A Bayesian hierarchical model to understand the effect of terrain on wind turbine power curves

Abhinav Prakash, Se Yoon Lee, Xin Liu, Lei Liu, Bani Mallick, and Yu Ding

***Abstract*—This paper is concerned with explicitly modeling the effect of terrain on wind power curves. Terrain characteristics are spatially-varying but temporally constant, whereas other power curve-affiliating variables such wind speed, temperature, and wind power vary both spatially and temporally. In order to effectively model such two modes of variation in the data, we employ a Bayesian hierarchical model (BHM) that connects the terrain characteristics with the parameters in a power curve. BHM jointly models the data from all turbines on a wind farm for attaining the turbine-specific, terrain-incorporating power curves. Our analysis shows that, out of the three terrain variables available in our data, ruggedness has the strongest effect on the power curve. We also evaluate the applicability of using the resulting power curve model for turbines on a different terrain and find that incorporating terrain information explicitly is beneficial. The specific BHM mechanism of using terrain information leads to over 7–10% improvement over the group averaging approach.**

***Index Terms*—Logistic curve, Multidimensional power curve, Power curve transferability, Spatio-temporal dataset, Wind farm.**

I. Introduction

WIND energy is one of the fastest-growing clean energy sources and a key to achieving carbon net zero for the global community. By the end of 2022, wind energy contributed to 10.2% of the total electricity production in the US [1]. Playing a critical role in wind energy operations and planning is the wind turbine power curve, which maps the relationship of wind speed and other environmental variables to wind power production [2], [3]. The physical law for wind power generation is given as follows:

$$y = \frac{1}{2} C_p \rho A V^3, \tag{1}$$

where, y is the wind power, C_p is called the power coefficient, ρ is the air density, A is the area swept by the turbine blades, and V is the wind speed. The physics-based model has its limitation mainly because the power coefficient C_p is not a fixed constant but depends on factors such as tip speed ratio, wind attack angles, as well as inflow conditions (which in turn depends on terrain). An analytical expression for C_p is not available for direct use, which means that C_p will have to be estimated through a data-driven approach anyway.

Because of such complexities and also to account for actual turbine operation in which active pitch control used in wind turbines alters the wind-power relationship, International Electrotechnical Commission (IEC), the international governing body, recommended a data-driven method to estimate the wind power curve from the turbine's operational data [4], nicknamed the *binning method* [2]. The binning method uses primarily the wind speed to estimate the power curve. Over the last decade, there have been many new developments in data-driven wind power modeling to include other environmental variables such as air density, turbulence intensity, and humidity in the power curve model (for example, [5]–[7]) and made significant improvements. However, there are other intrinsic (turbine-related) and extrinsic (environment-related) factors that affect the power production, but has not been studied in the data science literature, mainly due to the lack of data availability. One such extrinsic factor is the terrain complexity at a wind turbine's location.

It is understood that the underlying local terrain for a wind turbine affects its power output; see [8]–[11] for details. The main reason for terrain entangling with power production is its effect on environmental variables, such as wind speed, wind shear, and turbulence intensity. Although the wind speed measurement is used in power curve estimation, it is only a point measurement. It does not capture fully the wind profile and inflow condition faced by the turbine blades. The terrain acts as an intermediate between wind and turbines—different terrains cause wind profiles and inflow conditions to be different. Thus, terrain modeling can provide valuable insight on wind power production.

The literature on understanding the effect of terrain predominantly focuses on physics-based computational models, which attempt to provide the understanding through modeling the atmospheric boundary layer under different terrains. Sempreviva et al. [8] study the effect of surface roughness on wind speed and wind shear using atmospheric internal boundary layer (IBL). Fragoulis and Fragoulis [9] study fatigue loading of wind turbines under different terrains. Tian et al. [11] conducted an experiment in a large-scale aerodynamic/atmospheric boundary layer (AABL) wind tunnel for understanding the effect of hilly terrain on wind flow characteristics. Tian et al. [11] concludes that on a high slope hilly terrain, the mean wind velocity dramatically decreases, whereas for a low slope hill, the change in mean wind velocities is small. Han et al. [10] combine both wake model

Abhinav Prakash is with Walmart Inc., Sunnyvale, CA 94086 and Se Yoon Lee is with Johnson & Johnson, Irvine, CA, 92618. Prakash and Lee worked on this paper during their doctoral study at Texas A&M University. Xin Liu and Lei Liu are with Beijing TianRun New Energy Investment Corporation Ltd., Beijing 100029, China. Bani Mallick is with Department of Statistics, Texas A&M University, College Station, TX 77843. Yu Ding is with the H. Milton Stewart School of Industrial and Systems Engineering at Georgia Institute of Technology, Atlanta, GA 30332. Corresponding author: Yu Ding (yu.ding@isye.gatech.edu). This work was partially supported by NSF under grants IIS-1741173l, CCF-1934904, and Goldwind under contract no. M1900794.

and terrain model to evaluate adequacy of a wind farm and reliability of its wind turbines. Although the aforementioned papers study the effect of terrain, they do so at a small scale or in an experimental set up. Thus, the literature lacks large-scale studies on the effects of terrain on operational wind farms, outside of experimental set up. Such large-scale studies, if available, would further expand the boundary of knowledge in terms of our understanding of terrain effects on power curve at a wind farm level. To the best of our knowledge, there is no data-driven/machine learning power curve model that explicitly incorporates terrain characteristics. By "data-driven/machine learning power curve," we refer to the type of studies that use site-specific operational data for modeling, such as [5], [6] or generally as those in Chapter 5 of [2], which are also coded in both `R` and `Python` packages [12], [13].

The use case for modeling terrain goes beyond just understanding its effect on power curve. Terrain modeling would also prove beneficial in transferring power curve to a new wind farm site for which one does not yet have wind and power data, but already have terrain measurements. Regressing on terrain measurements helps re-calibrate wind power production using power curves estimated from other sites and thus optimize turbine placement and operation for the new site. Hammer and Barber [14] recently attempted to understand how the power curve can be reused (or transferred) from one site to another. They did not model the terrain explicitly but used an *ad hoc* approach, which is to take the average power curve of a few nearby turbines and transfer it to a new location. Hammer and Barber [14] tested their hypothesis that the accuracy of power prediction decreases when the distance between the training and test turbines increases, but did not find anything conclusive.

The recent interest in reusing power curves on different wind farms (with various terrain characteristics) and the lack of a consistent mechanism to improve reusability of power curves further increases the importance of terrain modeling. When we contemplate why terrain-incorporating power curve models are rare, despite the strong interest, a reason coming immediately to mind is the absence of open source terrain data of wind farms in the existing Open Data Resources [15]. We are lucky to have a dataset embodying terrain measurements in addition to wind and power.

Of course, having terrain data is only a necessary condition. To incorporate the terrain characteristics in a power curve model needs data science innovation. Unlike wind and power data used in power curve modeling, which vary both in time and with turbine sites, terrain measurements do not change in time but only with sites. For a specific turbine, its surrounding terrain remains a constant, so that the machine learning routines used to learn the wind-to-power relationship cannot be used to learn the terrain-to-power relationship.

To solve this type of problem, two sets of statistical learning methodologies are useful: the Bayesian hierarchical models [16] and the mixed effect models [17]. Between them, Bayesian hierarchical model provides uncertainty quantification and has been found to provide a more tractable inference in many problem settings [16], [18], [19]. For this reason, we devise a Bayesian hierarchical model, consisting of a logistic curve, which has wind and other environmental variables as inputs, for modeling individual power curves, and then let the terrain variable influence the parameters of the logistic power curve, so that the terrain effect is also incorporated. We use an onshore wind farm dataset of 66 turbines with terrain measurements to estimate the turbine-specific, terrain-incorporating power curve model. We further study the usability of the resulting power curve model on a different terrain as well as the impact of the terrain characteristics on a turbine's wind power production.

The main innovation in this work is to jointly model wind power curves and their terrain information using a Bayesian hierarchical model (BHM). To our best knowledge, our BHM-based power curve model is the first turbine-specific and terrain-incorporating power curve model. This model allows us to understand quantitatively the effect of terrain on power curves. Compared with conventional farm-based averaging approaches, we found that including the terrain information in the power curve model can improve the accuracy of wind power estimation, by as much as 10%, when the power curve is used on a set of hold-out turbines that are on different terrains. Our analysis results also indicate that the ruggedness index (RIX) near a turbine is the most influential terrain factor affecting the wind power curve, out of the three terrain variables measured in this study. The RIX effect is negative, which is to say, as the terrain becomes more rugged, the power productivity of a turbine decreases.

The rest of the paper is organized as follows. Section II describes the dataset in detail. Section III introduces the model and the inference technique. Section IV contains the results of our study. Section V concludes the paper.

II. Data description and preprocessing

The data comprises of 66 turbines from an onshore wind farm. Each turbine consists of high-frequency observations (6 to 8 observations per minute) of wind speed (V), wind direction (D), and ambient temperature (T). Using the industry practice, we convert the high-frequency data to 10-minute averages. We also compute the turbulence intensity (I) of the wind speed and standard deviation of wind direction (sdD) using the high-frequency data. In this study, we use one year worth of data for each turbine, which is sufficient to learn a good representation of the power curve. After converting one year data to 10-minute averages, we have about 35,000 to 45,000 data points for each turbine. The number of data points are fewer than the total number of 10-minute intervals in a year (which is 52,560), as there are many time points (10-min intervals) where the turbine was not operating (no power production).

A. Terrain data

The terrain characterization was conducted by the wind farm's owner/operator, following the international standards in IEC 61400-12-2 [20], and then the data was given to our team for the modeling use. To explain the data briefly, the terrain is characterized by three variables: slope of the terrain (in degrees), ruggedness index or RIX (dimensionless),

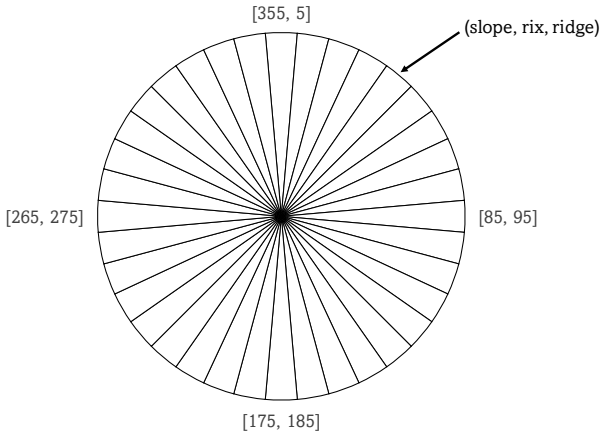

Fig. 1. A bird's eye view on how the terrain is characterized at a single turbine. The three terrain variables—slope, RIX, and ridge—are characterized for each ten-degree sector.

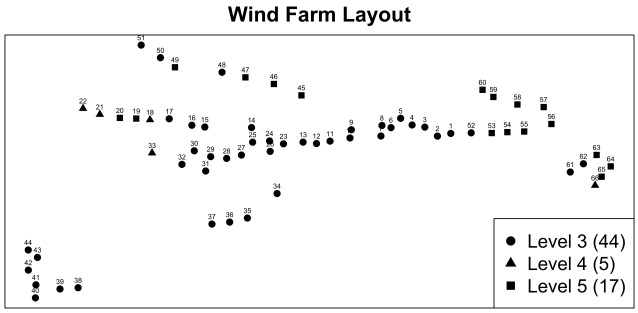

Fig. 2. Layout of the wind turbines with their respective categorical terrain level. The numbers in the parenthesis in the legend denote the number of turbines in a given terrain category.

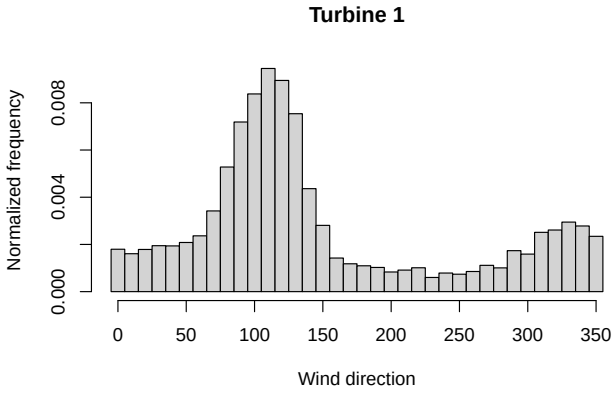

Fig. 3. The frequency of wind direction data for Turbine 1.

and maximum height of the ridge (in meters). The slope of the terrain measures the inclination of the terrain in degrees. The ruggedness index of the terrain is the fraction of the surrounding terrain that is steeper than a given critical slope; please see Mortensen et al. [21] for a detailed description of RIX measurement. Maximum ridge height, as the name suggests, is the height (in meters) of the highest ridge on the terrain.

The terrain variables are mapped for each turbine and expressed at different levels of details. At the most detailed level, the surrounding at a given turbine is divided into 36 sectors of ten degrees angle each and the three terrain variables are measured in every sector. Fig. 1 shows the top-view of a single turbine with its terrain-measurement sectors. For a turbine having altogether 36 such sectors, its terrain is fully characterized by a 36×3 matrix.

At the next level of details, for each of the terrain variables, the corresponding values in the 36 sectors is averaged to get an overall score for that terrain variable. After such aggregation, the terrain is characterized by a 3×1 vector which has the average values of slope, RIX, and ridge. At the coarsest level of representation, the terrain variable average scores are further aggregated to place the corresponding terrain into a category represented by an integer; the larger, the more complex the terrain is. The specific rules for aggregation and terrain classification can be found in Section 6.3.4 of IEC 61400-12-2 [20]. The 66 turbines on the specific wind farm we studied are associated with three of the categories, taking the values of 3, 4 and 5, respectively. In summary, we have the terrain data at three different granularities—the sector-wise terrain characteristics, variable-wise terrain scores, and the overall terrain category. The layout of the wind farm with each turbine's overall terrain category is shown in Fig. 2.

We admit that we did not quite figure out how to use the finest terrain information, i.e., the sector-wise measurements expressed in the 36×3 matrix, in a power curve model. Instead, we use the sector-wise measurements to compute a weighted average measurement for each terrain variable and use the resulting 3×1 vector (the second level of details) in our model. The weights are based on the frequency of wind direction data for each sector. The intuition behind this weighting is that the proportion of time for which the wind is blowing over a certain sector is different. Thus, in order to give a proper representation to the terrain characteristics of every sector, we weight the terrain measurements based on the frequency of data flowing through that sector. Let B be the set of all the direction sectors, i.e., B has 36 elements, one corresponding to a sector as shown in Fig. 1. Let $f_{b,i}$ be the normalized frequency of wind direction when the direction is in sector $b \in B$ and for Turbine i, such that $\sum_{b \in B} f_{b,i} = 1$. The wind direction data for each turbine is used to calculate $f_{b,i}$ for that turbine. The normalized frequency $f_{b,i}$ for one of the turbines (Turbine 1) is shown in Fig. 3.

Let $\boldsymbol{r}_{b,i}$ be the 3×1 vector of terrain measurement for direction sector b for Turbine i. Then, the weighted terrain measurement is given by:

$$\boldsymbol{r}_i = \sum_{b \in B} f_{b,i} \times \boldsymbol{r}_{b,i}, \tag{2}$$

where $\boldsymbol{r}_i$ is the weighted terrain vector with three components for Turbine i.

We also experimented by taking a simple average (unweighted) of the terrain variables and found an insignificant reduction in prediction accuracy when predicting at a test site using leave-one-out cross-validation (LOO-CV); this is to be

TABLE I
THE FIVE-FOLD CROSS-VALIDATION RMSE USING DIFFERENT COMBINATION OF INPUTS FOR TURBINE 1. THE NUMBERS IN THE PARENTHESIS INDICATE THE STANDARD ERROR FOR RMSE COMPUTED OVER FIVE FOLDS. THE INPUT PAIR (V, T) IS THE MOST PARSIMONIOUS MODEL THAT IS WITHIN ONE STANDARD ERROR FROM THE LOWEST RMSE MODEL (V, T, D).

Inputs	RMSE (Standard Error)
V	0.0467 ($\pm$0.0010)
$\mathbf{V}, \mathbf{T}$	0.0344 ($\pm$0.0018)
V, T, D	0.0334 ($\pm$0.0017)
V, T, D, sdD	0.0345 ($\pm$0.0013)
V, T, D, sdD, I	0.0377 ($\pm$0.0014)

explained in detail in a later section. Therefore, we use the weighted terrain averages as the default terrain representation for the rest of the analysis. The weighted measurement reflects the *effective* terrain measurement at a turbine.

B. Important input variables

We conduct a variable selection to find the important subset of inputs among the five inputs described earlier. To this end, we use a forward-stepwise subset selection using the kNN power curve function in the `DSWE R` package [12]. The choice of kNN is unimportant. Had we used other power curve models in the `DSWE` package, the important inputs selected remain the same.

The model is evaluated based on root mean square error (RMSE) criterion. Hastie et al. [22, Chapter 3] recommend using the least complex model that is within one standard error of the best model as per the cross-validation RMSE, where the standard error is the standard deviation of the RMSE based on different folds for a particular model (combination of inputs). Table I shows the five-fold cross-validation error for different subset of inputs for Turbine 1; other turbines also show a similar pattern. The values next to the RMSE in the parenthesis are the standard error for that particular subset of inputs. We find that the wind speed and ambient temperature are two most important variables for power curve modeling. Adding a third variable sometimes helps but generally does not result in significant reduction in RMSE. Hence, for our modeling purpose, we choose the wind speed and ambient temperature as the input variables.

C. Autocorrelation

A common assumption underlying most statistical learning models is that the data is independent and identically distributed (i.i.d.) [22]. In the case of wind power curve data, we observe that the data is serially correlated in time, which invalidates the i.i.d. assumption.

Prakash et al. [7] used a Bayesian sub-sampling scheme known as thinning to reduce the autocorrelation in the data and found improvement in accuracy of their model. The thinning approach removes data points from the dataset such that the remaining data can be treated as independent. The time lag between the indices of two data points in the thinned dataset is referred to as the thinning number. For example, if the original dataset is indexed as $\{(x_1, y_1), (x_2, y_2), (x_3, y_3), \dots\}$, using a thinning number of T would result in a thinned dataset with the following data points: $\{(x_1, y_1), (x_{T+1}, y_{T+1}), (x_{2T+1}, y_{2T+1}), \dots\}$. Prakash et al. [7] propose a rule to compute the thinning number for a given dataset based on the strength of the autocorrelation in the data. Using their proposed approach, we compute the thinning number for each of the turbine in the wind farm, and then use the final thinning number as the average of the individual thinning numbers rounded up to the next integer. The average thinning number turns out to be 15. This resulting thinned dataset for each of the turbines is then used for subsequent modeling. Fig. 4 shows the autocorrelation function (ACF) plot for wind speed before and after thinning for Turbine 1 of our dataset. The figure clearly shows that wind speed is autocorrelated but the strength of autocorrelation is significantly reduced after thinning the dataset. Other turbines in our dataset show a similar pattern.

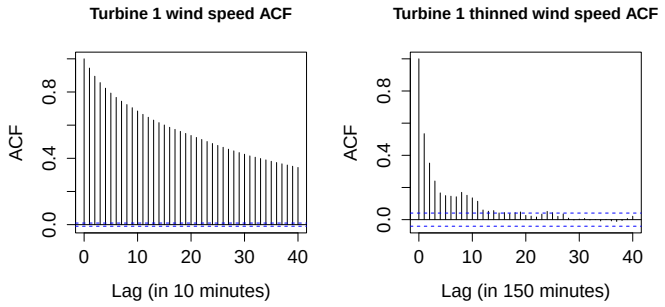

Fig. 4. Autocorrelation function (ACF) plot of wind speed for Turbine 1 before and after thinning.

III. BAYESIAN HIERARCHICAL MODEL

We need to model all the wind turbines jointly in order to estimate the effect of terrain on the power curves. Bayesian hierarchical models (BHM) [16] provides a natural way to model group data, where by "group", we refer to the grouping of power curves in accordance to the turbines associated with their respective terrain.

A. Model

We start off by defining the notations for the data. Let $\boldsymbol{x}_{i,j} = (v_{i,j}, t_{i,j})^\top$ be the jth input point (a vector with two components—wind speed and temperature, respectively) for Turbine i. Let N denote the total number of turbines (66), that is, $i = 1, \dots, N$. Let M be the total number of data points for a turbine after thinning, that is, $j = 1, \dots, M$. Therefore, the total number of data points considering all the turbines is $M \times N$. Let $y_{i,j}$ be the normalized power response corresponding to $\boldsymbol{x}_{i,j}$. The normalization is done by dividing the power by the rated power of the turbines. Let $\boldsymbol{r}_i$ be the vector of weighted terrain measurements for Turbine i; see Section II-A for details.

The Bayesian hierarchical model is given as follows:

$$y_{i,j} = p(\boldsymbol{x}_{i,j}; \boldsymbol{\theta}_i) + \epsilon_{i,j} \ , \tag{3}$$

$$\boldsymbol{\theta}_i = \boldsymbol{q}(\boldsymbol{r}_i) + \boldsymbol{\epsilon}_{\theta,i} \ , i = 1, \cdots, N, \ j = 1, \cdots, M. \tag{4}$$

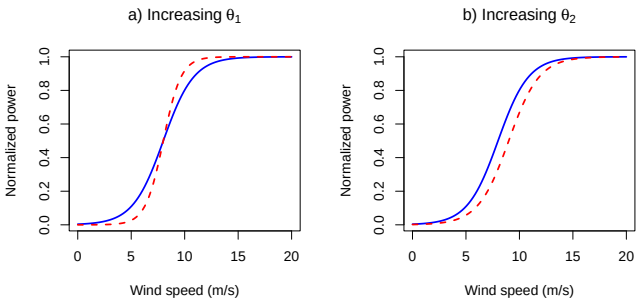

Fig. 5. A logistic power curve and the effect of increasing: a) θ_1 (slope of the inflection point); and b) θ_2 (location of the inflection point). The blue solid line represents the power curve before the change, whereas the red dashed line represents the power curve after the change.

where $p(\cdot)$ is the multi-dimensional power curve characterized by a vector of turbine-specific parameters $\boldsymbol{\theta}_i$ for Turbine i, $\boldsymbol{q}(\cdot)$ is a vector-valued function of the terrain variables $\boldsymbol{r}_i$, mapping the terrain characteristics to the power curve parameters $\boldsymbol{\theta}_i$, and $\boldsymbol{\epsilon}_\theta$ is a random vector of the same size as $\boldsymbol{\theta}_i$, representing the variance in $\boldsymbol{\theta}_i$'s that is not explained by the terrain variables. The size of $\boldsymbol{\theta}_i$, and thus that of $\boldsymbol{q}(\cdot)$ and $\boldsymbol{\epsilon}_\theta$, depends on the choice of the function, $p(\cdot)$.

We model the power curve function, $p(\cdot)$, as a two-parameter logistic curve defined as follows:

$$p(v;\theta_1,\theta_2) = \frac{L}{1+\exp(-\theta_1(v-\theta_2))}\ , \tag{5}$$

where L is the upper asymptote of the logistic curve, θ_1 and θ_2 are the parameters of the logistic curve representing the slope and location of the inflection point of the curve, respectively, and v is the wind speed. We would fix the upper asymptote to $L = 1$, as our dataset contains normalized power that remains in $[0, 1]$. The logistic model is a widely used parametric function to model power curves [23]–[27]. The advantage of using the logistic curve is two-fold: 1) It mimics the S-shape of a power curve while keeping the number of parameters low; 2) The parameters θ_1 and θ_2 also have physical meaning in a power curve. The effect of changing the parameters on the curve is shown in Fig. 5. Increasing θ_1, which denotes the slope of the inflection point of the logistic curve, would result in an increased power productivity. On the other hand, increasing θ_2 would shift the power curve towards the right, decreasing the power productivity.

Most of the current power curve models are nonparametric in nature, including the binning method (the industry standard) [4], the tempGP model that was recently proposed [7], and many others described in Chapter 5 of [2]. For this study, however, we chose a parametric power curve model—the logistic curve. The reason of our choice lies in the fundamental difficulty in modeling a group of nonparametric power curves as a function of the terrain variables. For a parametric function, the parameters of that function are explicitly given. As such, one can regress the parameters further over the terrain characteristics. A nonparametric function, on the other hand, does not have the type of model parameters as in parametric functions. In fact, the degrees of freedom of a nonparametric function increases with data, meaning that the model's effective number of parameters is changing and thus making it nontrivial to model the dependency of nonparametric power curves on terrain characteristics.

Given the underlying difficulty in employing a nonparametric model, the logistic curve makes a reasonable candidate for modeling the individual power curves in a BHM. The low number of parameters in individual power curves keeps the computation tractable even when we are jointly modeling all the turbines. The physical significance of the parameters helps us understand the effect of terrain on the power curves, as the parameters are modeled as a function of the terrain.

We model the dependence of the power curve parameters on the terrain variables as a linear function in terrain, given as follows:

$$\theta_1 = \alpha_1 + \boldsymbol{\beta}_1^\top \boldsymbol{r} + \epsilon_{\theta_1};\quad \theta_2 = \alpha_2 + \boldsymbol{\beta}_2^\top \boldsymbol{r} + \epsilon_{\theta_2}\ , \tag{6}$$

where α's denote the mean-level, $\boldsymbol{\beta}$'s are the coefficients of the linear model, and ϵ_θ's represent the variance in the parameters not explained by the terrain. Combining the linear model for the terrain with the logistic model for power curve also assigns the physical meaning to the coefficients, $\boldsymbol{\beta}$'s.

One of the drawbacks of the logistic curve when used as-is for modeling power curves is that it only models one input variable of wind speed. Wind power, however, is known to also depend on other input variables, in particular on ambient temperature, as established in Section II-B. Thus, it is desirable to have a multi-dimensional power curve model. We modify the one-dimensional logistic curve to include ambient temperature in the following way:

$$p(v,t;\boldsymbol{\theta},\boldsymbol{\eta}) = \frac{1}{1+\exp(-(\theta_1+\eta_1 t)(v-(\theta_2+\eta_2 t)))}\ , \tag{7}$$

where η's are the coefficients for ambient temperature, t. The advantage of Equation (7) is that it still maintains the S-shaped curve in wind speed, and changing the temperature only shifts the location and slope of the inflection point of the S-shaped curve. As such, the power curve now varies with temperature as well. Fig. 6 shows the power curve with two inputs—wind speed and ambient temperature.

Taken together, the final model is given as follows:

$$\begin{aligned}
y_{i,j} &= \frac{1}{1+\exp(-(\theta_{1,i}+\eta_{1,i}t_{i,j})(v_{i,j}-(\theta_{2,i}+\eta_{2,i}t_{i,j})))} + \epsilon_{i,j}\\
\theta_{1,i} &= \alpha_1 + \boldsymbol{\beta}_1^\top \boldsymbol{r}_i + \epsilon_{\theta_1,i},\\
\theta_{2,i} &= \alpha_2 + \boldsymbol{\beta}_2^\top \boldsymbol{r}_i + \epsilon_{\theta_2,i},\\
\epsilon_{i,j} &\sim \mathcal{N}(0,\sigma_\epsilon^2),\quad \epsilon_{\theta_k,i} \sim \mathcal{N}(0,\sigma_{\theta_k}^2)\ \mid\ k=1,2.
\end{aligned} \tag{8}$$

Equation (8) highlights the uniqueness of this model in terms of wind power curve modeling. In the literature, power curves are modeled for individual turbines, and hence, the inputs and the outputs, that is, v, t, and y vary only with time. In the proposed model, the inputs and the output of the power curve, v, t, and y, vary with both space and time—the index j is the time index and indicates the time varying variable, whereas the parameters of the model, $\theta_{1,i}$ and $\theta_{2,i}$, are spatially varying but temporally constant due to their dependence on the terrain.

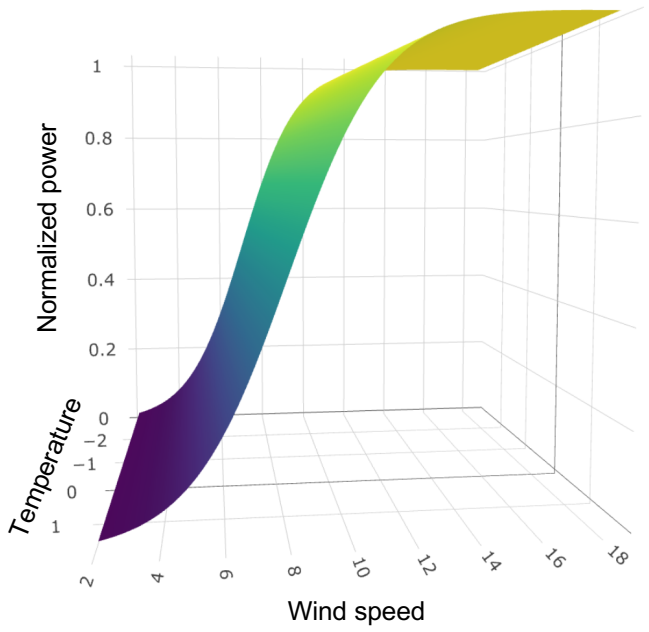

Fig. 6. A two-dimensional logistic power curve with wind speed and ambient temperature as the inputs.

This combination of spatiotemporal characteristics makes the proposed model novel as compared to the conventional power curve models.

B. Prior specification

We need to specify the prior distribution for each of the following parameters: (σ^2_ϵ, η_1, η_2, α_1, α_2, $\boldsymbol{\beta}_1$, $\boldsymbol{\beta}_2$, $\sigma^2_{\theta_1}$, $\sigma^2_{\theta_2}$). We do not need priors for θ's , as given the distributions for α's, $\boldsymbol{\beta}$'s and σ^2_θ's, the distribution for θ's are fully specified.

Table II provides the specification for the priors for each of the parameters. For most of the parameters, we choose an uninformative prior. For $\boldsymbol{\beta}$'s, we choose a ridge prior, which is also commonly used in Bayesian linear models [28] and works well in practice. The prior for $\boldsymbol{\beta}$'s introduces a new parameter τ, which itself is assumed to follow a Half-Cauchy distribution. The prior for η's are assumed to be Gaussian with mean γ and variance σ^2_η. The parameters γ and σ^2_η are modeled using uninformative priors.

Parameter	Prior
σ^2_ϵ	$\pi(\sigma^2_\epsilon) \propto 1/\sigma^2_\epsilon$
α_1, α_2	$\pi(\alpha) \propto 1$
$\sigma^2_{\theta_1}$, $\sigma^2_{\theta_2}$	$\pi(\sigma^2_\theta) \propto 1/\sigma^2_\theta$
η_1, η_2	$\pi(\eta) = \mathcal{N}(\gamma, \sigma^2_\eta)$
γ_1, γ_2	$\pi(\gamma) \propto 1$
$\sigma^2_{\eta_1}$, $\sigma^2_{\eta_2}$	$\pi(\sigma^2_\eta) \propto 1/\sigma^2_\eta$
$\boldsymbol{\beta}_k \ \vert\sigma^2_{\theta_k}, \tau^2_k$	$\mathcal{N}(\mathbf{0}, \sigma^2_{\theta_k}\tau^2_k\mathbf{I}),\ k = 1, 2$
τ^2_k	$\mathcal{C}^+(0, 1),\ k = 1, 2$

TABLE II
PRIOR FOR THE PARAMETERS

With the prior specified, we express the posterior distributions of all the parameters per Bayes' rule as follows, where $\mathcal{L}$ denotes the likelihood function:

$$\begin{aligned}&\pi(\sigma^2_\epsilon, \alpha_1, \alpha_2, \sigma^2_{\theta_1}, \sigma^2_{\theta_2}, \eta_1, \eta_2, \boldsymbol{\beta}_1, \boldsymbol{\beta}_2 | y_{i,j}, \boldsymbol{x}_{i,j}, \boldsymbol{r}_i)\\ &= \mathcal{L}(\sigma^2_\epsilon, \alpha_1, \alpha_2, \sigma^2_{\theta_1}, \sigma^2_{\theta_2}, \eta_1, \eta_2, \boldsymbol{\beta}_1, \boldsymbol{\beta}_2 | y_{i,j}, \boldsymbol{x}_{i,j}, \boldsymbol{r}_i)\\ &\times \pi(\sigma^2_\epsilon)\pi(\alpha_1)\pi(\alpha_2)\pi(\sigma^2_{\theta_1})\pi(\sigma^2_{\theta_2})\pi(\eta_1)\pi(\eta_2)\pi(\boldsymbol{\beta}_1)\pi(\boldsymbol{\beta}_2)\end{aligned} \tag{9}$$

We can further factorize the terms in Equation (9). However, we postpone that when we talk about the sampling scheme. Please note that once the priors and the model for the data have been defined, the posterior is fully specified. Next we explain the numerical procedure to sample from these posteriors.

C. Sampling scheme

We employ a Gibbs sampler [28], a type of Markov chain Monte Carlo (MCMC) method, to sample from the posterior distribution of each of the parameter. The Gibbs sampler is used to sample the parameters sequentially from the full-conditional distribution of the parameters, which is to sample one parameter conditioned on all other parameters and the data. We factorize the terms in Equation (9) and collect the appropriate terms to define the sampling steps as follows:

- Sample from $\pi(\sigma^2_\epsilon | y_{i,j}, p(\boldsymbol{x}_{i,j}; \boldsymbol{\theta}_i, \boldsymbol{\eta}_i)) \propto \prod_{i=1}^N \prod_{j=1}^M \mathcal{N}(y_{i,j} | p(\boldsymbol{x}_{i,j}; \boldsymbol{\theta}_i, \boldsymbol{\eta}_i), \sigma^2_\epsilon) \times \frac{1}{\sigma^2_\epsilon}$
- Sample from $\pi(\theta | y_{i,j}, p(\boldsymbol{x}_{i,j}; \boldsymbol{\theta}_i, \boldsymbol{\eta}_i), \sigma^2_\epsilon, \alpha, \boldsymbol{\beta}, \boldsymbol{r}, \sigma^2_\theta) \propto \prod_{i=1}^N \prod_{j=1}^M \mathcal{N}(y_{i,j} | p(\boldsymbol{x}_{i,j}; \boldsymbol{\theta}_i, \boldsymbol{\eta}_i), \sigma^2_\epsilon) \times \mathcal{N}(\theta | \alpha + \boldsymbol{\beta}^\top \boldsymbol{r}, \sigma^2_\theta)$
- Sample from $\pi(\alpha | \theta, \boldsymbol{\beta}, \boldsymbol{r}, \sigma^2_\theta) \propto \mathcal{N}(\theta | \alpha + \boldsymbol{\beta}^\top \boldsymbol{r}, \sigma^2_\theta) \times 1$
- Sample from $\pi(\boldsymbol{\beta} | \theta, \alpha, \boldsymbol{r}, \sigma^2_\theta, \tau^2) \propto \mathcal{N}(\theta | \alpha + \boldsymbol{\beta}^\top \boldsymbol{r}, \sigma^2_\theta) \times \mathcal{N}(\boldsymbol{\beta} | \mathbf{0}, \tau^2\sigma^2_\theta \mathbf{I})$
- Sample from $\pi(\sigma^2_\theta | \theta, \alpha, \boldsymbol{\beta}, \boldsymbol{r}) \propto \mathcal{N}(\theta | \alpha + \boldsymbol{\beta}^\top \boldsymbol{r}, \sigma^2_\theta) \times \frac{1}{\sigma^2_\theta}$
- Sample from $\pi(\eta | y_{i,j}, p(\boldsymbol{x}_{i,j}; \boldsymbol{\theta}_i, \boldsymbol{\eta}_i), \sigma^2_\epsilon, \gamma, \sigma^2_\eta) \propto \prod_{i=1}^N \prod_{j=1}^M \mathcal{N}(y_{i,j} | p(\boldsymbol{x}_{i,j}; \boldsymbol{\theta}_i, \boldsymbol{\eta}_i), \sigma^2_\epsilon) \times \mathcal{N}(\eta | \gamma, \sigma^2_\eta)$
- Sample from $\pi(\gamma | \eta, \sigma^2_\eta) \propto \mathcal{N}(\eta | \gamma, \sigma^2_\eta) \times 1$
- Sample from $\pi(\sigma^2_\eta | \eta, \gamma) \propto \mathcal{N}(\eta | \gamma, \sigma^2_\eta) \times \frac{1}{\sigma^2_\eta}$
- Sample from $\pi(\tau | \boldsymbol{\beta}, \sigma^2_\theta) \propto \mathcal{N}(\boldsymbol{\beta} | \mathbf{0}, \tau^2\sigma^2_\theta \mathbf{I}) \times \mathcal{C}^+(\tau; 0, 1)$

Except for the full conditional distributions of θ, η, and τ—denoted by $\pi(\theta|-)$, $\pi(\eta|-)$, and $\pi(\tau|-)$, respectively, for simplicity—all other posteriors have closed form, as they have conjugate priors. We do not have conjugate priors for $\pi(\theta|-)$ and $\pi(\eta|-)$ because of the use of the logistic model. For such cases, we can employ the elliptical slice sampler [29] to sample from the posterior distribution. The elliptical slice sampler works when the prior of a parameter, say $\boldsymbol{a}$, is Gaussian and the likelihood has any general form $\mathcal{L}(\boldsymbol{a})$. This is exactly the case for θ's and η's. Thus, the elliptical slice sampler is well suited for the task.

For sampling from $\pi(\tau|-)$, we could use the Metropolis Hastings (MH) algorithm [28]; however, doing so would require us to come up with and tune a proposal distribution. In order to avoid that, we employ another sampling algorithm called slice sampling [30], which does not require a proposal distribution. Slice sampling can sample any random variable for which we know either its probability density function (PDF) or a function that is proportional to its PDF. For τ, we can evaluate the value of $\mathcal{N}(\boldsymbol{\beta}|\mathbf{0}, \tau^2\sigma^2_\theta\mathbf{I}) \times \mathcal{C}^+(\tau; 0, 1)$,

which is proportional to its posterior PDF. Next, we apply our model and sampling scheme to the wind farm dataset.

IV. Case Study

In this case study, the first goal is to understand how terrain variables affect the power curve and the second goal is to evaluate the transferability of the power curve to a new location. For the first goal, we train the BHM on all the turbines. Because the first goal focuses on inference instead of on prediction, we do not need any hold-out turbine. For the second goal, it requires that we hold out a few turbines during training and use the learned BHM to predict the power curve of the hold-out turbines using their terrain information. We perform a leave-one-out cross-validation (LOO-CV), that is, we hold-out one turbine for testing and train the model on the rest of the turbines. We compute the RMSE for the test turbine using the BHM predictions and the actual power data recorded for the respective test turbine. We repeat this process for all the turbines on the wind farm. We compare the BHM with a few alternative methods (described later).

For all analyses, we fix the following specification of the BHM. We collect 15,000 MCMC samples for each of the parameters. The first 5,000 samples are removed as *burn-in* samples. We further thin the remaining 10,000 samples to reduce autocorrelation in the posterior samples.

A. Effect of terrain variables

The importance of individual terrain variables can be assessed by the value of their coefficients, $\boldsymbol{\beta}$. A fair comparison is only possible when the scales of the variables are the same. For this reason we first standardize all the terrain variables by subtracting their respective mean and dividing by their standard deviation.

Next, we train the BHM and collect MCMC samples. Before we delve into the inference, it is important to ensure that MCMC sampling has converged to a stationary distribution. One common approach to assess the convergence is to plot the scans for the parameters, which are the MCMC samples plotted against the sample index (up to 15,000). We did conduct this diagnostics check for all parameters and found (1) all MCMC samples did converge to a stationary distribution and (2) the burn-in samples of 5,000 seems to be sufficient to ensure stationarity of the distribution. Because such diagnostics check is rather standard in Bayesian sampling, we omitted the plots to save space.

After completing the MCMC diagnostics, we remove the burn-in samples and thin the remaining MCMC samples. Fig. 7 presents the boxplots of the three terrain variables in $\boldsymbol{\beta}$ for both the inflection slope (θ_1) and inflection location (θ_2), and Table III presents the mean posterior values for the same coefficients. We note the following from the figure and the table. RIX is the most influential terrain variable impacting the power curve because it has the largest coefficients. Furthermore, an increase in ruggedness of the terrain negatively impacts the power curve, that is, decreasing the productivity of the associated wind turbine. This is so because an increase in ruggedness decreases the coefficient for inflection slope and increases the coefficient for inflection location—both decrease the power productivity. This outcome makes intuitive sense.

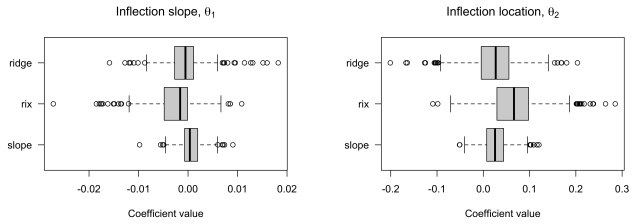

Fig. 7. Box plots for the coefficients of the terrain variables for BHM.

TABLE III
THE MEAN VALUES OF TERRAIN COEFFICIENT VECTORS $\boldsymbol{\beta}_1$ AND $\boldsymbol{\beta}_2$.

Coefficient	Slope	RIX	Ridge
$\boldsymbol{\beta}_1$	0.00063	-0.00273	-0.00065
$\boldsymbol{\beta}_2$	0.02674	0.06884	0.02432

We also notice that the variance of the terrain coefficients for RIX and ridge are larger than terrain slope's coefficient; see Figure 7. We found that the variables RIX and ridge are highly correlated (with a correlation of 0.94). Thus, the model could suffer from variance inflation [22], which often happens when highly correlated inputs are included. Such increase in the variance of estimates does not affect the prediction performance, as an increase in the value of the coefficient for one variable (RIX or ridge) will be compensated by a decrease in the value of the second one. But, using the two highly correlated variables together may result in an inaccurate inference on the importance of individual variables.

In order to better understand the terrain importance, we did further experiments. We trained multiple models with different subsets of the terrain variables and compare the mean and variance of the terrain coefficients. Specifically, we train BHM using one terrain variable at a time and a pair of variables at a time. As such, we have a total of seven models—three one-variable models, three two-variable models and one three-variable model (i.e., the original BHM). For each of the terrain variables—slope, RIX, ridge—we plot the boxplots of their coefficients in Figure 8. All the subplots on the left hand side (LHS) of Figure 8 are for θ_1, whereas the subplots on the right hand side (RHS) are for θ_2. The individual boxplots denote the mean and variance of the terrain coefficient for one terrain variable. For example, the bottom RHS subplot shows the coefficient of terrain ridge under different models with separate subset of terrain variables. Analyzing the bottom plot on the RHS, we note that when ridge and RIX are used together in a model, the variance of ridge's coefficient becomes significantly larger than when ridge is used without RIX. We can make a similar observation for RIX by reading the plot the middle plot on the RHS in Figure 8. All this shows that using multiple correlated variables together in a model is prone to variance inflation. The overall message regarding the dependence of power curve on terrain characteristics, that is, whether an increase in a particular terrain variable increases or decreases the power productivity, remains the

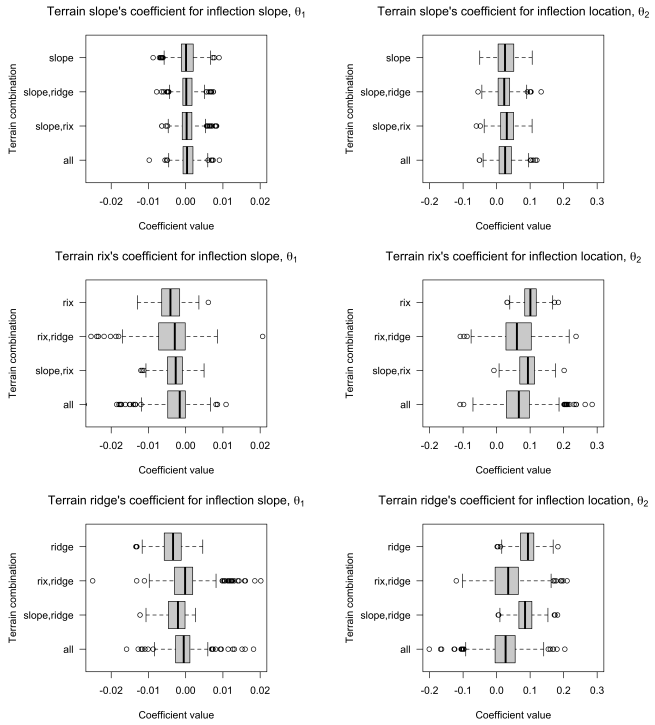

Fig. 8. Comparison of the coefficient's mean and variance for models with different terrain subsets. The left hand side (LHS) plots are for logistic curve parameter θ_1 and the right hand side (RHS) plots are for θ_2.

same for different combinations of terrain inputs. However, the strength of these dependence is impacted (reduced) by using two correlated terrain variables. Thus, one should consider using either RIX or ridge along with slope to get an accurate understanding of the importance of the variables for the terrain modeling task. Given the definition of RIX and ridge in Section II A, we believe that RIX, when compared to ridge, captures a holistic profile of the entire terrain. The maximum ridge height, on the other hand, is based on one point in the terrain profile.

B. Prediction on a test site

The second goal of this case study is to evaluate the transferability of BHM to an unseen turbine. As described earlier, we employ a LOO-CV to assess the performance of BHM on an unseen turbine. The BHM model uses the terrain characteristics at a hold-out (unseen) turbine used to predict its power curve as per Equation (8).

To compare the performance of BHM with binning (the industry benchmark method), we propose multiple *ad hoc* schemes that can be used to transfer the binning power curves to the unseen/test turbine. One of the most straightforward ways to transfer the power curve is to take the average of all the training power curves as the power curve of the test turbine, as done in [14]. This approach, however, does not explicitly take into account terrain information. We denote this method as *Avg-binning* in the results.

The second way to transfer the power curve is to use the overall terrain category information, as described in Section II-A, with the binning power curve. For every test turbine, we find the subset of training turbines that have the same overall terrain category as that of the test turbine. We then use the average power curve of this subset of training turbines as the predicted power curve for the test turbine. This method is denoted as *Terrain-binning* in the results.

The third way to transfer the binning power curve is by considering the k-nearest neighbors of the test turbine, and using the average power curve of the k-nearest neighboring turbines. We denote this method as *kNN-binning*. While implementing this method, we try two different values of k: 10, 20. Thus, we actually have two results for this method, namely *10NN-binning* and *20NN-binning*. We want to, however, stress the following—a critical limitation for this nearest neighborhood approach is that it is only applicable when one is trying to transfer the power curve *within an existing wind farm*, where all the turbines have been constructed and operated, and thus the data are available to fit their power curves. For testing the transferability of power curve to a new wind farm where no turbine has been constructed, such method does not work.

To test the transferability of power curves to a different wind farm, a more realistic neighborhood proxy is to use the *farthest* k ($k = 10$ or 20) neighbors from the target test turbine and use the farthest neighbor to mimic the circumstances where training and test turbines are on different wind farms. This approach is referred to as *10FN-* or *20FN-binning*, where FN means farthest neighborhood.

The results for all the aforementioned methods are tabulated in Table IV. The results show the following. Using terrain information improves the prediction accuracy over other methods. Secondly, using the terrain information in an *ad hoc* way, such as in the method Terrain-binning, can also provide an improvement (about 3%) over just using the average (Avg-binning). Comparing BHM with Terrain-binning, BHM with terrain can deliver an improvement 6–7%, meaning that the use of terrain information by BHM is more effective than simply using it for group averaging. Figure 9, top panel, presents the estimated wind power curve and observed sample data from Turbine 1. Since BHM uses two dimensional inputs—wind speed and temperature, we fix the temperature to a small interval $[-15, -14]$ to get the power-vs-wind speed plot that is comparable with other methods. We observe that the BHM predictions (denoted by hollow circles) are much closer to the data (denoted by $*$). Such a good fit is in fact true for BHM by and large. To visualize the goodness-of-fit across all temperature values in a dimension agnostic way, we plot, in the bottom panel of Figure 9, the actual-vs-predicted power for a set of data points sampled across all temperature values from Turbine 1. We note that the prediction output from BHM are in general closer to the 45° dashed diagonal line, confirming a better fit of BHM.

The nearest neighbors of a turbine carry valuable information about an unseen location, and as a result, their prediction accuracy can be almost as good as that of a terrain-incorporating power curve. What this suggests is that suppose one has a set of turbines already constructed and wants to know how a supposed turbine would behave anywhere in that neighborhood, one does not have to *explicitly* worry about the

TABLE IV
THE AVERAGE LEAVE-ONE-OUT CROSS-VALIDATION (LOO-CV) RMSE AND MAE USING DIFFERENT METHODS.

Method	Average LOO-CV RMSE	Increase Over Best (RMSE)	Average LOO-CV MAE	Increase Over Best (MAE)
BHM	**4.241**	0.00%	**2.788**	0.00%
Avg-binning	4.613	8.77%	3.086	10.68%
Terrain-binning	4.485	5.75%	2.984	7.04%
10NN-binning	4.329	2.07%	2.831	1.56%
20NN-binning	4.411	4.01%	2.920	4.73%
10FN-binning	5.139	21.17%	3.502	25.60%
20FN-binning	5.065	19.43%	3.449	23.73%

terrain characteristics. The neighborhood turbines *implicitly* incorporate such information in their power curves.

The more practical need for power curve transfer is to a different location. The Terrain-binning can be used for such purpose. But if one really wants to use a neighborhood method, then a more realistic comparison is with the farthest neighborhood method. As shown in the comparison between BHM and 10FN/20FN-binning, the FN methods are not competitive—if the training turbines are far away (although in the same wind farm), the prediction accuracy can drastically deteriorate as compared to a terrain-based model like BHM. One would be better off simply using the Terrain-binning approach.

In summary, BHM performs better than all the *ad hoc* methods used in this case study both in terms of RMSE and MAE. Comparing with the methods of averaging all binning-based power curves or those power curves associated with a specific terrain, respectively, BHM is about 6–9% better in terms of RMSE or 7–11% better in terms of MAE. Such consistent results confirm that using the terrain characteristics provides valuable information for predicting the power curve of an unseen site. The way that BHM uses the terrain information is better than the *ad hoc* approaches

C. A two-step point estimation for fast computation

In this work, we considered a Bayesian hierarchical model to jointly model all the turbines. The advantage of the model is its ability to naturally quantify uncertainty and model multiple turbines jointly. The disadvantage of such a model is its long computation. The Bayesian models are generally slower to compute than their frequentist counterpart. Using a computer with Apple silicon chip, running an instance of our proposed model took about 7,200 seconds (2 hours) on a single core, whereas the binning method took 15 seconds for one turbine, or 15 minutes for all the turbines.

In order to provide a faster solution for practical purposes, we experimented by converting our model to a two-step point estimation using a frequentist approach. In particular, we estimate the parameters of the logistic curves, $\theta_{1,i}, \theta_{2,i}$, separately for each turbine i using a least-squares approach, that is, minimize the sum residuals to obtain θ's:

$$(\hat{\theta}_{1,i}, \hat{\theta}_{2,i}) = \arg\min \sum_{j=1\ldots M} (y_{i,j} - p(\boldsymbol{x}_{i,j}; \theta_{1,i}, \theta_{2,i}))^2 \quad (10)$$

These separately estimated parameters θ's are then passed to a linear regression model as response, with terrain variables

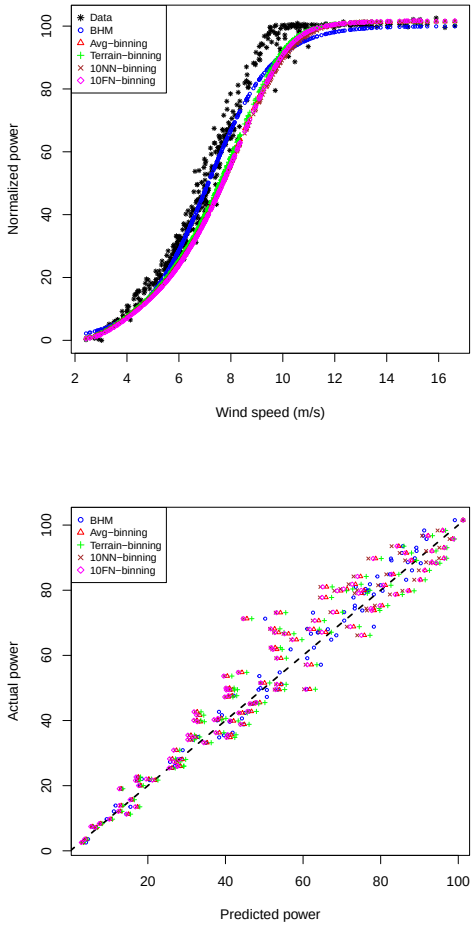

Fig. 9. Goodness-of-fit plot for various power curve methods. Top panel: power-vs-wind speed plot after fixing temperature interval. Bottom panel: actual-vs-predicted power plot across all temperature values, and the closer a data point is to the 45° dashed diagonal line, the better fit is the corresponding method.

$\boldsymbol{r}_i$ as the inputs for estimating the coefficients $\boldsymbol{\beta}$'s. When we used this two-step model to conduct a LOO-CV evaluation, as we have done in the previous section, we found that the prediction accuracy of the two-step approach is about one percentage worse than the original BHM. In general, a two-step procedure is known to introduce bias in the estimates and also provides only a point estimate. Given the large amount data in our study, the two-step approach performs reasonably well while reducing the computation time significantly. The entire computation of the frequentist two-step approach took 15 minutes on a single CPU, very much comparable to those binning-based alternatives in Table IV . Hence, when there is a shortage of computational resources, practitioners can resort to the two-step approach for ensuring power curve transferability with help of the terrain information.

V. CONCLUSION

Wind turbines are installed over terrains with various degrees of complexity. However, their power productivity as a function of the terrain is not yet well quantified using actual

power production data. This work is a step towards filling this gap in knowledge. The main innovation is the Bayesian hierarchical power curve model that incorporates terrain information and jointly models all the turbines on a wind farm. This new approach is in contrast with existing wind power curves in the literature, which model each turbine separately and do not consider explicitly the terrain information.

By developing the Bayesian hierarchical model, the paper presents one of the first terrain-incorporating power curves, leading to a better understanding of the relationship between wind turbine power generation and terrain characteristics. Such model produces more accurate power predictions when being reused for turbines on a different terrain. Specifically, the explicit use of terrain information through BHM improves the prediction performance by 7–10% in terms of MAE as compared to a naive transfer of the binning power curve that is the average of the power curves either over the whole wind farm or in the respective terrain category. Our model also demonstrates that the ruggedness of the terrain seems to have the largest impact on power productivity and a rugged terrain reduces the power production efficiency of a wind turbine.

A few assumptions that we make in this study are: employing a parametric power curve power and a linear terrain dependence on the power curve parameters. Parametric models (specifically, the linear models) are good for model interpretability, but have low degrees of freedom and cannot model complex relationships between variables. In order to check if the terrain variables have a more complex relationship with the power curve than what a linear model can capture, a natural future work would be to extend the current work to include nonlinear relationship between the terrain variables and the power curve parameters. A further enhancement is to use nonparametric power curve models. The main challenge is how to incorporate the hierarchical terrain information in nonparametric power curve models. Such a model would have large degrees of freedom, which makes the inference less tractable and causes the resulting model prone to overfitting. New research is needed to overcome these challenges to render a nonparametric hierarchical terrain-incorporating power curve model practical and useful.

REFERENCES

[1] EIA, "US Energy Information Administration," Webpage: https://www.eia.gov/energyexplained/wind/electricity-generation-from-wind.php, Accessed on: 2023-06-23, 2023.

[2] Y. Ding, *Data Science for Wind Energy*. Boca Raton, FL: Chapman & Hall, 2019.

[3] Y. Ding, N. Kumar, A. Prakash, A. E. Kio, X. Liu, L. Liu, and Q. Li, "A case study of space-time performance comparison of wind turbines on a wind farm," *Renewable Energy*, vol. 171, pp. 735–746, 2021.

[4] International Electrotechnical Comission (IEC) 61400-12-1, *Wind Energy Generation Systems—Part 12-1: Power Performance Measurements of Electricity Producing Wind Turbines (Edition 2.0)*, Geneva, Switzerland, 2017.

[5] R. J. Bessa, V. Miranda, A. Botterud, J. Wang, and E. M. Constantinescu, "Time adaptive conditional kernel density estimation for wind power forecasting," *IEEE Transactions on Sustainable Energy*, vol. 3, no. 4, pp. 660–669, 2012.

[6] G. Lee, Y. Ding, M. G. Genton, and L. Xie, "Power curve estimation with multivariate environmental factors for inland and offshore wind farms," *Journal of the American Statistical Association*, vol. 110, no. 509, pp. 56–67, 2015.

[7] A. Prakash, R. Tuo, and Y. Ding, "The temporal overfitting problem with applications in wind power curve modeling," *Technometrics*, vol. 65, no. 1, pp. 70–82, 2023.

[8] A. M. Sempreviva, S. E. Larsen, N. G. Mortensen, and I. Troen, "Response of neutral boundary layers to changes of roughness," *Boundary-Layer Meteorology*, vol. 50, no. 1, pp. 205–225, 1990.

[9] A. Fragoulis and A. Fragoulis, "The complex terrain wind environment and its effects on the power output and loading of wind turbines," in *35th Aerospace Sciences Meeting and Exhibit*, 1997, p. 934.

[10] X. Han, J. Guo, and P. Wang, "Adequacy study of a wind farm considering terrain and wake effect," *IET generation, transmission & distribution*, vol. 6, no. 10, pp. 1001–1008, 2012.

[11] W. Tian, A. Ozbay, and H. Hu, "Terrain effects on characteristics of surface wind and wind turbine wakes," *Procedia Engineering*, vol. 126, pp. 542–548, 2015.

[12] N. Kumar, A. Prakash, and Y. Ding, *DSWE: Data Science for Wind Energy, `R` Package*. The Comprehensive `R` Archive Network (`CRAN`) https://cran.r-project.org/web/packages/DSWE/index.html, 2020.

[13] P. Kumar, A. Prakash, and Y. Ding, *DSWE: Data Science for Wind Energy, `Python` Package*. The `Python` Package Index (`PyPI`) https://pypi.org/project/dswe/, 2022.

[14] F. Hammer and S. Barber, "Transferability of site-dependent wind turbine performance predictions using machine learning," *Journal of Physics: Conference Series*, vol. 2151, no. 1, p. 012006, 2022.

[15] IEA, "International Energy Agency Wind Task 43 Open Data Resources," Webpage: https://www.ieawindtask43.org/proceedings-work-products/open-data-resources, Accessed on: 2022-12-18, 2022.

[16] P. D. Congdon, *Applied Bayesian Hierarchical Methods*. Chapman and Hall/CRC, 2010.

[17] J. Jiang and T. Nguyen, *Linear and Generalized Linear Mixed Models and Their Applications*. Springer, 2007.

[18] S. Y. Lee and B. K. Mallick, "Bayesian hierarchical modeling: Application towards production results in the eagle ford shale of south texas," *Sankhya B*, pp. 1–43, 2021.

[19] S. Y. Lee, "Bayesian nonlinear models for repeated measurement data: An overview, implementation, and applications," *Mathematics*, vol. 10, no. 6, p. 898, 2022.

[20] International Electrotechnical Comission (IEC) 61400-12-2, *Wind Turbines—Part 12-2: Power Performance of Electricity-Producing Wind Turbines Based on Nacelle Anemometry (Edition 1.0)*, Geneva, Switzerland, 2013.

[21] N. G. Mortensen, A. Tindal, and L. Landberg, "Field validation of the rix performance indicator for flow in complex terrain," in *2008 European Wind Energy Conference and Exhibition*, Brussels, Belgium, 2008.

[22] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed. New York, NY: Springer, 2009.

[23] A. Kusiak, H. Zheng, and Z. Song, "On-line monitoring of power curves," *Renewable Energy*, vol. 34, no. 6, pp. 1487–1493, 2009.

[24] M. Lydia, A. I. Selvakumar, S. S. Kumar, and G. E. P. Kumar, "Advanced algorithms for wind turbine power curve modeling," *IEEE Transactions on Sustainable Energy*, vol. 4, no. 3, pp. 827–835, 2013.

[25] D. Villanueva and A. E. Feijóo, "Reformulation of parameters of the logistic function applied to power curves of wind turbines," *Electric Power Systems Research*, vol. 137, pp. 51–58, 2016.

[26] D. Villanueva and A. Feijóo, "Comparison of logistic functions for modeling wind turbine power curves," *Electric Power Systems Research*, vol. 155, pp. 281–288, 2018.

[27] A. A. Ezzat, M. Jun, and Y. Ding, "Spatio-temporal short-term wind forecast: A calibrated regime-switching method," *The Annals of Applied Statistics*, vol. 13, no. 3, pp. 1484–1510, 2019.

[28] P. D. Hoff, *A First Course in Bayesian Statistical Methods*. New York, NY: Springer, 2009.

[29] I. Murray, R. Adams, and D. MacKay, "Elliptical slice sampling," in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*. JMLR Workshop and Conference Proceedings, 2010, pp. 541–548.

[30] R. M. Neal, "Slice sampling," *The Annals of Statistics*, vol. 31, no. 3, pp. 705–767, 2003.