



Short Term Wind Turbine Power Output Prediction

Sándor Kolumbán¹, Stella Kapodistria^{2(✉)}, and Nazanin Noorae³

¹ Department of Mathematics and Computer Science of the Hungarian Line, Babeş-Bolyai University, Cluj-Napoca, Romania

² Department of Mathematics and Computer Science, Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands
s.kapodistria@tue.nl

³ MSD, Eindhoven, The Netherlands

Abstract. In the wind energy industry, it is of great importance to develop models that accurately forecast the power output of a wind turbine, as such predictions are used for wind farm location assessment or power pricing and bidding, monitoring, and preventive maintenance. As a first step, and following the guidelines of the existing literature, we use the *supervisory control and data acquisition* (SCADA) data to model the *wind turbine power curve* (WTPC). We explore various parametric and non-parametric approaches for the modeling of the WTPC, such as parametric logistic functions, and non-parametric piecewise linear, polynomial, or cubic spline interpolation functions. We demonstrate that all aforementioned classes of models are rich enough (with respect to their relative complexity) to accurately model the WTPC, as their mean squared error (MSE) is close to the MSE lower bound calculated from the historical data. However, all aforementioned models, when it comes to forecasting, seem to have an intrinsic limitation, due to their inability to capture the inherent auto-correlation of the data. To avoid this conundrum, we show that adding a properly scaled ARMA modeling layer increases short-term prediction performance, while keeping the long-term prediction capability of the model. We further enhance the accuracy of our proposed model, by incorporating additional environmental factors that affect the power output, such as the ambient temperature and the wind direction.

Keywords: Wind turbine power curve modeling · parametric and non-parametric modeling techniques · probabilistic forecasting · SCADA data

The authors acknowledge the Daisy4offshore consortium for the provision of the data. The work of Sándor Kolumbán and Stella Kapodistria is supported by NWO through the Gravitation-grant NETWORKS-024.002.003. Sándor Kolumbán also received funding from the Institute for Complex Molecular Systems (ICMS) in Eindhoven. Furthermore, the research of Stella Kapodistria was partly done in the framework of the TKI-WoZ: Daisy4offshore project.

1 Introduction

Wind turbine power curves (WTPC) are used for the modeling of the power output of a single wind turbine. Such models are needed in i) Wind power pricing and bidding: Electricity is a commodity which is traded similarly to stocks and swaps, and its pricing incorporates principles from supply and demand. ii) Wind energy assessment and prediction: Wind resource assessment is the process by which wind farm developers estimate the future energy production of a wind farm. iii) Choosing a wind turbine: WTPC models aid the wind farm developers to choose the generators of their choice, which would provide optimum efficiency and improved performance. iv) Monitoring a wind turbine and for preventive maintenance: A WTPC model can serve as a very effective performance monitoring tool, as several failure modes can result in power generation outside the specifications. As soon as an imminent failure is identified, preventive maintenance (age or condition based) can be implemented, which will reduce costs and increase the availability of the asset. v) Warranty formulations: Power curve warranties are often included in contracts, to insure that the wind turbine performs according to specifications. Furthermore, service providers offer warranty and verification testing services of whether a turbine delivers its specified output and reaches the warranted power curve, while meeting respective grid code requirements. See, e.g., [13,21,27] and the references therein for the aforementioned applications. Thus it is pivotal to construct accurate WTPC models. However, this is a difficult problem, as the output power of a wind turbine varies significantly with wind speed and every wind turbine has a very unique power performance curve, [14].

In this paper, we create an accurate WTPC model based on a real dataset and suggest practical and scientific improvements on the model construction. We initially construct a *static* model (in which the regressor(s) and the regressand(s) are considered to be independent identically distributed (i.i.d.) random variables) for the WTPC and demonstrate how various parametric and non-parametric approaches are performing from both a theoretical perspective, and also with regard to the data. In particular, we explore parametric logistic functions, and the non-parametric piecewise linear interpolation technique, the polynomial interpolation technique, and the cubic spline interpolation technique. We demonstrate that all aforementioned classes of models are rich enough to accurately model the WTPC, as their mean squared error (MSE) is close to a theoretical MSE lower bound. Within each non-parametric model class, we select the best model by rewarding MSEs close to the theoretical bound, while simultaneously penalizing for overly complicated models (i.e., models with many unknown parameters), using the Bayesian information criterion (BIC), see [20]. We demonstrate that such a static model, even after incorporating information on the wind speed and the available environmental factors, such as wind angle and ambient temperature, does not fully capture all available information. To this end, we propose in this paper, to enhance the static model with a

dynamic layer (in which the regressor(s) is considered to be inter-correlated e.g., time series or stochastic processes), based on an autoregressive-moving-average (ARMA) modeling layer.

1.1 Contribution of the Paper

In this paper, based on a real dataset, we explore a hybrid approach for the wind turbine power output modeling consisting of the static model plus the dynamic layer. This approach: i) provides a very accurate modeling approach; ii) is very useful for accurate short and long-term predictions; iii) indicates that, within the cut-in wind speed (3.5 m/s) and the rated output wind speed (15 m/s), the conditional distribution of the power output is Gaussian. We consider that points i)–ii) mentioned above will contribute directly to the practice, as the accurate modeling and forecasting capabilities are of utter importance. We follow the straightforward and industrially accepted approach of first estimating the WTPC, then extending this model with additional modeling layers. Although, there are various ways of modeling the WTPC, we argue, that above a given level of modeling flexibility, the exact choice is not important. The presented approach relies on available predictions of the environmental conditions. Furthermore, point *iii*) mentioned above will greatly benefit the literature, as it is the first stepping stone towards proving that random power injections from wind energy in the electric grid can be accurately modeled using a Gaussian framework, see [15, 16]. All in all, in this paper, we provide a new dataset collected from a wind turbine, and use it to show how to accurately model and forecast power output. Since the features of the used SCADA data are common among other types of turbines, the conclusions of the paper remain valid for other turbines as well. The analysis presented in this paper is scientifically and practically relevant, and contributes substantially from both the modeling and forecasting aspect, while providing a thorough overview of sound statistical methods. All results presented in the paper are motivated scientifically (when appropriate and possible) and are supported by real data.

1.2 Paper Outline

In Sect. 2, we treat the WTPC modeling: First, in Sect. 2.1, we present a simple static WTPC modeling approach, which models the power output as a function of the wind speed, using both parametric and non-parametric approaches; parametric logistic models (Sect. 2.1), non-parametric piecewise linear models (Sect. 2.1), polynomial models (Sect. 2.1), and spline models (Sect. 2.1). In Sect. 2.2, we compare the various modeling classes and determine criteria for model selection. In Sect. 3, we enhance the static model at hand by incorporating additional factors, such as the wind angle and the ambient temperature. Analyzing the residuals of the enhanced static model, we are motivated to introduce a dynamic Gaussian layer in our model, cf. Sect. 4, which produces very accurate short-term predictions, cf. Sect. 4.3. We conclude the paper with some remarks in Sect. 5. Finally, in the appendix, Sect. A, we describe the raw data and provide all information on how the data was cleaned.




Fig. 1. WTPC of the V80-2.0MW (picture from [26])

2 Power Curve Modeling and Its Limitations

When it comes to pricing wind power, assessing the possible location for a wind turbine installation or to forecasting short-term (expected) power generation for supply purposes, the main tool proposed in the literature is the WTPC. Such a curve is used to describe the relationship between the steady wind speed and the produced power output of the turbine. The shape of the WTPC for the type of wind turbines of interest to this study is depicted in Fig. 1, while the fitted curve based on the cleaned data is depicted in Fig. 2.




Fig. 2. The power output p against wind speed w for the cleaned training dataset together with an estimated WTPC

The WTPC, in ideal (laboratory) conditions, is given by the International Standard or the manufacturer, cf. [6, 26], but such curves can change over time due to environmental changes or due to component wear. This makes it paramount to estimate the power curve for each turbine individually, so these tailor-made WTPCs may be used for power generation forecasting, decision under uncertainty, and monitoring.

The literature dealing with the topic of WTPC modeling techniques is extensive and covers many fields, see, e.g., [3, 10, 12, 13, 23] and the references therein. The majority of this work is focused on obtaining the best parametric or non-parametric model for the power curve of a turbine based on the available data. To this purpose, different approaches are compared using various criteria. However, as we point out in this paper by comparing the various model classes proposed in the literature, obtaining a parametric or non-parametric estimate of the WTPC is of limited value. Although Fig. 2 indicates an “appropriate” fitted model to the power output using the wind speed, there are apparent remaining residuals that are not explained by the fitted power curve. Investigating the statistical properties of the residuals reveals features that should be taken into account in the modeling. The remaining residuals can be explained by the fact that it might be needed to use additional regressors (besides the wind speed) to explain the power output, and can be attributed to the fact that the homoscedasticity assumption is not valid, i.e. the variance of the residuals is not constant and the residuals are highly correlated. These two issues can be partly overcome by considering machine learning and artificial intelligence approaches, see, e.g., [5, 17, 25] and the references therein. However, the drawback in these approaches is that the models are not interpretable, that they require a trove of data for their training, and that they cannot be easily transferred to model the WTPC of another wind turbine. For these reasons, we strongly believe that, contrary to the existing literature, the focus should shift to obtaining models that are interpretable, that can be easily extended to various regressors, and that can capture the heteroscedastic nature of the data. Such models should not only be ranked according to the regular modeling power, but also according to their computational complexity and their numerical robustness.

2.1 Power Curve Modeling Classes

In this section, we present and compare various model classes proposed in the literature. Firstly, we introduce some notation in Sect. 2.1 that allows us to describe in a uniform manner the models belonging to different model classes. Thereafter, we describe how to calculate the estimates for each model class. For all model classes, we assume that the value of the power curve is constant below 3.5 m/s taking the estimated power output value at 3.5 m/s. Similarly, in the wind speed range between 15 m/s and 25 m/s, the power output curve is constant taking the estimated power output value at 15 m/s. Above the cut-out speed 25 m/s, the power output is set to zero as the turbine should not operate. Thus, this part of the curve is not estimated, as in addition the cleaned dataset does not contain any observations in this range. These limitations need to be incorporated into the estimation procedure of the specific models, the details of which are presented in Sect. 2.1. We note that the chosen 15 m/s as an upper bound, for the rated speed, is lower than the value suggested by the manufacturer (see Sect. A), but the data validate our choice.

We consider both parametric and non-parametric models and compare the various model classes using the mean squared error (MSE) value, while within a class (for the non-parametric models) we select a model taking into account the complexity associated with it; non-parametric model classes (e.g. polynomial or spline models) have a nested structure, where the nesting levels correspond to complexity levels within the class (e.g. degree of polynomial or knot points of splines). The selection procedure of a model within a class is presented in Sect. 2.1.

Modeling and Least Squares Estimation. A power curve is a functional relation between the wind speed, w , and the power generation, p . We define this functional relation as

$$p = \mathcal{M}_\theta(w), \quad (1)$$

where $\mathcal{M}_\theta(\cdot)$ is a function belonging to the model class parametrized by a vector of parameters $\theta \in \mathbb{R}^{n_\theta}$, with n_θ the dimension of the parameter vector depending on the model class \mathcal{M} . In the next sections, we consider various parametric and non-parametric model classes: logistic models $\mathcal{G}_\theta(\cdot)$ in Sect. 2.1, piecewise linear models $\mathcal{L}_\theta(\cdot)$ in Sect. 2.1, polynomial models $\mathcal{P}_\theta(\cdot)$ in Sect. 2.1, and spline models $\mathcal{S}_\theta(\cdot)$ in Sect. 2.1.

Given a dataset containing a series of wind speed and power output pairs, $(w_k, p_k)_{1 \leq k \leq N}$ with N the total number of observations, we define the least squares estimate within a model class as

$$\hat{\theta} = \arg \min_{\theta} \frac{1}{N} \sum_{k=1}^N (p_k - \mathcal{M}_\theta(w_k))^2. \quad (2)$$

In order to shorten notation, the power output given by the least-squares estimated model at a given time is going to be denoted as $\hat{p} = \mathcal{M}_{\hat{\theta}}(w)$, where the model class \mathcal{M} is always going to be clear from the context.

One important conclusion of the paper is that WTPC modeling has significant limitations. In order to show this, we introduce some elementary facts about least squares estimates related to quantized data (as the SCADA data are quantized to one decimal digit, as described in Sect. A).

Proposition 1 (Lower bound for MSE). *Irrespective of the model structure that is used to fit a model to the training data, if the training data are quantized in the regressor then there is a minimal attainable MSE and that can be calculated based on the data.*

Let the samples be $(x_k, y_k)_{1 \leq k \leq N}$ and consider a model $y = \mathcal{M}_\theta(x)$, with least squares estimate

$$\hat{\theta} = \arg \min_{\theta} \frac{1}{N} \sum_{k=1}^N (y_k - \mathcal{M}_\theta(x_k))^2.$$

Let \mathcal{X} be the set of all appearing values of x , i.e. $\mathcal{X} = \bigcup_{k=1}^N \{x_k\}$, then the minimal attainable MSE value can be calculated as

$$\min_{\theta} \frac{1}{N} \sum_{k=1}^N (y_k - \mathcal{M}_{\theta}(x_k))^2 \geq \frac{1}{N} \sum_{x \in \mathcal{X}} \sum_{k=1}^N \mathbb{1}_{\{x_k=x\}} (y_k - \bar{y}_x)^2, \quad (3)$$

with

$$\bar{y}_x = \frac{\sum_{k=1}^N \mathbb{1}_{\{x_k=x\}} y_k}{\sum_{k=1}^N \mathbb{1}_{\{x_k=x\}}}, \quad (4)$$

and $\mathbb{1}_{\{\cdot\}}$ an indicator function taking value 1, if the event in the brackets is satisfied, and 0, otherwise.

Proof. The MSE can be written as

$$\frac{1}{N} \sum_{k=1}^N (y_k - \mathcal{M}_{\theta}(x_k))^2 = \frac{1}{N} \sum_{x \in \mathcal{X}} \sum_{k=1}^N \mathbb{1}_{\{x_k=x\}} (y_k - \mathcal{M}_{\theta}(x))^2.$$

The right hand side of the above equation can be bounded by calculating lower bounds to each group of summands involving the same regressor value x . Given x , let $z_x = \mathcal{M}_{\theta}(x)$, then the corresponding group of summands can be written as

$$S_x := \frac{1}{N} \sum_{k=1}^N \mathbb{1}_{\{x_k=x\}} (y_k - z_x)^2.$$

The derivative of S_x with respect to z_x is

$$\begin{aligned} \frac{\partial}{\partial z_x} S_x &= -\frac{2}{N} \sum_{k=1}^N \mathbb{1}_{\{x_k=x\}} (y_k - z_x) \\ &= 2\frac{z_x}{N} \sum_{k=1}^N \mathbb{1}_{\{x_k=x\}} - \frac{2}{N} \sum_{k=1}^N \mathbb{1}_{\{x_k=x\}} y_k. \end{aligned}$$

Solving the optimality condition $\frac{\partial}{\partial z_x} S_x = 0$ for z_x reveals that the minimum is obtained at (4). Thus, the lower bound is attained if $\forall x \in \mathcal{X} : \mathcal{M}_{\theta}(x) = \bar{y}_x$.

The lower bound given in (3) is always true, but it is not necessarily a tight bound. If every regressor's value, x , appears only once in the data, then this bound would be 0, which is trivial for a sum of squares. The bound will give a nonzero value in the case of observations with $|\mathcal{X}| < N$, where $|\mathcal{X}|$ denotes the cardinality of the set \mathcal{X} . In our case, when considering $\mathcal{X} = \{3.5, 3.6, \dots, 14.9, 15\}$, with $|\mathcal{X}| = 116$ and $N = 12849$, the bound is non-zero.

Constrained Model. In order to keep the notation and the calculations simple, without loss of generality, we estimate the corresponding parameters and choose the best WTPC model (within a class), only for wind speeds in the range [3.5, 15].

To this end, we consider a slightly modified model: For a model \mathcal{M}_θ determined by the parameter vector θ , from the model class \mathcal{M} , we define the constrained model $\overline{\mathcal{M}}_\theta$ as

$$\begin{aligned} \overline{\mathcal{M}}_\theta(w) = & \mathcal{M}_\theta(3.5 \cdot \mathbf{1}\{w < 3.5\} + w \cdot \mathbf{1}\{3.5 \leq w < 15\} \\ & + 15 \cdot \mathbf{1}\{15 \leq w < 25\}) - \mathcal{M}_\theta(0) \cdot \mathbf{1}\{25 \leq w\}. \end{aligned} \quad (5)$$

The argument of \mathcal{M}_θ is constructed such that for wind values smaller than 3.5 the model $\overline{\mathcal{M}}_\theta$ will result in the same power output values as $\mathcal{M}_\theta(3.5)$, for values $w \in [3.5, 15)$ the model $\overline{\mathcal{M}}_\theta$ results in the same power output as \mathcal{M}_θ , for values $w \in [15, 25)$ the model $\overline{\mathcal{M}}_\theta$ results in the same power output as $\mathcal{M}_\theta(15)$, and for wind values above 25 the predicted power output is zero. Considering the constrained model, we can estimate the parameters of the model as usual after a slight transformation of the training data: for all observations with $w < 3.5$ the value of w is changed to 3.5, for all observations with $w \in [15, 25)$ the value of w is changed to 15, and all observations with $w \geq 25$ are ignored. Then, this transformed dataset is used for the parameter estimation.

Logistic Models. Logistic models have been widely used in growth curve analysis and their shape resembles that of a WTPC under the cut-out speed. For this reason, they were recently applied to model WTPCs, see [8, 13]. [12] present an overview of parametric and non-parametric models for the modeling of the WTPC, and state that the 5-parameter logistic (5-PL) function is superior in comparison to the other models under consideration. However, as we show in the sequel, this statement should be viewed with skepticism and perhaps should be interpreted as the result of a comparison only within models with the same number of parameters (parametric models) or same level of complexity (non-parametric models).

In this section, we apply a different formulation of the logistic model from the one used in [12], so as to improve fitness. The 5-PL model used in [13] is given as follows

$$p = \theta_5 + \frac{\theta_1 - \theta_5}{\left(1 + \left(\frac{w}{\theta_2}\right)^{\theta_3}\right)^{\theta_4}}. \quad (6)$$

In this model, parameters θ_1 and θ_5 are the asymptotic minimum and maximum, respectively, parameter θ_2 is the inflection point, parameter θ_3 is the slope and θ_4 governs the non-symmetrical part of the curve. However, this type of 5-PL does not describe the asymmetry as a function of the curvature, see [18]. As an alternative, [24] proposed a technique which can handle the curvature in the extreme regions. We apply this technique with a slight modification to fit a logistic model to the WTPC. The general form of the model is presented in Eq. (7)

$$\begin{aligned}
p &= \mathcal{G}_\theta(w) \\
&= \theta_1 + \frac{\theta_4 - \theta_1}{1 + \exp\left[-\left\{\theta_2(w - \theta_3) + \theta_\ell(w - \theta_3)^2 \mathbb{1}_{[3.5, \theta_3)}(w) + \theta_u(w - \theta_3)^2 \mathbb{1}_{[\theta_3, 15]}(w)\right\}\right]},
\end{aligned} \tag{7}$$

with $\mathbb{1}_A(w)$ denoting the indicator function taking value 1 when $w \in A$, for a given set A , and zero otherwise. We substitute the term $(w - \theta_3)^2 \mathbb{1}_{[3.5, \theta_3)}(w)$ in (7) with $(w - \theta_3)^4 \mathbb{1}_{[3.5, \theta_3)}(w)$, so as to capture more accurately the curvature in the left tail of the WTPC, cf. Fig. 2. For this reason, we refer to this model as the modified Stukel model (mStukel).

In Table 1, we present the MSE and BIC for the 5-PL model and the mStukel model according to (5). As it is evident from the results presented in Table 1, the mStukel model drastically improves the fitness of the WTCP. The parameter estimates, $\hat{\theta}^{(g)}$, of the mStukel model with the corresponding standard errors are given in Table 2.

Table 1. MSE and BIC for the fitted models on the training and validation datasets

Models	Training dataset		Validation dataset
	MSE	BIC	MSE
5-PL	1554.2700	131000	1650.7300
mStukel	884.4321	123710	1020.3800

Table 2. Parameter estimates (StandardError) for the mStukel logistic model on the training dataset

Parameter	Training set
$\hat{\theta}_1$	-30.8580 (1.3187)
$\hat{\theta}_2$	0.5845 (0.0010)
$\hat{\theta}_3$	9.6481 (0.0032)
$\hat{\theta}_4$	2010.46 (1.2119)
$\hat{\theta}_u$	0.1602 (0.0019)
$\hat{\theta}_\ell$	-0.0010 (0.00004)

Piecewise Linear Model Class. Piecewise linear models are not particularly appealing for practical use for many reasons, but they are very useful as benchmarks. We include piecewise linear models so they can serve as a benchmark non-parametric model class and because, as it is shown in the sequel, cf. Proposition 2, this class can achieve the bound of the MSE.

Let the piecewise linear function be defined as follows

$$p = \mathcal{L}_{\boldsymbol{\theta}^{(\ell)}}(w) = \theta + \sum_{k=0}^{m-1} \mathbb{1}\{s_k \leq w\}(w - s_k)\theta_k. \quad (8)$$

The parameter vector $\boldsymbol{\theta}^{(\ell)}$ consists of the (height) parameter θ and the segment slope parameters θ_k , $k = 0, 1, \dots, m-1$. The splitting points s_0, \dots, s_{m-1} should be defined beforehand. Throughout the paper, we use equidistant splitting points on the interval $[3.5, 15]$ and we estimate the parameters of the constrained model $\tilde{\mathcal{L}}_{\boldsymbol{\theta}^{(\ell)}}$ defined in Sect. 2.1.

The piecewise linear model can achieve the bound of the MSE on the training data. This is due to the quantized nature of the values of the data to one decimal digit. Thus, using 116 splitting points for the piecewise linear model, we can cover the entire range of wind values in $[3.5, 15]$. In this case, the least-squares estimate of the power output is given as the average of the power values of samples given the value of the wind speed, thus attaining the lower bound of the MSE on the training data.

Proposition 2 (Piecewise linear model attaining the MSE lower bound). *For a scalar dataset with one dimensional regressors with $|\mathcal{X}| = m+1$ a piecewise linear model of order m with split points \mathcal{X} attains the minimal MSE bound given in Proposition 1.*

Proof. For $|\mathcal{X}| = 1$ the only parameter to be estimated is θ , which should be chosen as \bar{y} , cf. Proposition 1.

The rest of the proof is based on induction on the cardinality of the set \mathcal{X} , denoted by $|\mathcal{X}|$. Let $(x^{(i)})_{0 \leq i \leq m}$ be the ordered values of \mathcal{X} , such that $x^{(0)} < x^{(1)} < \dots < x^{(m)}$ and lets assume that the parameters $\theta, \theta_0, \dots, \theta_{m-2}$ are chosen such that the linear model with these parameters attains the minimal MSE on the restricted dataset having regressors $\mathcal{X} \setminus \{x^{(m)}\}$. To prove the statement, we need to show that θ_{m-1} can be chosen such that $\mathcal{L}_{\boldsymbol{\theta}}(x^{(m)}) = \bar{y}_{x^{(m)}}$. From the definition of the piecewise linear function

$$\mathcal{L}_{\boldsymbol{\theta}^{(\ell)}}(x^{(m)}) = \theta + \sum_{k=0}^{m-2} (x^{(m)} - x^{(k)})\theta_k + (x^{(m)} - x^{(m-1)})\theta_{m-1}.$$

Solving this equation for θ_{m-1} , we get that

$$\theta_{m-1} = \frac{\bar{y}_{x^{(m)}} - \theta - \sum_{k=0}^{m-2} (x^{(m)} - x^{(k)})\theta_k}{x^{(m)} - x^{(m-1)}},$$

which concludes the proof.

It can be stated in general that once a model structure has enough degrees of freedom to assign the estimates $\mathcal{M}_{\boldsymbol{\theta}}(w)$ independently to every wind value $w \in \mathcal{X}$, then the lower bound for the MSE value can be attained.

Piecewise linear model classes, with a fixed number of splitting points equidistantly chosen in an interval, are not properly nested, if $m = 1, 2, \dots$. Proper nesting is achieved, if $m = 2^0, 2^1, 2^2, \dots$, or if some other exponential series is chosen. The parameters $\hat{\theta}_m^{(\ell)}$ of a model belonging to the fixed choice of m can be estimated for different values of m , but then the problem reduces to optimally choosing m , which is a model selection problem.

The other reason, why it is instructive to examine the properties of the piecewise linear model structure, is that, assuming Gaussian residuals, the combined variance of the estimated parameters can be calculated analytically. This is visualized in Fig. 3 as the trace of the estimated covariance matrix of the parameters is shown against the complexity of the model class m . This shows the generic features of model selection problems.

For very small values of m the modeling error is big, so the estimated variance of those few estimated parameters is going to be big (combination of modeling error and variance from noise), so the sum is going to be a sum of few but large in absolute value entries. Values of m that correspond to a model class that can properly model the data will result in a sum that contains more summand terms, but with smaller in absolute value entries. The variance of the parameters in this case is expected to be small for two reasons: *i*) the modeling error is reduced or eliminated; *ii*) a small number of parameters needs to be estimated from the data. For higher values of m the number of summands will increase and so will the corresponding absolute values of the entries. This is because the modeling error was already minimized and a higher number of parameters needs to be estimated from the data, which increases their variance. This heuristic results in a quasiconvex shape of the MSE as a function of the complexity parameter m (a.k.a. the model order). The goal of model selection is to define how an optimal model order \hat{m} should be chosen. This question arises in the case of all non-parametric model classes and our approach is based on the BIC, see Sect. 2.1.

As it can be seen in Fig. 3, the trace of the estimated covariance matrix of the parameters is quasiconvex shaped, as it decreases in the beginning and then it increases rapidly when the model order is increased. If this is compared to the decrease of the MSE shown in Fig. 4, we see that going above a given complexity level just adds unnecessary uncertainty to the estimation without improving the modeling precision. This trade-off should be balanced by the model selection algorithm. Using model selection based on the BIC, see Sect. 2.1, the optimal number of linear segments turns out to be $m = 13$. The parameters of the estimated model are given in Table 3. The MSE of this model on the training data is 815.5127, while on the validation set it is 974.6084.

Polynomial Model Class. A univariate polynomial model of degree m of the power function is given as

$$p = \mathcal{P}_{\theta}(w) = \sum_{i=0}^m \theta_i w^i. \quad (9)$$




Fig. 3. The sum of the parameter variances for the piecewise linear model assuming Gaussian residuals

Table 3. Estimated parameters of a piecewise linear model with $m = 13$ segments

$t_k = 3.5 + k \frac{15-3.5}{13}, k = 0, \dots, 12$			
Parameter	Training set	Parameter	Training set
$\hat{\theta}$	-8.3398	$\hat{\theta}_0$	0.8929
$\hat{\theta}_1$	95.4169	$\hat{\theta}_2$	17.8948
$\hat{\theta}_3$	42.1545	$\hat{\theta}_4$	58.0053
$\hat{\theta}_5$	38.8957	$\hat{\theta}_6$	60.8006
$\hat{\theta}_7$	47.2667	$\hat{\theta}_8$	6.8485
$\hat{\theta}_9$	-68.7682	$\hat{\theta}_{10}$	-207.2188
$\hat{\theta}_{11}$	-94.4818	$\hat{\theta}_{12}$	4.8853

The formulation given in (9) should be adapted to take into account the constrained model $\bar{\mathcal{P}}$ defined in Sect. 2.1. However, even after the transformation to the constrained model, and the reduction of the wind range to practically [3.5, 15], we have to note that estimating the parameters $\theta_i, i = 0, \dots, m$, of the polynomial model is a numerically difficult problem. This is due to the fact that, e.g., for a polynomial model of degree $m = 14$, the coefficient matrix of the parameters includes entries corresponding to values $1, 15, 15^2, \dots, 15^{14}$. Inverting such a matrix is numerically unstable due its high condition number, cf. [2].

To overcome the numerical stability issues, one of the simplest techniques is to rescale the argument w of the polynomial, so higher powers of the argument will still remain numerically tractable. With this change, we redefine the polynomial model as

$$\mathcal{P}_{\theta^{(p)}}(w) = \bar{p} + d_p \sum_{i=0}^m \theta_i \left(\frac{w - \bar{w}}{d_w} \right)^i, \tag{10}$$

where the polynomial parameter vector $\theta^{(p)}$ contains the coefficients of the polynomial \mathcal{P} as well as the scaling parameters $\bar{w}, \bar{p}, d_w, d_p$. \bar{w} and \bar{p} denote the

sample averages of the wind speed (w), and the sample average of the power output (p), respectively, while d_w and d_p denote the sample standard deviation of the wind speed (w), and of the power output (p), respectively. The model order parameter for polynomial models is the degree of the polynomial, m .

Polynomial models are not performing well according to the literature. This is the result of a combination of factors: Firstly, they are not capable of capturing the flat plateau on the left and the right tail of WTPC. Once this obvious drawback is compensated by considering the constrained model, polynomial models drastically increase their fitness. Secondly, there are various numerical difficulties associated with the estimation of the parameters of polynomial models. Unfortunately, this issue constitutes a significant drawback especially at higher model orders, as we show in Sect. 2.2.

Estimating (in the least squares sense) the coefficients of a polynomial with degree $m = 14$ results in the parameters presented in Table 4. The choice of degree $m = 14$ is explained in Sect. 2.1. The MSE of this model on the training data is 812.2287, while on the validation set it is 969.8870.

Note, that the polynomial coefficients in Table 4 are reported with 15 decimal digits, as we take into account the support of the wind values $[3.5, 15]$ and the maximum degree of the polynomial model: this is due to the fact that for example for the polynomial of degree $m = 14$ the leading coefficient of the polynomial, $\hat{\theta}_{14}$, is multiplied with $\left(\frac{w-7.241154953692900}{3.092009009051451}\right)^{14}$ for $w \in [3.5, 15]$. This illustrates that the estimation of the polynomial coefficients is numerically sensitive, which is not the case for the other discussed non-parametric model classes.

Spline Model Class. Splines provide a universal family for approximating smooth functions. A spline is defined by a series of knot points and by polynomials representing its value between the knot points in a continuous way [19]. Formally a cubic B-spline is given by a triplet of parameters $\theta^{(s)} = (m, \mathbf{k}, \alpha)$,

Table 4. Estimated parameters of a polynomial model with degree $m = 14$

Par.	Training set	Par.	Training set
\hat{p}	1012.7	\hat{w}	7.241154953692900
\hat{d}_p	601.0210490157367	\hat{d}_w	3.092009009051451
$\hat{\theta}_0$	-1.083983804472287	$\hat{\theta}_8$	0.913785265115761
$\hat{\theta}_1$	1.027493542215327	$\hat{\theta}_9$	0.158488326138962
$\hat{\theta}_2$	0.437620131289974	$\hat{\theta}_{10}$	-0.462562253267288
$\hat{\theta}_3$	-0.258311524269187	$\hat{\theta}_{11}$	0.100868720012586
$\hat{\theta}_4$	0.152839963020718	$\hat{\theta}_{12}$	0.068782606353010
$\hat{\theta}_5$	0.837258874326937	$\hat{\theta}_{13}$	-0.034900832592028
$\hat{\theta}_6$	-0.693269004241413	$\hat{\theta}_{14}$	0.004495461408312
$\hat{\theta}_7$	-0.791866808461089		

where $m \in \mathbb{N}_+$ is the number of basis functions used, $\mathbf{k} \in \mathbb{R}^{m+4}$ is a vector of knot points in nondecreasing order, $\boldsymbol{\alpha} \in \mathbb{R}^m$ is the vector of coefficients for the basis functions $B_{i,\mathbf{k},3}$ defined by the Cox-de Boor recursion [4]

$$\begin{aligned} B_{i,\mathbf{k},0}(x) &= \mathbb{1}_{[k_i, k_{i+1})}(x), \quad i = 1, \dots, m+3, \\ B_{i,\mathbf{k},d}(x) &= \frac{x - k_i}{k_{i+d} - k_i} B_{i,\mathbf{k},d-1} \\ &\quad + \frac{k_{i+d+1} - x}{k_{i+d+1} - k_{i+1}} B_{i+1,\mathbf{k},d-1}, \quad d = 1, 2, 3, \\ &\quad i = 1, \dots, m+3-d. \end{aligned} \tag{11}$$

A cubic B-spline model for the WTPC is given as

$$p = \mathcal{S}_{\theta^{(s)}}(w) = \sum_{i=1}^m \boldsymbol{\alpha}_i B_{i,\mathbf{k},3}(w). \tag{12}$$

The complexity of cubic spline models is defined by the number of basis functions m . If the knot points \mathbf{k} are fixed then the parameters $\boldsymbol{\alpha}$ can be estimated analytically in the least-squares sense, but this cannot be done simultaneously with the location of the knots [7]. We use a simple suboptimal procedure to find the estimates, which performs the estimation in two rounds. In the first round the knot points are equidistantly chosen in the [3.5, 15] interval and the parameters $\boldsymbol{\alpha}$ are estimated. In the second round, new knot points are calculated, based on the data and the first round estimates, using the MATLAB[®] routine *newknt*, which reallocates the knot points to allow a better estimation of $\boldsymbol{\alpha}$. Then $\boldsymbol{\alpha}$ is estimated for the second time.

The estimated parameters $\hat{\boldsymbol{\theta}}^{(s)}$ of a cubic spline using $m = 17$ B-splines are

$$\begin{aligned} \hat{\boldsymbol{\alpha}} &= [-8.0336698 \quad -7.2559215 \quad -23.865741 \\ &\quad 78.529492 \quad 156.55003 \quad 272.98557 \\ &\quad 452.80144 \quad 690.69908 \quad 1022.923 \quad 1400.7208 \\ &\quad 1721.1444 \quad 1921.2212 \quad 1998.4378 \quad 1992.549 \\ &\quad 2005.308 \quad 1997.8069 \quad 2000.3969] \end{aligned} \tag{13}$$

with knot points

$$\begin{aligned} \hat{\mathbf{k}} &= [3.5 \quad 3.5 \quad 3.5 \quad 3.5 \quad 4.4247 \quad 5.2668 \quad 5.9855 \\ &\quad 6.7569 \quad 7.6994 \quad 8.6481 \quad 9.7265 \\ &\quad 10.8994 \quad 11.6831 \quad 12.3575 \quad 12.9990 \\ &\quad 13.6470 \quad 14.323515 \quad 15 \quad 15 \quad 15]. \end{aligned} \tag{14}$$

An interesting feature of the resulting $\hat{\mathbf{k}}$ is that its first four entries and last four entries coincide. As it can be deduced from Eq. (11), the multiplicity of the knot points shows how smooth is the function at the specific knot point.

The two endpoints (due to their high multiplicity) indicate that the higher order derivatives are zero at the endpoints of the support [3.5, 15], so the estimated WTPC is flat at the left and right tails of the support. This is expected and desired, since the support was chosen so that the WTPC outside this support is constant (left and right tail of the WTPC).

B-splines are zero outside the range defined by the knot points, so a proper power function estimate is obtained by transforming \mathcal{S} to the constrained model $\bar{\mathcal{S}}$ defined in Sect. 2.1. The MSE of the model $\bar{\mathcal{S}}$, with the parameters given above, evaluated on the training data, is 811.9171, while on the validation set, it is 969.5854.

We note that [22] developed a much more evolved procedure for the selection of the number of the knot points, as well as for the selection of the location of the knot points, but such a complicated model choice does not improve more than 1% the modeling fit, which is insignificant if compared to the improvements achieved by incorporating the wind direction and the ambient temperature, and by the addition of the dynamic layer.

Model Selection Based on BIC. In the case of non-parametric models, the model class consists of subclasses indexed by the complexity parameter m (a.k.a. model order), i.e., piecewise linear models with an increasing number of segments, polynomial models with an increasing degree, or spline models with an increasing number of basis functions. The appropriate model order should be selected in a way that adheres to the principle of parsimony: Goodness-of-fit must be balanced against model complexity in order to avoid overfitting—that is, to avoid building models that in addition to explain the data, they also explain the independent random noise in the data at hand, and, as a result, fail in out-of-sample predictions.

There are several approaches for selecting a model, among others the AIC [1] or the BIC [20]. Although AIC can be asymptotically optimal under certain conditions, BIC penalizes the model complexity stronger. Therefore we use the BIC for the selection of the models reported in the previous sections. The BIC is defined as

$$\text{BIC}(\hat{\theta}_m) = \ln(N)n_{\hat{\theta}} - 2 \ln(\hat{L}),$$

where N is the number of data samples used to estimate $\hat{\theta}_m$, $n_{\hat{\theta}}$ is the number of estimated parameters and \hat{L} is the estimated likelihood of the observations assuming the model with estimated parameters $\hat{\theta}_m$.

Assuming a Gaussian noise model $p_k = \mathcal{M}_{\hat{\theta}}(w_k) + \varepsilon_k$, $k = 1, \dots, N$, where $\varepsilon_k \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2)$, we get that the BIC can be written as

$$\text{BIC}(\tilde{\theta}) = \ln(N)n_{\tilde{\theta}} + N \ln(2\pi\sigma^2) + \frac{\sum_{i=1}^N \varepsilon_i^2}{\sigma^2}.$$

If the parameters $\hat{\theta}_m$ of a model from the subclass with complexity m are estimated using the training data, then the MSE on the training data, say MSE_m ,

is an asymptotically unbiased estimator for the unknown variance σ^2 . Thus, evaluating the BIC on the training data yields that

$$\text{BIC}(\hat{\theta}_m) \approx \ln(N)n_{\hat{\theta}_m} + N \ln(\text{MSE}_m) + N \ln(2\pi) + 1.$$

Models with different complexity are compared using the $\text{BIC}(\hat{\theta}_m)$ and the model complexity is estimated as

$$\hat{m} = \arg \min_m \text{BIC}(\hat{\theta}_m),$$

resulting in the final estimate

$$\hat{\theta} = \hat{\theta}_{\hat{m}}.$$

Using this procedure, we obtain that for the piecewise linear models, the optimal number of segments is 13, for the cubic spline models the optimal number of basis functions is 17, while for the polynomial models the optimal degree is 14.

2.2 Comparison of Models from Different Classes

Figures 4 and 5 depict the behavior of the MSE as a function of the complexity for the different model structures on the training and the validation datasets, respectively. The goal of this section is to summarize the remarks that can be made based on these figures. Since, the mStukel logistic model has a fixed number of complexity parameters m (a.k.a. fixed model order), its MSE is depicted in Figs. 4 and 5 as a constant function in m , taking values 884.4321 and 1020.3800, on the training and validation sets, respectively.

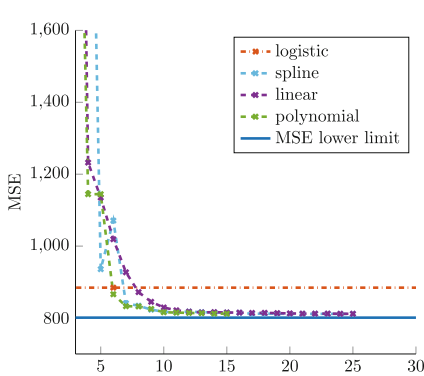


Fig. 4. The MSE on the training set

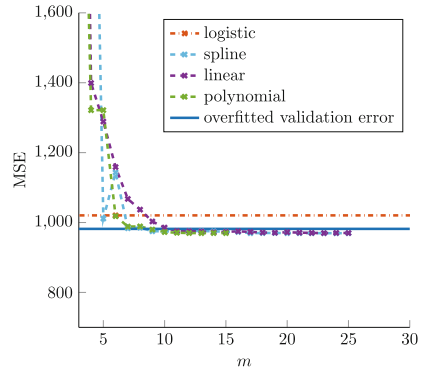


Fig. 5. The MSE on the validation set

The logistic model class, due to its parametric nature, only contains models that have a specific shape similar to what is expected from the WTPC. This

is the reason why it performs relatively well, if compared to other models with matching complexity ($m = 6$), especially when compared to the piecewise linear and the spline models. The main advantage of parametric models is that they have a pre-defined shape that matches the data, and they can describe the model with a much smaller number of parameters. As a result they can be used in case the data sparsely covers the support. However, in our case, due to the large amount of data covering densely the full support of the wind values, such advantage does not become apparent.

As stated in Proposition 1, we can calculate the MSE lower bound based on the data. Regardless of the model class, the MSE converges to that limit, as the complexity parameter tends to infinity, $m \rightarrow \infty$. Moreover, for some of the model classes we investigate, convergence will occur with finite complexity, e.g., piecewise linear models converge at $m = 116$, cf. Sect. 2.1. Similar complexity values can be calculated for the other model classes. It is important to note that the MSE converges rapidly to the lower bound for small values of m , while for large values of m , convergence seems to slow down significantly. As it can be seen in Fig. 4, this happens in the range $m \in [10, 15]$, depending on the model class.

As expected, when considering a very complicated model, then the validation error has the tendency to increase in comparison to the optimal complexity model. The solid line in Fig. 5 depicts the validation error of the model that attains the lower limit of the estimation error. As it is visible in the figure, this overfitting error is really small in comparison to the validation error of model orders around the optimal order (the validation error of the most overfitted model is 955.9658 that is approximately 1% worse than the validation errors of the different models). This shows that the data are fully covering the support $[3.5, 15]$ of the wind speed for the constrained model, and that, in our case, overfitting issues are of minor importance, as such overfitting does not impact significantly the validation error.

The model orders selected by the BIC, cf. Sect. 2.1, are all under $m = 20$, so they result in relatively simple models. They require more parameters than the logistic model, but the comparison based on the BIC indicates that the extra flexibility of these models is needed. This is not unexpected, given the provided improvement in terms of the modeling error.

Here we can underline one of the main messages of the paper: non-parametric models seem to be more suitable for the WTPC modeling than parametric models. This is mainly due to the relatively simple shape of the WTPC and the large amount of available data that can be used for the estimation of models with high complexity.

Another advantage for non-parametric models is that their shape does not depend significantly on the knowledge of the cut-in speed and the rated speed. As long as a lower bound for the cut-in speed and an upper bound for the rated speed are known, these models obtain the right shape between the two values. Whereas, for parametric models these two values are either chosen beforehand, which

impacts the model significantly, or they need to be estimated, which impacts significantly the difficulty of estimation (e.g., non-convexity, non-uniqueness, etc.).

Based on the above remark, that non-parametric models provide a better fit for the WTPC, we now turn our attention to the natural question on how to choose between the various classes of non-parametric models. This question is treated in the sequel in more detail.

In what follows, the goal is to show that in theory it should not matter which non-parametric model class we choose to estimate the WTPC, however practical considerations can still result in arguments against particular model classes. The main objective, when considering a model class, should be the numerical robustness of the estimation procedure that can provide the corresponding estimates.

The polynomial model structure is evaluated only up to degree $m = 15$, cf. Figs. 4 and 5. This is because estimating higher order polynomials is numerically infeasible, as we already mentioned in Sect. 2.1. When it comes to estimation of splines with fixed knot points \mathbf{k} , the estimates of the coefficients α can be obtained in a numerically reliable way. Similarly for piecewise linear models, given the split points (t_k) , the estimation is numerically reliable. Due to the simple shape of the WTPC, the allocation of these points is not particularly important. What could be gained by the optimal choice of these points is on the one hand negligible, as it can be seen from Fig. 4, and on the other hand it can be mitigated by adding extra parameters.

In order to illustrate that the choice of the model class is almost irrelevant, we define a measure of comparison for models from different model classes, say $\mathcal{M}_{\theta_i}(\cdot)$, $i = 1, 2$, as follows

$$\Delta_{\theta_1, \theta_2} = \frac{\mathbb{E}(\mathcal{M}_{\theta_1}(W) - \mathcal{M}_{\theta_2}(W))^2}{\min\{\mathbb{E}(P - \mathcal{M}_{\theta_1}(W))^2, \mathbb{E}(P - \mathcal{M}_{\theta_2}(W))^2\}}, \quad (15)$$

with W and P denoting the random wind speed and the random power output, respectively. $\Delta_{\theta_1, \theta_2}$ measures what is the expected difference between predictions made by the two models θ_1 and θ_2 relative to the modeling error of the better of the two. This quantity is evaluated empirically both on the training and on the validation data using the models estimated earlier and the obtained values are reported in Table 5. The difference between the logistic and the selected spline model $\Delta_{\hat{\theta}_{(g)}, \hat{\theta}_{(s)}}$ is approximately 10%, the difference between the piecewise linear and spline models $\Delta_{\hat{\theta}_{(l)}, \hat{\theta}_{(s)}}$ is under 1%, and $\Delta_{\hat{\theta}_{(p)}, \hat{\theta}_{(s)}}$ gets even smaller when it comes to the polynomial and spline models. This indicates that optimizing the model selection with regard to the class is not expected to provide significant improvements. As there is no significant difference between the various non-parametric WTPC model classes, from this point onward, we restrict our analysis to the spline model given in Sect. 2.1.

In the next sections, we address two points of concern: *i*) we discuss how to improve the WTPC model by incorporating more environmental variables, such

Table 5. The relative difference between models of different classes

Parameter	Training set	Validation set
$\Delta_{\hat{\theta}^{(g)}, \hat{\theta}^{(s)}}$	0.0884	0.0715
$\Delta_{\hat{\theta}^{(\ell)}, \hat{\theta}^{(s)}}$	0.0059	0.0048
$\Delta_{\hat{\theta}^{(p)}, \hat{\theta}^{(s)}}$	0.0005	0.0004

as the relative wind angle and the ambient temperature, and *ii*) we explore if the residuals of the model are Gaussian and investigate how to incorporate the natural autocorrelation of the data into the model by specifying that the power output variable depends linearly on its own previous values and on a stochastic term. In Sect. 3, we discuss the results of incorporating more environmental variables into the power estimation, while in Sect. 4, we explore the possibility of estimating the power output based on previous measurements in time.

3 Including More Physical Parameters

From a physical perspective the power output can be model as

$$p = \frac{1}{2} \rho \pi R^2 C_p(\lambda, \beta) w^3, \quad (16)$$

with p the power captured by the rotor of a wind turbine, ρ the air density, R the radius of the rotor determining its swept area, C_p the power coefficient which is a function of the blade-pitch angle β and the tip-speed ratio λ , and w the wind speed, see [9, Eq. (2)]. Thus, although wind speed is the most relevant factor determining the power output, it is evident from (16) that other environmental or turbine specific factors impact the power output. One way to improve the modeling and forecasting capabilities of the WTPC model is to incorporate additional relevant parameters according to the physical first principle arguments. In accordance to our available data, we illustrate the additional benefits of incorporating two additional environmental parameters: the relative incidence angle of the wind with respect to the rotor plane, say ϕ , and the ambient temperature recorded on the exterior of the wind turbine nacelle, say T .

The new signals are incorporated into (12) as follows

$$p = \mathcal{F}_{\theta^{(f)}}(w, \phi, T) = \mathcal{S}_{\theta^{(s)}}(w \cdot |\cos(\phi)|^{c_\phi}) (1 + c_T (T - \bar{T})), \quad (17)$$

with ϕ and T the relative incidence angle and the ambient temperature, \bar{T} the average temperature obtained by the training data, and $c_\phi \geq 0$ and c_T are parameters to be estimated from the data.

The inclusion of these factors can be argued based on heuristic arguments as follows: As a rough approximation, it can be stated that power generation is only achieved by the perpendicular component of the wind speed to the rotor plane of the turbine. This perpendicular component is mathematically represented by $w \cdot \cos(\phi)$. The introduction of the absolute value of the cos term ensures that the

direction of the wind is not changed. Furthermore, the inclusion of the $c_\phi \geq 0$ parameter in the $|\cos(\phi)|^{c_\phi}$ term makes sure that the wind is not amplified (i.e., the wind speed cannot get a multiplier greater than one). Regarding the inclusion of the temperature factor, this is motivated by the inherent physical relation of the temperature and the air density, as well as the prominent role of air density in the physical expression of the power output, cf. (16). Without assuming any specific functional form for this dependence, the parameter c_T can be thought of as the partial derivative of this relationship around the average temperature \bar{T} .

Using the B-spline WTPC model $\hat{\theta}^{(s)}$ with complexity $m = 17$, with estimated parameters $\hat{\mathbf{k}}$ and $\hat{\boldsymbol{\alpha}}$ given in Sect. 2.1, we can estimate the value of c_ϕ and c_T in the least squares sense. In Table 6, we provide the MSE values of the model (17), where the parameters are estimated or fixed to zero in different combinations. Fixing either c_ϕ or c_T is equivalent to omitting the corresponding modeling aspect. This allows us to see the impact of the different environmental parameters on the power generation. The first line contains the baseline, the MSE value of the WTPC model with only the wind factor. The other lines contain the MSE values corresponding to WTPC models generalized to include only the incidence angle; only the relative temperature; or both the angle and the temperature. In the remainder of the paper, $\hat{\theta}^{(f)}$ denotes the combination of the B-spline WTPC model $\hat{\theta}^{(s)}$ generalized to include the two environmental factors, with estimates $\hat{c}_\phi = 1.0115$ and $\hat{c}_T = -0.0050$. For this reason, when needed, we shall illustrate this by writing $\hat{\theta}^{(f)} = (\hat{\theta}^{(s)}, \hat{c}_\phi, \hat{c}_T)$.

Table 6. The MSE values of models given in (17) with different modeling complexity

c_ϕ parameter	c_T parameter	MSE on training set	MSE on validation set
$c_\phi = 0$	$c_T = 0$	811.9171	969.5854
$\hat{c}_\phi = 0.4279$	$c_T = 0$	810.1959	943.8046
$c_\phi = 0$	$\hat{c}_T = -0.0047$	690.5758	798.9763
$\hat{c}_\phi = 1.0115$	$\hat{c}_T = -0.0050$	681.6206	753.7712

From Table 6, it is evident that the inclusion of the incidence angle does not improve significantly the model. This is evident by the difference between the MSE values of the two models captured in the first and second line, respectively, of the table, or similarly in the third and fourth line. This difference is smaller than one percent. This is not because the incidence angle is not relevant for power generation, but because the automatism in the turbine keeps the nacelle facing the most beneficial direction with regard to the power output. Figure 6 depicts the empirical density function of the incidence angle ϕ and it shows that the incidence angle is tightly concentrated around 0° , which indicates that the wind is almost always nearly perpendicular to the rotor plane. Contrary, the inclusion of the temperature, with estimated parameter c_T , adds more than 15%

in the modeling precision. Figure 7 depicts the values of the ambient temperature for the training set (red) and the validation set (blue). Moreover, the negative sign of \hat{c}_T matches the physical insight that increasing the temperature leads to a decrease of the air density at constant pressure.




Fig. 6. Empirical density of the incidence angle in degrees




Fig. 7. The temperature variation during the training and the validation period (Color figure online)

All in all, the impact of the temperature is much larger than that of the incidence angle because the turbine mechanisms cannot influence the temperature, as they can the incidence angle. According to (16), besides the wind speed and the air density (in the form of temperature in our case, due to temperature data availability) there seem to be no other important environmental factors, that may affect the power generation significantly and that can be predicted well. In the next section we investigate the residuals of the model with the environmental factors and model the power output by adding a time series layer that allows for a significant short-term forecasting improvement.

4 Predicting Power Output Using Time Series

The effect of the modeling error can be compensated, to some extent, by modeling the residual power output with a stochastic process that has a time scale, which is much slower than that of the wind turbine mechanics. This is based on the intuition that the power output reacts quickly to environmental changes (time scale of minutes), but the environment changes in a slower rate (if there is strong wind, that will last for a few hours with a high probability).

In order to model the power output more accurately, we need to understand the statistical properties of the residuals, which are obtained as follows

$$r = p - \mathcal{F}_{\hat{\theta}_{(f)}}(w, \phi, T) = p - \hat{p}, \quad (18)$$

where \hat{p} is a brief notation for $\mathcal{F}_{\hat{\theta}_{(f)}}(w, \phi, T)$.

The standard deviation of the distribution of the residual, conditioned on the wind speed value, is shown in Fig. 8. It is apparent from the figure, that the standard deviation of the residuals is relatively small below the cut-in speed and above the rated speed, but this is not the case between these values. Thus, the WTPC model (even enhanced with the environmental factors) does not explain fully the power output leaving only the inherent randomness (captured by the residuals). Therefore, we need to further enhance the WTPC model. To this end, we concentrate in the range of wind values for which the standard deviation of the residuals seems to be large.

Let σ_w denote the conditional standard deviation of the residuals r restricted to observations with wind values equal to w

$$\sigma_w = \frac{\left(\sum_{k=1}^N \mathbb{1}_{\{w_k=w\}} (p_k - \mathcal{F}_{\hat{\theta}(f)}(w_k, \phi_k, T_k))^2\right)^{1/2}}{\left(\sum_{k=1}^N \mathbb{1}_{\{w_k=w\}}\right)^{1/2}}.$$

The rescaled residual signal r' is defined as

$$r' = \frac{r}{\sigma_w}.$$

Figure 9 depicts the conditional distribution of the rescaled residual samples in the full wind speed range, while Fig. 10 concentrates on the range [3.5, 15]. There are some remarkable features that should be emphasized.




Fig. 8. The standard deviation of the residuals $p - \hat{p}$ with respect to the wind speed




Fig. 9. Wind support [0, 25]




Fig. 10. Wind support [3.5, 15]

In Fig. 9, it is notable that the rescaled residuals outside the range defined by the cut-in speed and the rated speed have distinctive patterns. These patterns are partly caused by the fact that the data are quantized and partly caused by the fact that we are rescaling the residuals, which in this case means that we are dividing with a standard deviation close to zero, see Fig. 8 for the values of the standard deviation of the residuals.

In Fig. 10, it is notable that the rescaled residuals have very similar conditional densities for different wind speed values between the cut-in speed and the rated speed, i.e. in the range [3.5, 15] the conditional distribution of the residuals given the wind speed is (approximately) independent of the wind speed. Thus, the wind speed is (approximately) independent from the rescaled residuals. This indicates that the WTPC captures most of the wind dependence in this range except for the variance.




Fig. 11. The marginal density (blue) of the rescaled residuals conditioned on $w \in [3.5, 15]$ and the density of the standard Gaussian distribution (red) (Color figure online)

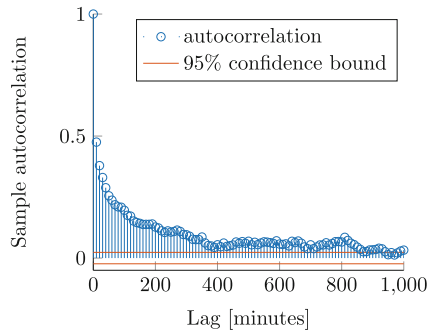


Fig. 12. The sample autocorrelation of the rescaled residual

As it can be seen in Fig. 11, the empirical density of the rescaled residuals inside the restricted wind speed interval resembles a Gaussian density. Combining this result with the fact that the residual is independent of the wind speed between the cut-in speed and the rated speed, we can reasonably assume

that the rescaled residuals can be modeled using a Gaussian stochastic process. Motivated by this result, in Sect. 4.1, we model the rescaled residuals r' using an autoregressive moving-average (ARMA) model. This model can capture the inherent autocorrelation observed in the data, cf. Fig. 12.

4.1 Dynamic Modeling

In all previous sections, for reasons of simplicity and for the clarity of the exposition, we suppressed the time index from the notation. This was also in accordance with the static models we investigated for modeling the WTPC. In this section, we further improve on the previous static models by incorporating a dynamic aspect satisfying also the Gaussian behavior of the rescaled residuals. This will additionally permit us to increase the short-term forecasting potential of our model. To this purpose, we reinstate the subscript t (indicating the time dependence) to all variables and consider for modeling purposes an ARMA(q_1, q_2) model, with q_1 autoregressive terms (with coefficients $a_i, i = 1, \dots, q_1$) and q_2 moving-average terms (with coefficients $c_i, i = 1, \dots, q_2$), i.e.,

$$r'_t = 1 + \varepsilon_t + \sum_{i=1}^{q_1} a_i r'_{t-i} + \sum_{i=1}^{q_2} c_i \varepsilon_{t-i}. \quad (19)$$

Taking into account the above model for the rescaled residuals, in the sequel, based on historical data till a given time τ , we develop the forecasting model for the power output

$$p_t = \mathcal{F}_{\theta(f)}(w_t, \phi_t, T_t) + \sigma_{w_t} r'_t, \quad t \geq \tau, \quad (20)$$

where the estimated parameters for the model $\mathcal{F}_{\theta(f)}(w_t, \phi_t, T_t)$, the conditional standard deviation σ_w , and the parameters of the ARMA model are estimated based on historical data (e.g. same season in previous year), while the estimated values of the driving noise ε_t depend on observations preceding t , but close to it in time, so this cannot be constructed based on historical data only.

We should note that the rescaled residuals seem to have a Gaussian density only inside the interval $[3.5, 15]$ of the wind speed values. Outside of this range of values, the variance of the residuals is constant and seemingly very small. Thus, similarly to Sect. 2.1, we define the constrained model with regard to the rescaled residuals as follows

$$\begin{aligned} p_t = & \mathcal{F}_{\theta(f)}(w_t, \phi_t, T_t) \\ & + \mathbb{1}\{g_\ell \leq w_t \leq g_u\} \sigma_{w_t} r'_t \\ & + (1 - \mathbb{1}\{g_\ell \leq w_t \leq g_u\}) e_t, \end{aligned} \quad (21)$$

with $3.5 \leq g_\ell < g_u \leq 15$. The values g_ℓ and g_u will be determined, so as to ensure that within the interval $[g_\ell, g_u]$ the conditional distribution of the rescaled residuals is Gaussian. Furthermore, e_t is a Gaussian noise source, such that for every finite set of indexes $\{t_1, \dots, t_N\}$ the corresponding components are

independent and identically normally distributed. Moreover, e_t is independent from the driving noise behind the rescaled residual signal, ε_t .

It needs to be mentioned that the noise ε_t cannot be obtained from the data using the model given in (21). To overcome this issue, we assume that the noise ε_t and the rescaled residual process starts at zero at the beginning of the measurement time line, i.e. ε_t and r'_t are zero, for $t < 0$, and assume that it stays “frozen” while the wind is outside the interval $[g_\ell, g_u]$. For the latter, we equivalently glue together consecutive periods of time for which the wind is within the desired range $[g_\ell, g_u]$.

The parameters of the constrained model described in Eq. (21) are: *i*) the parameters of the $\mathcal{F}_{\theta^{(f)}}(w_t, \phi_t, T_t)$ model; *ii*) the parameters of the wind speed dependent residual rescaling factors σ_w ; *iii*) the parameters of the ARMA(q_1, q_2) model, say θ_{ARMA} : $\{a_1, \dots, a_{q_1}\}$ and $\{c_1, \dots, c_{q_2}\}$. So the full parameter vector θ of the model consists of $\theta^{(f)}$, σ_w , $\{a_1, \dots, a_{q_1}\}$ and $\{c_1, \dots, c_{q_2}\}$.

As it is visible in Fig. 9, the conditional distribution of the rescaled residuals conditioned on the wind speed is not Gaussian on the full wind speed support. In order to determine the lower and upper limits of the Gaussian range, g_ℓ and g_u , we performed the Anderson-Darling test to find the p -values for the conditional distributions that show how likely it is that the rescaled residuals, given the wind speed values, are samples from a standard normal distribution. Figure 13 shows the p -values of the test, along with the wind speed value boundaries that we used for later calculations. In particular, with relatively high confidence (p -value > 0.05) we cannot reject the hypothesis that the samples come from a standard normal distribution for $g_\ell = 5.4$ and $g_u = 13.6$, while outside these bounds the hypothesis can be rejected with extremely high confidence (p -value ≈ 0).




Fig. 13. The p -value of the Anderson-Darling test for standard normality performed on the conditional distribution of the rescaled residuals

Having decided on the values g_ℓ and g_u , we describe below the procedure for estimating the parameters θ of the model described in Eq. (21). First, we estimate the parameters of the $\mathcal{F}_{\theta^{(f)}}(w_t, \phi_t, T_t)$ model, next we estimate the parameters σ_w of the rescaled residuals, last we estimate the parameters of the

Table 7. ARMA model parameters belonging to different orders

Model order	model parameters				
$q_1 = 0, q_2 = 5$	$c_1 = 0.4188$	$c_2 = 0.2941$	$c_3 = 0.2379$	$c_4 = 0.1738$	$c_5 = 0.1085$
$q_1 = 5, q_2 = 0$	$a_1 = 0.4125$	$a_2 = 0.1271$	$a_3 = 0.0824$	$a_4 = 0.0375$	$a_5 = 0.0248$
$q_1 = 5, q_2 = 5$	$a_1 = 1.3982$	$a_2 = -0.3649$	$a_3 = -0.1556$	$a_4 = 0.3187$	$a_5 = -0.2043$
	$c_1 = -0.9894$	$c_2 = 0.0847$	$c_3 = 0.1463$	$c_4 = -0.2927$	$c_5 = 0.0901$

ARMA model. This is indicated by the data, as the distribution of the rescaled residuals r'_t is symmetric around zero and it is independent of the wind w_t . All in all, this procedure is mathematically described as

$$\hat{\theta} = \arg_{\theta} \left(\min_{\theta_{\text{ARMA}}} \min_{\sigma_w} \min_{\theta^{(f)}} \frac{1}{N} \sum_{k=1}^N (p_k - \hat{p}_k)^2 \right), \quad (22)$$

where the prediction \hat{p}_k is obtained in accordance to Eq. (21) and N is the total number of observations. Note that as $N \rightarrow \infty$, $\hat{\theta}$ converges to the least square estimate obtained by optimizing for every parameter simultaneously, instead of the proposed sequential optimization. This follows from the fact that the residuals $p_k - \mathcal{F}_{\theta^{(f)}}(w_t, \phi_t, T_t)$ are symmetric around zero.

The estimated parameters $\hat{\theta}$ are calculated as follows: For $\hat{\theta}^{(f)} = (\hat{\theta}^{(s)}, \hat{c}_{\phi}, \hat{c}_T)$, we use the estimated parameters for the B-spline model $\hat{\theta}^{(s)}$ given in Sect. 2.1, and the environmental coefficients \hat{c}_{ϕ} and \hat{c}_T given in Sect. 3. For the rescaled residuals, $\hat{\sigma}_w$ is estimated in a non-parametric way for every appearing wind value in the dataset and the values are shown in Fig. 8. For the parameters of the ARMA model, we refer to Table 7. For the calculation of the parameters, we use the System Identification Toolbox [11] of Matlab.

4.2 Comparison of Model's Forecasting Capabilities

This section describes the forecasting capabilities of the models outlined above. In particular, we use the simple B-spline WTPC model, cf. Eq. (12), as a baseline to underline the improvement offered by utilizing the additional environmental regressors, cf. Eq. (17), as well as modeling the variance and the correlations remaining in the residuals of the model, cf. Eq. (21). The corresponding parameters for the models under consideration are presented in Sects. 2.1, 4.1, and 3, respectively. We depict the MSE of the three models in Figs. 14 and 15 as a function of the forecasting horizon. For the calculation of the MSE as a function of the forecasting horizon, we need to note that although the sampling frequency of the data is every $\delta = 10$ min, the forecasting horizon can receive any positive continuous value. Keeping this in mind, we define the MSE, given the forecasting horizon h , as follows

$$\text{MSE}_h = \frac{1}{N - \lceil h/\delta \rceil} \sum_{k=\lceil h/\delta \rceil}^N (p_k - \hat{p}_{k-\lceil h/\delta \rceil})^2, \quad (23)$$

where for the prediction of the k -th value, $\hat{p}_{k|k-\lceil h/\delta \rceil}$, it is required to provide as an input the wind speed values w_1, w_2, \dots, w_k , the temperature values T_1, T_2, \dots, T_k , the angle values $\phi_1, \phi_2, \dots, \phi_k$, and the power output values $p_1, \dots, p_{k-\lceil h/\delta \rceil}$, i.e., we predict from the $k - \lceil h/\delta \rceil$ power output values the future, given perfect future information of the explanatory values.

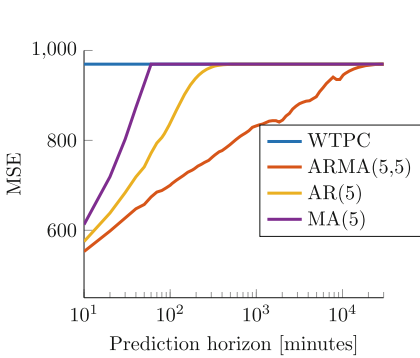


Fig. 14. Simple WTPC

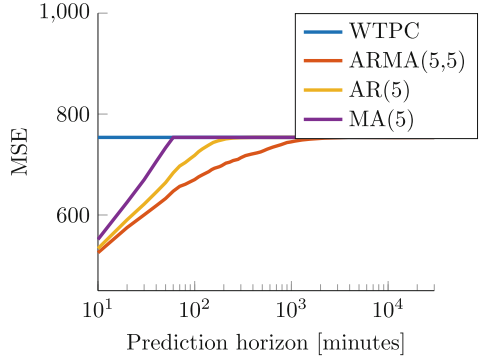


Fig. 15. WTPC enhanced with temperature and angle

As shown in Figs. 14–15, the two static WTPC models, cf. Eqs. (12) and (17), given the wind speed, angle and temperature, have a constant MSE regardless of the forecasting horizon. Contrary, the dynamic ARMA(q_1, q_2) models permit a significant reduction of the MSE, especially for short-term forecasting in the range of 1 to 50, 1 to 10^2 , or 1 to 10^3 min, depending on the values of the (q_1, q_2) parameters. Naturally, as the prediction horizon increases the added benefit of knowing the power output values from the past is getting less and less valuable. Furthermore, the least effective is the moving-average MA(q_2) model structure, since it utilizes information only from the past q_2 samples. So if the prediction horizon reaches this limit, no past information is used. As the unconditional expectation of the zero mean Gaussian process is zero, the expectation of the rescaled residuals will also be zero. This can be seen in Figs. 14 and 15, in which the MSE value of the MA(5) model becomes equal to that of the corresponding static WTPC model for prediction horizons $h > 5 \cdot \delta = 50$ min. While, for the same value of the forecasting horizon, the ARMA(5, 5) model has a reduced MSE value by 17% compared to the corresponding static WTPC model. In contrast to the short-term forecasting characteristics, the advantages of the dynamic models versus the static models disappear in the long-term. Thus, the dynamic WTPC model with the ARMA layer can be used for both short-term and long-term forecasting, as for short horizons it outperforms the static WTPC, while for longer horizons has the same performance as the static counterpart model. This result is evident in light of Figs. 14 and 15 as the various MSE values of the dynamic models converge to the MSE value of the corresponding static WTPC model.

We need to note that since the autocorrelation of the rescaled residuals (calculated using the enhanced WTPC model), cf. Fig. 12, is significant for a long period of time (more than 400 min), estimating higher order ARMA models would reduce the MSE value in comparison to the corresponding static model but this effect will vanish for time horizons longer than the autocorrelation length of the rescaled residuals.

Comparing Figs. 14 and 15, we note that enhancing the static model with the wind direction and the ambient temperature has two effects: the MSE drops significantly, but the added benefit of the dynamic layer vanishes faster (10^3 minutes instead of 10^4). This is due to the fact that the estimated ARMA model in the case of the simple WTPC was additionally trying to capture the autocorrelation structure of the temperature, which is persistent for lengthy lags. While, in the enhanced model, the temperature is provided as a regressor and therefore the ARMA model only needs to capture the remaining residual, whose autocorrelation vanishes for lower lag lengths in comparison to the temperature.

4.3 Forecasting Confidence

In this section, we are interested in investigating the performance of the dynamic model in terms of its forecasting ability. To this purpose, we visualize a power output trajectory in Fig. 16 and depict the difference between the prediction and the actual measurements in Fig. 17. In both figures, we define time 0 to be the starting point of the forecasting horizon and we assume that the wind speed, temperature and relative angle values are known also during the forecasting period, while the power output values are known only till time 0. For the creation of the figures, we consider both the static and the dynamic WTPC model and plot their predictions together with the corresponding prediction intervals. The prediction interval of the static WTPC model has a constant width, while the dynamic model has a varying width depending on the value of the wind speed. The 95% confidence band for predicted power production trajectory $\hat{p}_{k|0}$ was calculated based on the ARMA model and the wind dependent scaling factor.

For the static model, the fixed width interval is a result of calculating the variance of the residuals using the static WTPC model with the $p - \hat{p} = p - \mathcal{F}_{\hat{\theta}(f)}(w, \phi, T)$, cf. (17), as two times the standard deviation of the residuals can be used as an approximation for the 95% confidence region of the prediction. This calculation on our data results in a standard deviation for the residuals equal to 12.4925. However, for this to hold it should be the case that the residuals are normally distributed and independent of the wind speed, but as we have already shown this is not the case. Note that the variance of the residuals for the static model is calculated using observations covering the support of the wind speed values $[0, 25]$, thus simultaneously taking into account the part for which the static model is very accurate, $[0, g_\ell) \cup (g_u, 25]$, and the part in which it is highly inaccurate, $[g_\ell, g_u]$. As a result, the estimate for the variance of the residuals is overly conservative in $[0, g_\ell) \cup (g_u, 25]$, while it seems to underestimate the variance in $[g_\ell, g_u]$. This is clearly visible in Fig. 17 as during the first half of the

forecasting horizon the wind speed was in $[0, g_\ell) \cup (g_u, 25]$, while in the second half it is $[g_\ell, g_u]$.

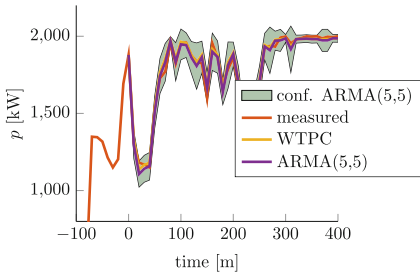


Fig. 16. Typical prediction trajectory with the corresponding 95% confidence bound




Fig. 17. The difference between the prediction and the actual measured power

For the dynamic model, cf. (21), the variance of the power estimate $\hat{p}_{k|k-\lceil h/\delta \rceil}$ can be calculated by considering the unknown random variables of $(\varepsilon_\ell)_{k-\lceil h/\delta \rceil < \ell \leq k}$, which drive the rescaled residual stochastic process $r'_{k|k-\lceil h/\delta \rceil}$, see (19). We know that the variance of $r'_{k|k-\lceil h/\delta \rceil}$ is monotonically increasing with the prediction distance $\lceil h/\delta \rceil$ and it has a bounded limit, since the ARMA model is stable (every solution $z \in \mathbb{C}$ of the $1 - \sum_{i=1}^{q_1} a_i z^{-i} = 0$ characteristic equation has absolute value less than 1). If we assume that ε_t is an i.i.d. Gaussian signal, then the variance of the power predictions can be calculated from the ARMA model and the wind dependent rescaling σ_{w_t} . The fact that the dynamic model results in a confidence band with varying width is initially surprising as the width might even shrink in size over time. The explanation for this result is that the proposed dynamic model contains a wind speed dependent scaling. This scaling factor has very small uncertainty when the wind speed is outside the interval $[g_\ell, g_u]$, i.e., for wind speed values outside the interval $[g_\ell, g_u]$, the variance is smaller in comparison to the corresponding value calculated over the full support. As it can be seen in Fig. 16, the wind value region, in which the width of the confidence band shrinks corresponds to wind speed values above the wind value g_u (the wind speed is not depicted directly but this can be inferred from the power curve and the shown measured power). In this region of wind values, predictions are more accurate and the confidence band becomes narrower than the corresponding confidence band of the static WTPC model. While, for wind values inside the interval $[g_\ell, g_u]$, it can be seen that the confidence band of the dynamic model gets wider and may even contain the confidence interval of the static WTPC model. The increased width of the confidence band is due to the combined effect of predicting values further into the future as well as the changes in the wind speed, that also effect the rescaling factor.

The confidence band of the static WTPC model contains 100% of the samples. To illustrate how imprecise the uncertainty estimate belonging to the

WTPC model is, we can calculate the maximal confidence level for which the empirical confidence is not 100%. The 41.0701% confidence band contains 99.9920% percent of the validation data, which shows a significant underestimation of uncertainty. The 95% confidence band corresponding to the dynamic model contains 96.0990% of the samples in the validation time line. This shows that the uncertainty of the predictions can be evaluated quite reliably using the dynamic model, as the empirical confidence level is relatively close to the theoretical one.

5 Concluding Remarks

The paper focuses on the short- and long-term power output forecast of a wind turbine based on past measurements of the wind speed, power output and other environmental factors, as well as perfect knowledge of the future wind speed, angle, and ambient temperature. We showed that the parametrization of the WTPC is not a key factor in improving prediction performance. The reason behind this is that for any model with a sufficiently rich structure (in the case of non-parametric models that implies sufficiently high complexity) we can achieve MSE values close to the lower bound, as long as we have sufficient data to estimate all the unknown parameters at hand. This is not a problem in our case, as we have a trove of data, making it more important to consider models that have a rich enough structure and the unknown parameters can still be estimated with high numerical accuracy, e.g. the polynomial based models suffer from numerical instability issues. Given the available data we have at our disposal, our conclusion is that the B-spline model with a sufficiently high number of knots provides a good modeling choice, as it can capture every detail of a WTPC and it can be estimated in a numerically stable way. Of course, if the data were sparse but the wind speed still covered the range of 3.5 m/s , 15 m/s , a better option would be a logistic model as such models maintain the shape of the WTPC.

The error between the actual power generation values and those predicted by the WTPC have special characteristics that open up possibilities for better modeling. Below the cut-in speed and above the rated speed of the turbine the predictions are quite accurate. However, in the middle range of the wind speed this is no longer true. We have shown that, on the dataset at hand, a proper rescaling of these residuals can transform the residual signal into a Gaussian signal. Modeling this Gaussian rescaled residual as a stochastic time series allowed us to improve significantly the predictions. The proposed model structure was able to improve up to 40% in predicting the power output of the wind turbine on short-term predictions, while the long-term prediction capabilities of the model are identical to that of the WTPC.

A Appendix: Data

The goal of this section is to describe the features of the data used in this study. The data was obtained by the supervisory control and data acquisition (SCADA)

system of a wind turbine operator in the Netherlands. The data was collected from an off-shore Vestas V80-2.0MW wind turbine, with a rated capacity of 2MW. Vestas V80-2.0MW joins the grid connection at a wind speed of 4m/s, has a rated actual power output of 2MW (typically achieved) at a wind speed of 16ms, and it is disconnected at a wind speed of 25ms. See Fig. 1 for a depiction of the theoretical WTPC. These are suggested values offered by the manufacturer, but they might change due to wear of the turbine or due to installation or geographic circumstances.

The data used for the analysis presented in this paper spans across two years and the dataset contains recordings of the environmental conditions, as well as the physical state, and power output of the turbine.

There are two important features of SCADA datasets, which are not specific to the data of our study but are common amongst SCADA datasets recorded throughout the wind industry. One of these is the 10 min reported frequency of the SCADA observations; although the signals of interest are collected at a relatively high frequency, only processed observations calculated on a 10 min window are recorded in the SCADA databases. These processed signals contain the average, maximum, minimum and standard deviation of the wind speed, and the power output amongst other quantities of interest. The second important feature is that the data are strongly quantized due to the rounding of the reported number. As a result, the observations are recorded up to one decimal digit. Some of our finding are consequences of these two properties which correspond to the quasi industry standard. Because of this, we expect that our results are also applicable to similar data coming from other wind turbine operators or wind turbine service providers.

A.1 Description of the Raw Dataset

All graphs and figures were produced using two seasonal parts of the available dataset. Throughout the paper, we refer to the data recorded between June 1, 2013, and August 31, 2013, as the training data, and the corresponding period of year 2014 as the validation data. Although, we have access to the full two year data set, we choose to restrict our analysis in a specific season of the year, as this reduces seasonality effects, while still maintaining a significant amount of data, and it permits a full decoupling between the training and the validation data. It is important to note that the results presented in the paper still hold when we perform the same analysis using the full year 2013 as training data and the data from 2014 for validation.

The dataset contains observations of various signals every 10 min. Some of the signals contained in the dataset are the ambient wind speed, say w_t , the relative direction of the wind speed with respect to the nacelle, say ϕ_t , the ambient temperature, say T_t , and the power output produced by the turbine, say p_t , at time t , $t \geq 0$. Besides the aforementioned continuous valued signals, there are some nominal variables with a discrete support, such as the variable pertaining to the different operational states of the turbine. Such variables help to identify time periods during which the turbine is out of use (maintenance,

free run, blades turned into low resistance position) or if the wind turbine is in a state different from normal operational condition.

In the first part of the paper, we suppress the subscript t as we deal with static models, while in the second part of the paper we deal with dynamic models, and we, therefore, reinstate the subscript t notation.




Fig. 18. The power output, p , against the wind speed, w , in the raw (red) and the cleaned (blue) dataset from 2013 (Color figure online)

A.2 Cleaning the data

The quality of the available SCADA data is extremely good, nevertheless it requires some pre-processing before creating the forecasting models. We list below the cleaning rules implemented in this study, according to which we disregard observations:

1. Missing entries (NAs): there are a few timestamps that are completely missing from the 10 min sampling sequence.
2. Incomplete entries (IN): if one or more of the signal values, e.g., the power output, the wind speed, etc., are missing from a data record, then the full record corresponding to this time stamp is discarded.
3. Not normal operation (NNO): based on the value of the state variables we can disregard states that do not correspond to normal operational conditions, e.g. free rotation of the wind turbine without connection to the grid, derated operation, etc.

Table 8. Summary report of the data cleaning

Year	2013	2014
Total number of observations	13248	13248
Number of NAs, IN, & NNO observations	255	445
Number of outliers	144	165
Total number of observations after cleaning	12849	12638
Percentage of observations kept after cleaning	97%	95.4%

4. Outliers: Firstly, all observations of wind power corresponding to the same wind speed are grouped together and the corresponding box plot is generated. Then, for every given wind speed value, all points with power generation outside the whiskers of the box plot (i.e., all observations falling outside the interval $(Q_1 - 3IQR, Q_3 + 3IQR)$) are discarded.

Table 8 contains the summary report of the data cleaning procedure. It shows that approximately 5% of the original data is discarded, still leaving a trove of data to be used for estimation purposes. The scatter plot of the power output, p , against the wind speed, w , is shown in Fig. 18. In this figure, we have color-characterized the training data by depicting in red the raw data, and in blue the cleaned dataset used for the analysis.

References

1. Akaike, H.: A new look at the statistical model identification. *IEEE Trans. Autom. Control* **19**(6), 716–723 (1974)
2. Belsley, D.A., Kuh, E., Welsch, R.E.: *Regression Diagnostics: Identifying Influential Data and Sources of Collinearity*, vol. 571. Wiley, Hoboken (2005)
3. Carrillo, C., Obando Montaño, A., Cidrás, J., Díaz-Dorado, E.: Review of power curve modelling for wind turbines. *Renew. Sustain. Energy Rev.* **21**, 572–581 (2013)
4. De Boor, C.: *A Practical Guide to Splines*. Applied Mathematical Sciences, vol. 27, 1st edn. Springer, New York (1978)
5. Pelletier, F., Masson, C., Tahan, A.: Wind turbine power curve modelling using artificial neural network. *Renew. Energy* **89**, 207–214 (2016)
6. IEC 61400–12-1, E.: Wind turbines - part 12-1: Power performance measurements of electricity producing wind turbines. Technical report, 88/244/FDIS (2005)
7. Kang, H., Chen, F., Li, Y., Deng, J., Yang, Z.: Knot calculation for spline fitting via sparse optimization. *Comput.-Aided Design* **58**, 179–188 (2015). *Solid and Physical Modeling 2014*
8. Kusiak, A., Zheng, H., Song, Z.: On-line monitoring of power curves. *Renew. Energy* **34**(6), 1487–1493 (2009)
9. Lee, G., Ding, Y., Genton, M.G., Xie, L.: Power curve estimation with multivariate environmental factors for inland and offshore wind farms. *J. Am. Stat. Assoc.* **110**(509), 56–67 (2015)
10. Li, S., Wunsch, D.C., O’Hair, E., Giesselmann, M.G.: Comparative analysis of regression and artificial neural network models for wind turbine power curve estimation. *J. Sol. Energy Eng.* **123**, 327–332 (2001)
11. Ljung, L.: *System Identification Toolbox User’s Guide*. The MathWorks Inc. (2010)
12. Lydia, M., Kumar, S.S., Selvakumar, A.I., Prem Kumar, G.E.: A comprehensive review on wind turbine power curve modeling techniques. *Renew. Sustain. Energy Rev.* **30**, 452–460 (2014)
13. Lydia, M., Selvakumar, A., Kumar, S., Kumar, G.: Advanced algorithms for wind turbine power curve modeling. *IEEE Trans. Sustain. Energy* **4**(3), 827–835 (2013)
14. Manwell, J.F., McGowan, J.G., Rogers, A.L.: *Wind Energy Explained: Theory, Design and Application*. Wiley, Hoboken (2010)
15. Nesti, T., Nair, J., Zwart, A.P.: Reliability of dc power grids under uncertainty: a large deviations approach. arXiv preprint [arXiv:1606.02986](https://arxiv.org/abs/1606.02986) (2016)

16. Nesti, T., Zocca, A., Zwart, A.P.: Line failure probability bounds for power grids. arXiv preprint [arXiv:1611.02338](https://arxiv.org/abs/1611.02338) (2016)
17. R. Lázaro, N.Y., Melero, J. (eds.): Wind turbine power curve modelling using Gaussian mixture copula, ANN regressive and BANN. In: The Science of Making Torque from Wind (TORQUE 2022) *Journal of Physics: Conference Series* (2022)
18. Ricketts, J., Head, G.: A five-parameter logistic equation for investigating asymmetry of curvature in baroreflex studies. *Am. J. Physiol.* **277**(2), 441–454 (1999)
19. Schumaker, L.: *Spline Functions: Basic Theory*. Cambridge Mathematical Library, 3rd edn. (2007)
20. Schwarz, G.: Estimating the dimension of a model. *Ann. Stat.* **6**(2), 461–464 (1978)
21. Shi, J., Qu, X., Zeng, S.: Short-term wind power generation forecasting: direct versus indirect ARIMA-based approaches. *Int. J. Green Energy* **8**(1), 100–112 (2011)
22. Shokrzadeh, S., Jafari Jozani, M., Bibeau, E.: Wind turbine power curve modeling using advanced parametric and nonparametric methods. *IEEE Trans. Sustain. Energy* **5**(4), 1262–1269 (2014)
23. Sohoni, V., Gupta, S.C., Nema, R.K.: A critical review on wind turbine power curve modelling techniques and their applications in wind based energy systems. *J. Energy* **2016**, 18 p., 8519785 (2016). <https://doi.org/10.1155/2016/8519785>
24. Stukel, T.: Generalized logistic models. *J. Am. Stat. Assoc.* **83**(402), 426–431 (1988)
25. Tümse, S., İlhan, A., Bilgili, M., Şahin, B.: Estimation of wind turbine output power using soft computing models. *Energy Sour. Part A: Recovery Util. Environ. Effects* **44**(2), 3757–3786 (2022)
26. Vestas: *Vestas Manual: V80-2.0 MW, V90-1.8/2.0 MW, V100-1.8 MW* (2011)
27. Widén, J., et al.: Variability assessment and forecasting of renewables: a review for solar, wind, wave and tidal resources. *Renew. Sustain. Energy Rev.* **44**, 356–375 (2015)