/we.2105

RESEARCH ARTICLE

A production economics analysis for quantifying the efficiency of wind turbines

Hoon Hwangbo, Andrew Johnson and Yu Ding
Department of Industrial and Systems Engineering, Texas A&M University, College Station, Texas 77840, USA

ABSTRACT

We quantify the productive efficiency of a wind turbine, using power output and environmental variable data, measured either at the turbine or at a meteorological mast near the turbine. The methods described can potentially help with decision makings in asset procurement, maintenance planning, or wind turbine control optimization. The current recommendation from the International Electrotechnical Commission regarding turbine performance evaluation is to use a power curve or power coefficient. What is commonly used in practice is the average performance power curve or power coefficient. When using the power curve to quantify productive efficiency, one crucial shortcoming is the lack of a common best performance benchmark, while the power coefficient approach uses an absolute efficiency measure that is not achievable. We introduce a new approach for efficiency quantification based upon production economics' concepts which provides estimates of a best performance benchmark. Our specific approach has two main components: (a) a best performance power curve is estimated and used together with the average performance curve to show how well a turbine has performed relative to its full potential; and (b) a covariate matching procedure is developed to control for environmental influences for the comparison of turbine performances over different periods. Through a simulation study, we demonstrate that the proposed efficiency is more sensitive to potential changes in the turbine. When analyzing multi-year wind turbine data, we observe that the turbine's efficiency is improving during the first 2 years of operation and then remains relatively constant during years 3 and 4. Copyright © 2017 John Wiley & Sons, Ltd.

KEYWORDS
efficiency analysis; environmental variables; non-parametric statistics; power curve; wind turbine performance

Correspondence
Yu Ding, Department of Industrial and Systems Engineering, Texas A&M University, College Station, Texas 77840, USA.
E-mail: yuding@tamu.edu

Received 24 November 2015; Revised 30 November 2016; Accepted 17 February 2017

1. INTRODUCTION

Wind energy is one of the fastest growing renewable energy sources. During the past decade, the cumulative installed capacity of wind energy in the USA has drastically increased from 6.7 GW in 2004 to nearly 66 GW in 2014.[1] This fast growth in wind energy capacity is not unique to the USA but is also a common trend shown globally.[2] Among the important issues to be addressed for making wind energy more competitive, one concerns performance quantification of wind turbines. Addressing this issue adequately, namely, quantifying a turbine's productive efficiency and understanding its change over time, helps guide numerous decisions for operating wind turbines; for instance, planning maintenance actions to counter degradation in turbine performance,[3] justifying costly retrofitting turbine upgrades,[4] or optimizing pitch and torque control[5] to prolong a turbine's service life. Performance quantification enables performance benchmarking of turbines from different manufacturers, which could also help with decision making in the asset procurement process.

For performance evaluation of wind turbines, International Electrotechnical Commission (IEC)[6] recommends to use (1) annual energy production (AEP), (2) power curve, or (3) power coefficient. The drawback of using power output directly (as in the case of AEP) is obvious, because wind power output is affected by wind input conditions, so that a fair comparison based on power output requires the input conditions to be set to comparable levels but doing so is not easy. On the other hand, once the power curves are estimated (e.g., Figure 1(a)), the relative positions on the power curve plot may suggest relative productive efficiency; see an example in Figure 1(b).

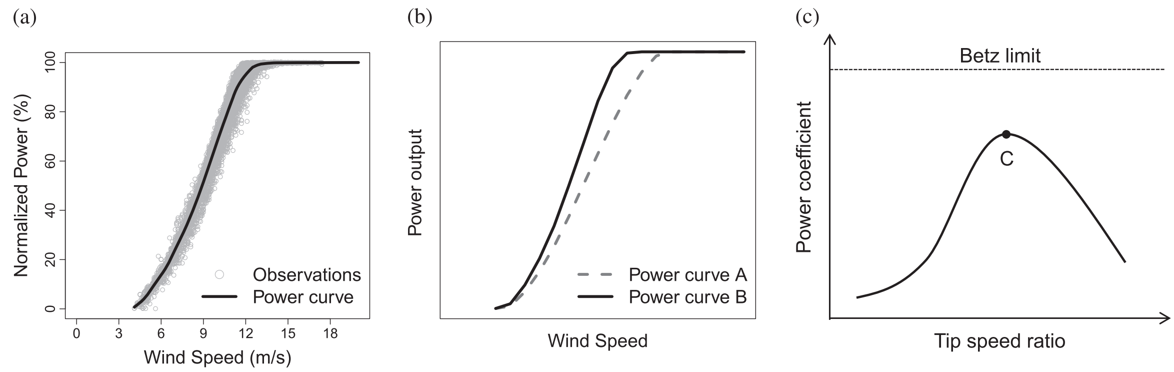

Figure 1. Illustration of turbine performance evaluation: (a) scatter plot of wind speed and power data, and the estimated power curve, (b) two power curves indicating relative efficiencies of wind turbines, in which curve B suggests a higher productive efficiency, (c) power coefficient curve and the Betz limit.

Power coefficient refers to the coefficient term, C_p, used in the power production equation $P = 0.5\rho A V^3 C_p$,[7] where P is the extractable wind power, V denotes wind speed, ρ denotes air density, and $A = \pi R^2$ is the rotor swept area for a rotor of radius R. The equation may leave readers the impression that C_p is a constant; in fact, it is not. C_p is typically modeled as a function of the tip speed ratio (i.e., the ratio between the tangential speed of the tip of a blade and the wind speed), attack angle (related to wind direction) and air density. This dependency of C_p on weather related inputs makes the power coefficient a *functional curve*, often plotted against the tip speed ratio; see an example in Figure 1(c). Same as the estimation of power curve, estimation of power coefficient curve averages the power coefficients, computed from the observational data through the power production equation. In practice, the largest power coefficient on the curve, for instance, point C in Figure 1(c), as the representative of the whole curve, is used for quantification of the aerodynamic efficiency.[8–10] The peak power coefficient is a popular efficiency measure, widely used to evaluate wind turbine designs[9,10] and various control schemes including pitch and torque controls.[11,12]

Considering that the role of the average power curve and power coefficient play in turbine performance evaluation, one may wonder whether either of them would be a good metric for the quantification of productive efficiency as well. To quantify a turbine's productive efficiency, one would need to estimate the best achievable performance as a benchmark, so that the ratio of the current performance to the best performance quantifies the degree to which the turbine has performed relative to its full potential. The difficulty in using a power curve or power coefficient is that they are both an average performance measure. As illustrated in Figure 1(a), the resulting power curve is a line passing through the middle of the observational data. A power curve informs us about the average behavior of the wind turbine; however, we would like to estimate a best performance benchmark.

In the case of power coefficient, practitioners use the theoretical upper limit, known as the Betz limit (=0.593),[13] as the best performance benchmark. But this upper limit is not practically achievable; generally power coefficients estimated are below 0.45. So when normalizing by the Betz limit, the corresponding efficiency measure never approaches one, rendering interpretation difficult. A more crucial limitation for power coefficient is that the efficiency quantification is based on a point representation of the power coefficient curve, so that two turbines with the same peak power coefficient values could still have different power coefficient curves, implying that their power productive efficiencies differ. We will present an example in Section 4 to illustrate this point.

These limitations of the existing performance benchmarks motivate us to look into the field of production economics[14] where efficiency analysis is a primary focus. In production economics, efficiency quantification is based on the estimation of a production function and the explicit modeling of systematic inefficiency, using input and output data for a set of production units, be it firms, factories, hospitals or power plants. In the context of wind energy, a wind turbine is a power production unit, wind speed is the dominating input driving power production, and the generated power is the output. In this paper, we introduce a production economics approach to estimate a performance benchmark for a wind turbine, which will be used together with the average performance power curve to quantify productive efficiency of the turbine.

Although wind speed is the dominating force driving wind power production, other environmental variables including wind direction, humidity, turbulence intensity and air density (as shown in the power production equation) may all affect wind power output.[15] Many of the environmental variables are measured at the meteorological mast. In order for the resulting productive efficiency measure to be used for performance comparison between different turbines or for the same turbine over different time periods, it is important to control for the environmental influences, so that the performance comparison quantifies the differences coming from a turbine's endogenous characteristics. For this reason, we develop a covariate matching procedure, allowing us to select a subset of the data, for which the probability distributions of the

Wind Energ. 2017; **20**:1501–1513 © 2017 John Wiley & Sons, Ltd.
DOI: 10.1002/we

environmental variables are matched. Note that controlling for environmental influences is also necessary when either AEP, power curve or power coefficient is used as a performance measure. Because of this, the covariate matching procedure is potentially widely applicable.

The rest of this paper is organized as follows. Section 2 starts with presenting the general concept and idea of production economics analysis, followed by discussing how to apply the ideas from production economics to the efficiency quantification of wind turbines. Section 3 establishes a method to control for environmental influences by matching the probability densities of the environmental covariates. Section 4 presents a case study applying our proposed efficiency measure to actual wind turbine data. Section 5 concludes the paper.

2. ESTIMATION OF THE BEST PERFORMANCE BENCHMARK

In this section, we apply a data-driven method, borrowing ideas from production economics, to estimate benchmark performance. We are interested in identifying a benchmark that is data driven and thus describes the efficient behavior of a wind turbine relative to the observed operational data. This section starts by introducing production economics, followed by a discussion describing standard efficiency estimation methods tailored to handle wind power production data.

2.1. Background of production economics analysis

Consider a set of production units (e.g., a wind farm) using x input (e.g., investment in a wind energy project) and producing y output (e.g. revenue from power generation). We can create a scatter plot of many x-y data pairs coming from different production units or the same production unit but over different periods; see Figure 2. Assuming no measurement errors associated with x and y, a common estimator in production economics, Data Envelopment Analysis (DEA),[16] estimates the efficient frontier enveloping all the observations.

The concept of an efficient frontier is understood as follows: a production unit whose input-output is on the frontier is more efficient than the production units whose input-output is being enveloped by the frontier. Consider observation D. Using the same input, the production unit associated with D produces less output than the production unit associated with point E; while to produce the same output, the production unit associated with D needs more input than the production unit associated with point F. So the production unit associated with D must be inefficient.

The efficient frontier is also called the production function, denoted by $f(x)$. The production function characterizes producible output given input x in the absence of inefficiency. Using the production function, the output of the inefficient production unit D can be expressed as

$$y_D = f(x_D) - u_D, \tag{1}$$

where $u_D \geq 0$ denotes the systematic inefficiency.

To estimate the production function $f(x)$, certain assumptions are made restricting the shape of the frontier. The most common assumption is that the frontier forms a monotone increasing concave function consistent with basic stylized characteristics of production.[17] When the data are assumed noise free, the tightest boundary enveloping all observations and maintaining monotonicity and concavity is a piece-wise linear function.

Convex or concave piecewise linear methods assuming noise-free data encounter some problems when applied to wind turbine data. The first is that the wind turbine data, like all other physical measurements, are inevitably contaminated by noises. The second difference is that the shape of the wind-power scatter plot is not concave. Instead, the data appears to follow an S-shape, as shown in Figure 3, comprising a convex region, followed by a concave region, and the two segments of curves are connected at an inflection point.

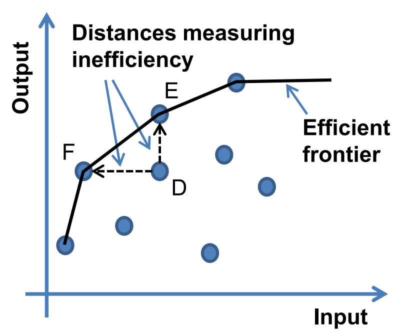

Figure 2. Production data and efficient frontier. [Colour figure can be viewed at wileyonlinelibrary.com]

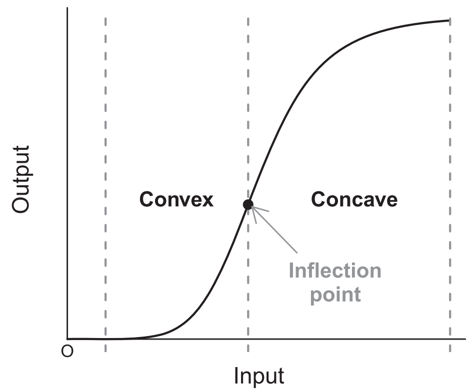

Figure 3. A 2-dimensional S-shaped curve: a region where the production function is convex followed by a region where the function is concave, connected with an inflection point.

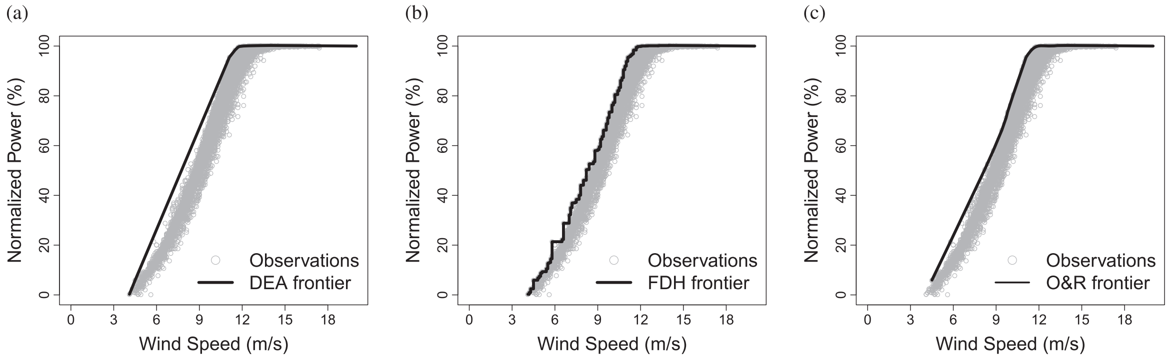

Figure 4. Various types of production function: (a) data envelopment analysis, (b) free disposal hull, and (c) deterministic S-shaped production function. FDH, free disposal hall; O&R, Olesen-Ruggiero.

Methods from production economics have recently been used in wind energy applications. Two studies[18,19] are noted but neither of them addressed simultaneously both aforementioned problems of convex/concave piecewise linear methods in dealing with wind turbine data. Carvalho *et al.* (2009)[18] simply applied the DEA approach to the wind-power data, which envelopes all observations from above with a piece-wise concave function. Pieralli *et al.* (2015)[19] applied a different approach, known as free disposal hall (FDH). FDH relaxes the concave function assumption but still assumes noise-free observations. Production economics researchers call this type of frontier analysis approach, assuming noise-free observations, *deterministic*. The problem with applying a deterministic approach to noisy wind production data is that it tends to overestimate the best performance benchmark because every observation is assumed to be achievable. Figure 4(a) and 4(b) show, respectively, the production frontiers estimated using the DEA and FDH approaches.

In production economics, the need to model noise is well established, promoting the subfield of stochastic frontier analysis (SFA).[20] The SFA model includes a random noise term ϵ to equation (1), as follows:

$$y = f(V) - u(V) + \epsilon. \tag{2}$$

Here and in what follows, we define the production function in terms of the wind turbine data and replace the input variable x with wind speed variable V, and consequently, y refers specifically to the power output produced by a turbine. To be consistent with the IEC standards,[6] we denote by V the wind speed normalized by air density, that is, $V = V_{10min} \cdot (\rho_{10min}/\rho_0)^{1/3}$ where V_{10min} and ρ_{10min} are, respectively, wind speed and air density averaged over 10-minute time intervals, and ρ_0 is the average of the measured air density at the test site during the periods of data collection. Hereinafter, we refer to this normalized wind speed as 'wind speed' unless otherwise stated. Random noise ϵ is assumed having a zero mean, while the systematic inefficiency term $u(V)$ is a non-negative random variable with positive mean, i.e., $\mu(V) := E[u(V)] > 0$. Note that $u(V)$ is a function of V, meaning that the amount of inefficiency varies as the input changes, which is referred to as a heteroskedastic inefficiency term.

While the current SFA research considers the noise effect in observational data, researchers typically do not address the second problem mentioned above, namely, the S-shaped curve exhibited in the wind turbine data; rather, researchers rely typically on parametric functional forms, such as Cobb-Douglas, that need not satisfy the S-shape for any parameter values. In fact, the S-shape constraint corresponds to the Regular Ultra Passum (RUP) law[14] in economics, which is motivated by production units having an increasing marginal rate of productivity followed by a decreasing rate of marginal productivity. Apparently, wind turbine power curves also satisfy the RUP law.

Very few production function estimators impose the RUP law explicitly. The exception is the DEA-type estimator developed by Olesen and Ruggiero (2014).[21] The frontier analysis employed by Olesen and Ruggiero (2014) is still deterministic, enveloping all observations from earlier, and consequently suffers from the overestimation that all other deterministic production function estimators suffered. Figure 4(c) presents an example of the production frontier estimated by the Olesen-Ruggiero method[21] when it is applied to a set of wind turbine data.

2.2. Estimation of average performance curve and best performance benchmark

We use shape constrained functional estimation to model the typical shape of wind data, which is consistent with the RUP law, and model noise. The estimator is described in detail in a previous study;[22] however, in the succeeding text, we summarize the key features.

The basic production function model in equation (2) is re-written as follows:

$$y = [f(V) - \mu(V)] + [\mu(V) - u(V) + \epsilon] = g(V) + e, \tag{3}$$

if we let $g(V) := f(V) - \mu(V)$ and $e := \mu(V) - u(V) + \epsilon$. The expression connects the power curve with the production function. For us to see this, consider the following. The term e is a redefinition of the error term with expectation zero. Because of this, $g(V)$ is the average practice power curve. As such, the production function $f(V)$ differs from the power curve $g(V)$ by the mean of the inefficiency varying by V.

This connection helps lay out the intuition behind the procedure of estimating $f(V)$. One would start with a power curve from the wind turbine data. Then estimate the mean function of the inefficiency and use it to rotate the average power curve to the new position to be the production frontier function.

We stress that because the final $f(V)$ needs to satisfy the RUP law (i.e., the S-shape constraint), the average performance power curve $g(V)$ that comes before the production function must satisfy the same shape constraint. This requirement makes our power curve estimation procedure different from those currently used in practice because none of them imposes the S-shape constraint explicitly. Common practice including the standard procedure recommended by IEC[4,6] estimates a power curve nonparametrically because the functional form of the power production equation is unknown. The resulting estimates still tend to approximate an S-shape curve, but with noticeable local differences from a strictly S-shaped curve. We will present an example in the succeeding text in Figure 5.

Estimating the shape constrained power curve $g(V)$ requires imposing convexity and concavity in the low and high wind speed regions, respectively. The convex segment should connect to the concave segment at the inflection point, which itself needs to be estimated from the data. The estimation of the convex segment or the concave segment can be carried out by using the method Convex Nonparametric Least Squares.[23] When the two segments are estimated simultaneously

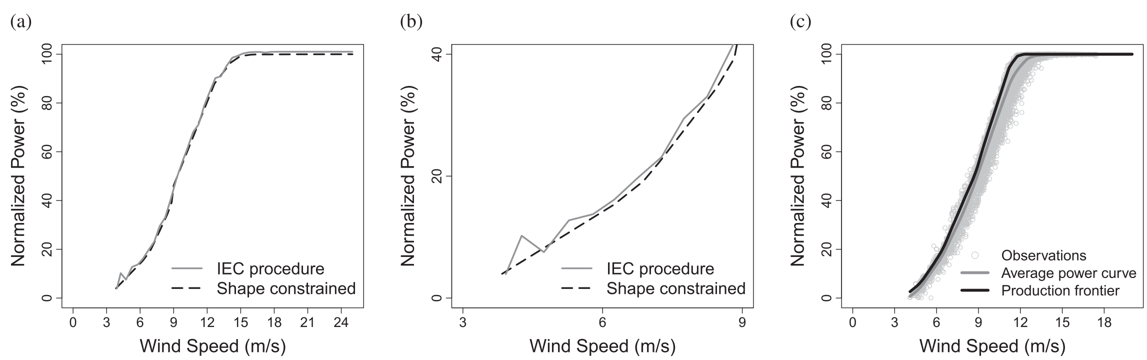

Figure 5. Illustration of stochastic S-shaped production function: (a) comparison with the International Electrotechnical Commission (IEC) standard procedure, (b) comparison with the IEC standard for the wind speed ranging from 3 m/s to 9 m/s, (c) estimates of average practice power curve and production frontier.

maintaining the shape constraints and the continuity at the inflection point, the final outcome of this step is the average performance power curve $g(V)$.

After $g(V)$ is estimated, we can take differences between the fitted power curve and the output y. According to the relationship in (3), the resulting residuals are the summation of two random components: $\mu - u$ and ϵ. Our modeling assumption states that u is non-negative, and ϵ is symmetrically distributed with respect to a zero mean. So, we expect to see a significant decrease in the density of the residuals at the value of μ. This understanding is used to estimate μ, which is unknown. If we can locate where the greatest decrease in the residual distribution occurs, it gives us an estimate of μ. Specifically, the technique in Hall and Simar[24] can be used for this estimation.

The following summarizes the steps in estimating the shape constrained stochastic production function $f(V)$:

1. Use the wind turbine data (wind speed and power) to estimate $g(V)$ while imposing the shape constraints and the continuity requirement at the inflection point; denote the estimated curve by $\hat{g}(V)$.
2. Estimate $\mu(V)$, the mean function of the inefficiency term;
3. Estimate $f(V)$ based on the relationship of $f(V) = g(V) + \mu(V)$; denote the estimated curve by $\hat{f}(V)$.

The steps of the procedure have been simply described here, and for technical details of the estimation procedure, see Hwangbo *et al.* (2015).[22]

Figure 5(a) presents two average performance power curves: one satisfies the shape constraint, obtained by the procedure outlined earlier, whereas the other obtained by the IEC's standard procedure does not. The two estimates are similar to each other but not the same (see the enlarged version in Figure 5(b)). Figure 5(c) presents the production frontier curve and the average performance curve, both satisfying the shape constraint. Compared with the deterministic estimators shown in Figure 4, one notices that this production frontier does not envelop all the observations. The observations beyond the frontier are affected by significant positive random noise.

With the average performance power curve and the best performance benchmark estimated, we propose the following efficiency measure, θ, which is the ratio of the energy produced under the average performance power curve over that under the best performance, integrated over the whole wind spectrum:

$$\theta = \frac{\int_{V_{ci}}^{V_{co}} \hat{g}(V)dV}{\int_{V_{ci}}^{V_{co}} \hat{f}(V)dV}, \tag{4}$$

where V_{ci} is the cut-in wind speed and V_{co} is the cut-out wind speed. Apparently, θ takes a value between 0 and 1; the closer θ is to 1, the closer the wind turbine performs to its full potential.

3. CONTROLLING FOR ENVIRONMENTAL INFLUENCES THROUGH COVARIATE MATCHING

We estimate both the best practice frontier curve and average performance power curve as a function of wind speed. Besides wind speed, air density and several other environmental variables, including wind direction, humidity, turbulence intensity and wind shear, all potentially affect the wind power production. These environmental influences are not controllable, but their existence does play a role affecting the inefficiency estimated from the power output data. Consequently, when comparing the productive efficiency of different turbines or the same turbine over different operational periods, practitioners often wonder what part of inefficiency is due to the turbine's intrinsic differences and what part of inefficiency comes from differences in environmental characteristics such as air dampness. This sort of ambiguity can be alleviated if the comparison periods have comparable environmental profiles. Creating comparable environmental profiles is what we try to accomplish in this section.

Let us consider monitoring a turbine's efficiency change over a number of time periods. The environmental variables are referred to as covariates in statistics. The covariate vector includes measurements from the wind farm as well as those computable using available measurements (such as wind shear), but does not include variables shown in previous studies to have little or no correlation to power output. We acknowledge that wind farms may gather different data on environmental measurements; for instance, one of the wind farms we worked with does not have humidity measurements. Nonetheless, our procedure presented here can be applied regardless of the number of variables included in the covariate vector.

We describe a method to match covariate vectors to make the environmental profiles across different time periods as similar as possible, thus removing the effect of environmental influences from the efficiency analysis. Suppose that we have p environmental variables in the covariate vector and t periods of operation. We assume that for different periods, the common p variables are available. If not, we reduce the set of variables included in the covariate vector to the subset common to all periods. We arbitrarily choose one of the periods as the reference period.

Let $\boldsymbol{X}^t = (x_1^t, \ldots, x_p^t)$ for $t = 1, \ldots, T$ denote the vector of covariates consisting of p environmental variables, including wind speed and air density, observed during the tth period and y^t be the corresponding power output. The reference period is denoted by $t' \in \{1, \ldots, T\}$. We recommend setting $t' = T$ so that the analysis is based on the most recent data.

Wind Energ. 2017; **20**:1501–1513 © 2017 John Wiley & Sons, Ltd.
DOI: 10.1002/we

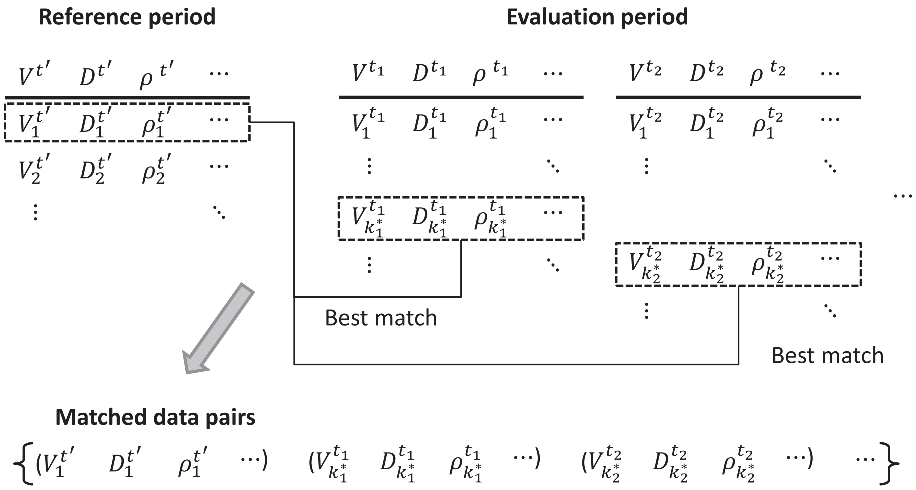

Figure 6. Procedure to construct a set of matched covariate vectors.

A non-reference period, also called an evaluation period, is denoted by $t \neq t'$. The data pairs in periods t' and t are represented, respectively, by $(\boldsymbol{X}_j^{t'}, y_j^{t'})$ for $j = 1, \ldots, n_{t'}$ and $(\boldsymbol{X}_k^{t}, y_k^{t})$ for $k = 1, \ldots, n_t$, where $n_{t'}$ and n_t are the number of observations in the two periods, respectively.

Our matching procedure starts with selecting a single observation in the reference period, comparing its environmental covariates (i.e., variables in $\boldsymbol{X}$) with the covariates of an observation in an evaluation period, and assessing their dissimilarity using a score defined in the succeeding text. Repeat this for all observations in the evaluation period and select the observation yielding the smallest dissimilarity score as the best match. If no observation in the evaluation period has a small enough score, the observation is removed from the reference period set. Otherwise, choose another evaluation period and find the best match to the observation in the reference period. As such, for a single observation in the reference period, there will be one matched observation from each of the evaluation periods. Altogether this produces a set of matched covariate vectors having similar environmental profiles. We then proceed with the same action for all observations in the reference period. Figure 6 illustrates this procedure.

Now let us define the dissimilarity score used in the matching process. Consider the j-th observation in the reference period t'. For the q-th variable in the covariate vector, we denote by $\mathcal{S}_{kq}$, $q \in \{1, \ldots, p\}$ and $k \in \{1, \ldots, n_t\}$, the dissimilarity score between this reference observation and the k-th observation in the evaluation period t. The dissimilarity score, $\mathcal{S}_{kq}$, is defined as follows:

$$\mathcal{S}_{kq} = \frac{|x_{jq}^{t'} - x_{kq}^{t}|/sd_q^{t'}}{x_{jq}^{t'}/mean_q^{t'}}. \tag{5}$$

The smaller the dissimilarity score is, more similar the two covariate vectors are. In the earlier definition, $mean_q^{t'}$ and $sd_q^{t'}$ are, respectively, the sample mean and the sample standard deviation of the q-th variable in period t', and their use is to normalize the scale of the preceding terms. A normalization is needed because two covariates can have different ranges both in an absolute value sense and in a percentage sense; for instance, air density varies by 15% from its mean value, while wind speed can vary up to 100%. Without the normalization, covariates with low variability will be labeled as matching, even though their density functions differ significantly.

The aforementioned formula can be used for almost any environmental variables, except for wind direction, which is a circular variable for which the value 0 and 360 are equivalent. For wind direction, we slightly modify the dissimilarity score as follows:

$$\mathcal{S}_{kq} = \frac{\min\left\{|x_{jq}^{t'} - x_{kq}^{t}|,\, 360 - |x_{jq}^{t'} - x_{kq}^{t}|\right\}}{x_{jq}^{t'}} \cdot \frac{mean_q^{t'}}{sd_q^{t'}}. \tag{6}$$

Once $\mathcal{S}_{kq}$ is calculated, we find the set of candidate best matches to the j-th observation in t' that satisfies $\mathcal{S}_{kq} \leq \omega$, where ω is a pre-specified threshold and usually is set to a small quantity, e.g., 0.25. The use of ω is to set a standard for eligible matches, such that any resulting matches are considered 'good enough'. If there is no observation in the evaluation period satisfying this dissimilarity constraint, then this j-th observation in period t' is skipped as having no matched record. On the other hand, if there are multiple of observations in this candidate set, we choose the best match, indexed as k^* that satisfies the following minimax criterion:

$$k^* = \underset{k \in \mathcal{K}}{\operatorname{argmin}} \left\{ \max_{q \in \{1,\ldots,p\}} \mathcal{S}_{kq} \right\}, \tag{7}$$

where $\mathcal{K} = \{k : \mathcal{S}_{kq} \leq \omega, \ \forall q = 1, \ldots, p\}$. Without the ω threshold and using the minimax criterion alone, one could end up with a match whose dissimilarity score may be uncomfortably large.

Once the maching process is done for all evaluation periods, for notational simplicity, we re-index the matched data pairs by i, such that $(\boldsymbol{X}_i^t, y_i^t)$, $i = 1, \ldots, n$ and for $t = 1, \ldots, T$ where n is the number of the matched data pairs.

Note that the earlier matching process does not produce an exact match but a good match, subject to the dissimilarity allowed by the threshold ω. To confirm the quality of the matches, we suggest plotting the probability density functions (pdf) of each environmental variable, empirically estimated from the data and visually inspected to assess how well the pdfs match across the comparison periods. Numerical examples will be presented in the case study section to further illustrate this point.

To evaluate the productive efficiency of a wind turbine controlling for the environmental influences, we use the matched data pairs, namely $\{(\boldsymbol{X}_i^t, y_i^t)\}$, to estimate the average practice power curve $g(V)$ and the best performance benchmark $f(V)$. Without prior knowledge of the period when a wind turbine shows the best performance, the best practice curve cannot be restricted to a specific t. For this reason, we pool the matched data pairs from all periods (including t') while estimating $f(V)$. In order to see how the turbine productive efficiency may have changed from period to period, we further estimate the average practice power curve for each t. As such, the efficiency measure in equation (4) can be re-expressed as follows:

$$\theta_t = \frac{\int_{V_{ci}}^{V_{co}} \hat{g}_t(V)dV}{\int_{V_{ci}}^{V_{co}} \hat{f}(V)dV}, \tag{8}$$

where $\hat{g}_t$ is the average practice curve of period t.

4. CASE STUDY

In this case study, we use data from two onshore wind turbines (WT1 and WT2) and two offshore wind turbines (WT3 and WT4). Table I summarizes the characteristics of these wind turbines; for certain entries, an approximation rather than the accurate value is given for the protection of the identities of the turbine manufacturers and wind farms. The wind turbine data include observations during the first 4 years of their operations. In particular, the onshore data are available for the period of 2008–2011, and the offshore data are from 2007 to 2010. The 4 years worth of data allows us to look into the performance change during the early stage of a turbine's operation. All measurements are the averages during 10-minute time intervals, a common practice for data arrangement in the wind industry.

We analyze the wind turbine data on an annual basis. In other words, we divide the 4-year data into four consecutive annual periods, namely that we have $T = 4$ and $t = 1, 2, 3, 4$. We evaluate turbine efficiency for each year because seasonal variations in atmospheric and meteorological conditions are significant, but yearly patterns are relatively stable.

The measurements we have for the offshore wind farm include power output (y), wind speed (V), wind direction (D), air density (ρ), turbulence intensity (I), wind shear (S) and humidity (H). For the onshore wind farm, however, the humidity measurements are not available. The power output is always measured on the wind turbine. Most of the environmental measurements are taken from a meteorological mast closest to the turbine, with the exception of wind speed and turbulence intensity which are measured on the wind turbine. The mast measurements are used either because some variables are only measured at the mast (such as air pressure and ambient temperature, which are used to calculate air density) or because the mast measurements are considered more reliable (such as wind direction). The wind speed measurements are taken from the nacelle anemometers, and they are further used to calculate the turbulence intensity. The industrial partners who provided the data told us that the wind speed data, measured by the nacelle anemometer, have been adjusted to be the free

Table I. Specification of the wind turbines.

	Onshore	Offshore
Location	USA	Europe
Cut-in wind speed (m/s)	3.5	3.5
Cut-out wind speed (m/s)	20	25
Rated wind speed (m/s)	approximately 13	approximately 15
Rated power (MW)	1.5–2	approximately 3
Initiation of operations (year)	2008	2007

Wind Energ. 2017; **20**:1501–1513 © 2017 John Wiley & Sons, Ltd.
DOI: 10.1002/we

stream equivalents in front of a turbine, rather than the raw measurements taken in the wake of a turbine's rotor. We use the nacelle anemometer measurements for wind speed to better differentiate the wind turbines nearby the same mast.

Prior to analyzing the data, we conducted some preprocessing, removing data records with missing values or data records taken while a turbine is unavailable, or excluding measurements such as negative power values. These observations seem to occur randomly, so we do not anticipate a sample selection bias.

After the preprocessing, we select the subset of data with comparable environmental profiles through the covariate matching method described in Section 3. For onshore wind turbines, the covariates to be matched include $\boldsymbol{X} = (V, D, \rho, I, S)$, whereas for offshore wind turbines $\boldsymbol{X} = (V, D, \rho, H, I)$. The reason that we did not include wind shear S for offshore wind turbines is because a previous study[15] found that conditioned on the inclusion of (V, D, ρ, H, I), the effect of S on power output appears negligible. We did test and see what if we included wind shear in the offshore turbine data matching process. It turns out that the results of that analysis produced the same insights described in the succeeding text.

For all turbine cases, we use $\omega = 0.25$ as the threshold assessing the dissimilarity. Before the covariate matching, the number of observations in each annual dataset ranges from 14,000 to 37,000, and these numbers reduce to 1400–2300 after the matching. The significant reduction in the number of observations indicates the importance of matching covariates. Had we used all the raw observations, the efficiency results would describe the differences in the operating environment across periods, rather than the intrinsic efficiency of the turbine. The matched data set still includes thousands of observation which is a large enough sample for estimating the best performance benchmark as well as the average performance curve.

Figure 7 and 8 present the pdfs of each environmental variable across the four comparison periods after the covariate matching; Figure 7 is for onshore turbine WT1, while Figure 8 is for offshore turbine WT3. We omit the plots for WT2 and WT4, which are similar, in the interest of space. One can notice that the choice of $\omega = 0.25$ leads to sufficiently good matching as demonstrated in the pdf plots.

Subsequently, we use the matched subset of data to estimate the productive efficiency measure for each comparison period, as defined in equation (8). Because of the randomness in the data, we add a confidence interval to the efficiency measure. To do that, we use a bootstrap procedure, which is to resample the data with replacement B times, and for each resampled dataset, compute the efficiency measure, which altogether results a total of B replications. Then, the confidence interval for the efficiency measure can be constructed using these B sample replications; for details about the bootstrap procedcure, please refer to a previous study.[25] In this case study, we used $B = 100$ and calculate 90% confidence intervals.

Figure 9 shows the productive efficiency θ_t and its confidence intervals for the four comparison periods, which are the first 4 years of a turbine's operation. Interestingly, one can notice that for all four turbines, their productive efficiency appears to have increased slightly, rather than deteriorated, during the early stage of operation. This pattern is more obvious for offshore turbines. This initial increase in efficiency was also recognized by Staffell and Green (2014).[26] Figure 9b in Staffell and Green (2014)[26] plots the fleet-level performance degradation of wind turbines over a 20-year period using the

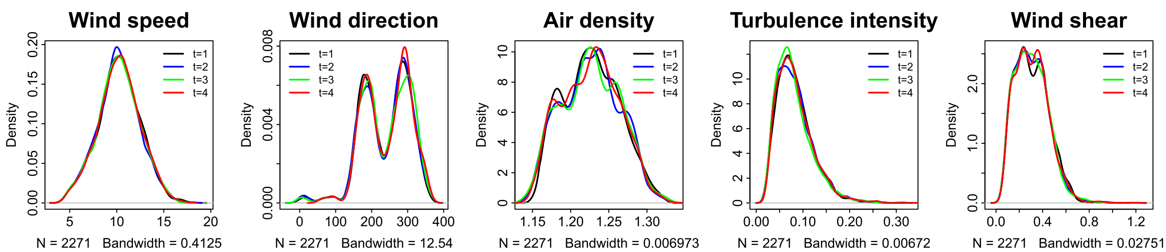

Figure 7. Probability density function plots of the matched covariates over the four comparison periods for onshore turbine WT1. [Colour figure can be viewed at wileyonlinelibrary.com]

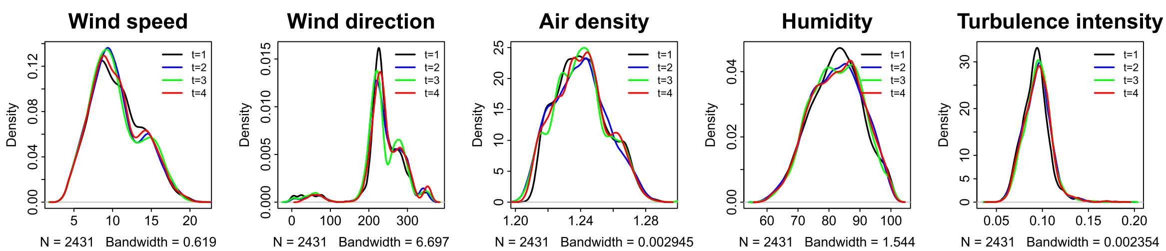

Figure 8. Probability density function plots of the matched covariates over the four comparison periods for offshore turbine WT3. [Colour figure can be viewed at wileyonlinelibrary.com]

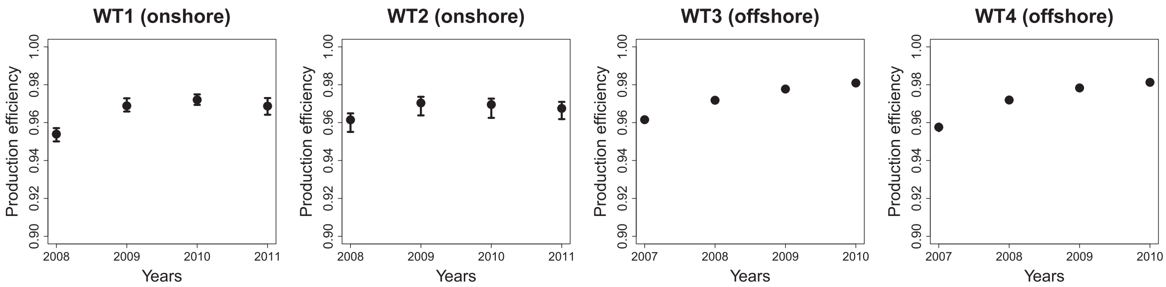

Figure 9. Productive efficiency θ_t, $t = 1, 2, 3, 4$. The bars represent 90% confidence intervals and the dots denote the mean values of the efficiency. For offshore wind turbines, the confidence intervals are very narrow, so that the bars are not shown explicitly.

Table II. Comparison between the productive efficiency θ_t and the (peak) power coefficient: the values represent the mean of the bootstrap estimates and the values in parenthesis are the respective 90% confidence intervals.

	Power coefficient				Productive efficiency θ_t			
	Year 1	Year 2	Year 3	Year 4	Year 1	Year 2	Year 3	Year 4
WT1	0.371	0.388	0.393	0.393	0.954	0.969	0.972	0.969
	(0.367, 0.377)	(0.386, 0.392)	(0.390, 0.397)	(0.389, 0.398)	(0.950, 0.957)	(0.966, 0.973)	(0.969, 0.975)	(0.964, 0.973)
WT2	0.444	0.466	0.463	0.462	0.962	0.970	0.970	0.968
	(0.439, 0.450)	(0.461, 0.470)	(0.460, 0.468)	(0.457, 0.467)	(0.955, 0.965)	(0.964, 0.974)	(0.963, 0.973)	(0.962, 0.971)
WT3	0.420	0.465	0.483	0.505	0.962	0.972	0.978	0.981
	(0.417, 0.423)	(0.461, 0.473)	(0.479, 0.488)	(0.497, 0.511)	(0.960, 0.963)	(0.971, 0.973)	(0.977, 0.979)	(0.980, 0.982)
WT4	0.417	0.473	0.484	0.506	0.958	0.972	0.978	0.981
	(0.417, 0.421)	(0.465, 0.484)	(0.477, 0.494)	(0.496, 0.516)	(0.956, 0.959)	(0.971, 0.973)	(0.977, 0.980)	(0.980, 0.983)

fleet's load factor as the performance measure. Staffell and Green (2014)'s study appears to suggest an initial period of 4–5 years before any noticeable degradation was witnessed, as well as an increase in turbine performance for the first one and half years, which is quite consistent with what we observed.

Next, we want to compare the proposed productive efficiency measure and power coefficient, given the popularity of power coefficient used in turbine performance evaluation. We calculate the peak power coefficient values for each comparison period, using the same matched subset of data and the power coefficient curves averaged for each annual period. We also apply the bootstrap procedure to compute the 90% confidence intervals of the (peak) power coefficient. The power coefficient values and the proposed productive efficiency values are presented in Table II, in which the values in the parenthesis are the respective confidence intervals.

As we mentioned before, the power coefficient itself is not a relative measure. One could divide a power coefficient by the Betz limit to get a similar interpretation as the productive efficiency value. Given that the annual power coefficient ranges from 0.371 to 0.506, the relative power coefficient efficiency would be between 63 and 85%. Please bear in mind that the Betz limit is a theoretical limit impractical to attain. So these low percentages should be taken into account with perspective; they should not be interpreted as saying that power production of the wind turbines is inefficient. If we look at the productive efficiency values, the wind turbine operations are actually reasonably efficient, relative to their full potentials.

Using the power coefficient values, we also notice a general upward trend and a leveling off. This message appears to reinforce what we found using the productive efficiency measure. In fact, there appears a fairly obvious positive correlation between the two measures; using all the values in Table II yields a correlation of 0.70 between power coefficient and the proposed productive efficiency. This positive correlation suggests that the proposed productive efficiency measures a turbine's performance on a broad common ground with the power coefficient.

One may wonder what is then the benefit of using the proposed productive efficiency measure rather than the power coefficient. To address this, we present a study below based on Khalfallah and Koliub (2007),[27] in which they investigate how dust accumulation on turbine blades affects turbine performance. They analyze wind turbines operated in Egypt where the air at the turbine site is very dusty. In Figure 10, we replot one of their graphs that compare the power production performance of a wind turbine when the blades are clean versus when they are exposed to dust accumulation for a month. Please note that with the dust accumulation, the power performance deteriorates more significantly for wind speed higher than 9 m/s than the lower wind speeds. Khalfallah and Koliub (2007) also compare the average power loss for a stall-regulated turbine and for a pitch-regulated and concluded that a pitch-regulated turbine suffers a smaller loss, around 3%.

Without the actual production data from the Egyptian turbines, we create a set of simulated data, mimicking the dust accumulation effect. We take the WT1 data measured during 2008 and modify it by decreasing the power output value by

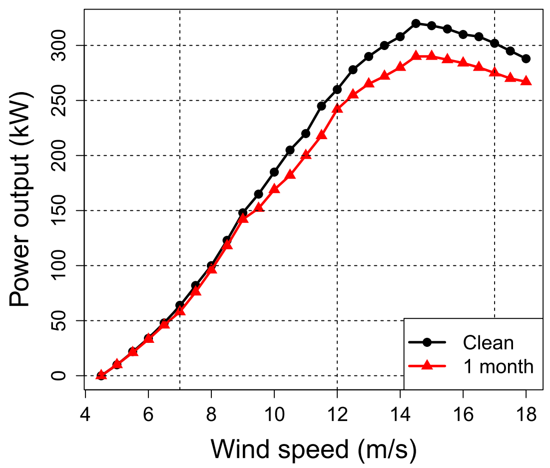

Figure 10. Effect of dust accumulation on turbine blades in Khalfallah and Koliub (2007).[27] [Colour figure can be viewed at wileyonlinelibrary.com]

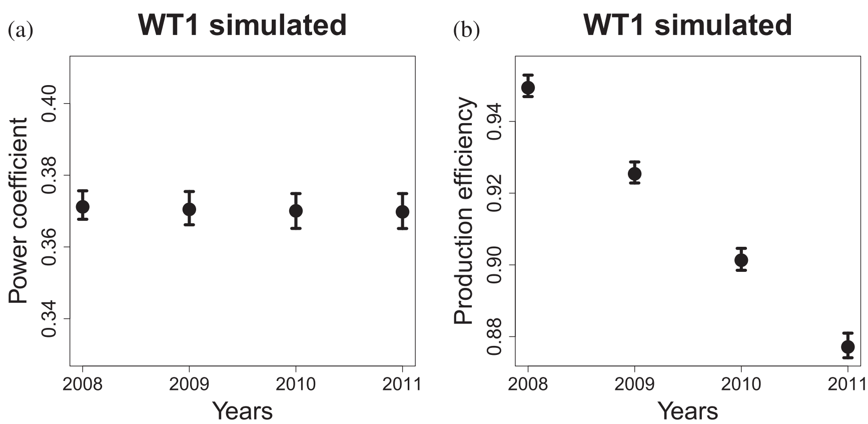

Figure 11. Change of wind turbine efficiency implied by (a) power coefficient and (b) productive efficiency. The bars represent 90% confidence intervals and the dots denote the mean values of the corresponding efficiency measures.

3% (because our turbine is pitch regulated) for those power outputs corresponding to wind speed of 9 m/s or higher. We treat the resulting data as for 2009. Then, we reduce the 2008 power data by 6% and 9% and use them as the substitute of 2010 and 2011 power data, respectively. All environmental data are left intact.

Figure 11 illustrates the change of the power coefficient (left panel) and the productive efficiency (right panel) over the 4-year period. Again, we use 100 bootstrap replications to compute the 90% confidence intervals, which show up as the bars in the plots. The use of the productive efficiency shows a clear trend, signifying the dust accumulation effect over the year. The average year-to-year decrease in the productive efficiency is about 2.4%. This magnitude of decrease seems reasonable as the 3% power reduction initially imposed is only applied to a subset of data (wind speed higher than 9 m/s). By contrast, the power coefficient does not show any trend in the change of power production ability. Recall that the power coefficient is the peak point representation of the power coefficient curve (the point C in Figure 1(c)) and, as such, using it could miss an underlying change that does not happen to the peak point area. For the original 2008 data, the peak value on the power coefficient curve is obtained at the wind speed ranging from 7.5 to 8.5 m/s, whereas the dust accumulation affects power production beyond that wind speed region. We acknowledge that the power coefficient metric could possibly reveal better than shown in this analysis if a larger range of wind speeds, instead of only the peak value, is used. The challenge is, of course, to find a proper method that aggregates the power efficient curve covering a large range of wind speeds for the purpose of characterizing a turbine's efficiency.

5. CONCLUDING REMARKS

Wind turbine operators often wonder how efficiently their turbine generators have been producing power relative to a practically attainable optimal case. For this purpose and taking advantage of ideas and methods in production economics, we introduce a method to estimate the best achievable performance in the operation of wind turbines. Determining such benchmark provides a reference for defining a normalized measure, quantifying the productive efficiency of a wind turbine. Compared with the current industrial practices, the proposed productive efficiency measure involves both the best performance benchmark and the average performance curve, whereas the power curves or power coefficients are average performance measure. Our case study shows that the proposed productive efficiency is more sensitive to a change in a turbine's production capability. When applied to the first 4 years of data on two pairs of turbines, we observe an increasing pattern in terms of productive efficiency in the initial operation of a turbine. This observation corroborates the findings in an independent study.

ACKNOWLEDGEMENT

Hwangbo and Ding's research was partially supported by the National Science Foundation under grant no. CMMI-1300560.

REFERENCES

1. US Department of Energy's Energy Efficience & Renewable Energy Website. *Windexchange: Installed wind capacity*, 2015. Available at http://apps2.eere.energy.gov/wind/windexchange/wind_installed_capacity.asp.
2. GWEC, *Global Wind Energy Outlook 2014*, Global Wind Energy Council, 2014.
3. Márquez FPG, Tobias AM, Pérez JMP, Papaelias M. Condition monitoring of wind turbines: techniques and methods. *Renewable Energy* 2012; **46**: 169–178.
4. Lee G, Ding Y, Xie L, Genton MG. A kernel plus method for quantifying wind turbine performance upgrades. *Wind Energy* 2015; **18**(7): 1207–1219.
5. Abdullah MA, Yatim AHM, Tan CW, Saidur R. A review of maximum power point tracking algorithms for wind energy systems. *Renewable and Sustainable Energy Reviews* 2012; **16**(5): 3220–3227.
6. International Electrotechnical Commission (IEC). *Iec 61400-12-1 Ed 1, Wind Turbines-Part 12-1: Power Performance Measurements of Electricity Producing Wind Turbines*. IEC: Geneva, Switzerland, 2005.
7. Boukhezzar B, Siguerdidjane H, Maureen Hand M. Nonlinear control of variable-speed wind turbines for generator torque limiting and power optimization. *ASME Journal of Solar Energy Engineering* 2006; **128**: 516–530.
8. Homola MC, Virk MS, Nicklasson PJ, Sundsbø PA. Performance losses due to ice accretion for a 5 mw wind turbine. *Wind Energy* 2012; **15**(3): 379–389.
9. Eriksson S, Bernhoff H, Leijon M. Evaluation of different turbine concepts for wind power. *Renewable and Sustainable Energy Reviews* 2008; **12**(5): 1419–1434.
10. Krogstad PÅ, Lund JA. An experimental and numerical study of the performance of a model turbine. *Wind Energy* 2012; **15**(3): 443–457.
11. Chen Z, Guerrero JM, Blaabjerg F. A review of the state of the art of power electronics for wind turbines. *IEEE Transactions on Power Electronics* 2009; **24**(8): 1859–1875.
12. Aho J, Buckspan A, Laks J, Fleming P, Jeong Y, Dunne F, Churchfield M, Pao L, Johnson K, *Tutorial of wind turbine control for supporting grid frequency through active power control: preprint*, National Renewable Energy Laboratory (NREL), Golden, CO. 2012.
13. Betz A. *Introduction to the Theory of Flow Machines*. Elsevier, 2014.
14. Hackman ST. *Production Economics: Integrating the Microeconomic and Engineering Perspectives*. Springer-Verlag: Heidelberg, 2008.
15. Lee G, Ding Y, Genton MG, Xie L. Power curve estimation with multivariate environmental factors for inland and offshore wind farms. *Journal of the American Statistical Association* 2015; **110**(509): 56–67.
16. Banker RD, Charnes A, Cooper WW. Some models for estimating technical and scale inefficiencies in data envelopment analysis. *Management Science* 1984; **30**(9): 1078–1092.
17. Varian HR. The nonparametric approach to demand analysis. *Econometrica* 1982; **50**: 945–973.
18. Carvalho A, Gonzalez MC, Costa P, Martins A. Issues on performance of wind systems derived from exploitation data. *Proceedings of the Industrial Electronics, 2009. IECON'09. 35th Annual Conference of IEEE*, Porto, Portugal, 2009; 3599–3604.

19. Pieralli S, Ritter M, Odening M. Efficiency of wind power production and its determinants. *Energy* 2015; **90**: 429–438.
20. Aigner D, Lovell CAK, Schmidt P. Formulation and estimation of stochastic frontier production function models. *Journal of Econometrics* 1977; **6**(1): 21–37.
21. Olesen OB, Ruggiero J. Maintaining the regular Ultra Passum law in data envelopment analysis. *European Journal of Operational Research* 2014; **235**(3): 798–809.
22. Hwangbo H, Johnson AL, Ding Y. Power curve estimation: functional estimation imposing the regular ultra passum law. *Working Paper* 2015. Available at SSRN: http://ssrn.com/abstract=2621033.
23. Kuosmanen T. Representation theorem for convex nonparametric least squares. *The Econometrics Journal* 2008; **11**(2): 308–325.
24. Hall P, Simar L. Estimating a changepoint, boundary, or frontier in the presence of observation error. *Journal of the American Statistical Association* 2002; **97**(458): 523–534.
25. Hastie T, Tibshirani R, Friedman J. *The Elements of Statistical Learning (2nd Edition, chapter 7)*. Springer-Verlag: New York City, 2009.
26. Staffell I, Green R. How does wind farm performance decline with age? *Renewable Energy* 2014; **66**: 775–786.
27. Khalfallah MG, Koliub AM. Effect of dust on the performance of wind turbines. *Desalination* 2007; **209**(1): 209–220.