



This is a repository copy of *Physically meaningful uncertainty quantification in probabilistic wind turbine power curve models as a damage-sensitive feature*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/199615/>

Version: Published Version

---

**Article:**

McClean, J.H. [orcid.org/0000-0003-0594-7634](https://orcid.org/0000-0003-0594-7634), Jones, M.R., O'Connell, B.J. et al. (2 more authors) (2023) Physically meaningful uncertainty quantification in probabilistic wind turbine power curve models as a damage-sensitive feature. *Structural Health Monitoring*, 22 (6). pp. 3623-3636. ISSN 1475-9217

<https://doi.org/10.1177/14759217231155379>

---

**Reuse**

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.



[eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk)  
<https://eprints.whiterose.ac.uk/>

# Physically meaningful uncertainty quantification in probabilistic wind turbine power curve models as a damage-sensitive feature

Jacques H Mclean<sup>1</sup> , Matthew R Jones<sup>1</sup>, Brandon J O'Connell<sup>1</sup>, Eoghan Maguire<sup>2</sup>  and Tim J Rogers<sup>1</sup>

Structural Health Monitoring

1–14

© The Author(s) 2023



Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/14759217231155379

journals.sagepub.com/home/shm



## Abstract

A wind turbines' power curve is an easily accessible form of damage-sensitive data, and as such is a key part of structural health monitoring (SHM) in wind turbines. Power curve models can be constructed in a number of ways, but the authors argue that probabilistic methods carry inherent benefits in this use case, such as uncertainty quantification and allowing uncertainty propagation analysis. Many probabilistic power curve models have a key limitation in that they are not physically meaningful – they return mean and uncertainty predictions outside of what is physically possible (the maximum and minimum power outputs of the wind turbine). This paper investigates the use of two bounded Gaussian processes (GPs) in order to produce physically meaningful probabilistic power curve models. The first model investigated was a warped heteroscedastic Gaussian process, and was found to be ineffective due to specific shortcomings of the GP in relation to the warping function. The second model – an approximated GP with a Beta likelihood was highly successful and demonstrated that a working bounded probabilistic model results in better predictive uncertainty than a corresponding unbounded one without meaningful loss in predictive accuracy. Such a bounded model thus offers increased accuracy for performance monitoring and increased operator confidence in the model due to guaranteed physical plausibility.

## Keywords

Gaussian process, uncertainty quantification, Bayesian, probabilistic, power curve, wind turbine

## Introduction

As global demand for wind turbines continues to grow, so too does the need for cost-effective methods of monitoring and maintaining these turbines. Both on-shore and off-shore wind turbines are typically installed in remote and difficult-to-reach locations. The nature of these environments means that cost and hazard of on-site monitoring and maintenance is high, and as such should be minimised where possible. Using autonomously collected data such as by the Supervisory Control and Data Acquisition (SCADA) system (which come installed on most modern wind turbines) to detect faults and optimise maintenance is therefore desirable.

The power curve is commonly used in Structural Health Monitoring (SHM) as a damage-sensitive feature, as seen in Papatheou et al.<sup>1</sup> and Gonzalez et al.<sup>2</sup> The power curve describes the fundamental relationship of a wind turbine between wind speed and power

output. The power curve can be modelled from SCADA data (which collects wind speed and power output data as a matter of course). New SCADA data that deviates from this power curve is a useful indication that something is wrong with the turbine and can be carried forward into further analysis and maintenance optimisation methods. Power curve modelling has benefits outside of SHM as well, such as use in

<sup>1</sup>Dynamics Research Group, Department of Mechanical Engineering, University of Sheffield, Sheffield, UK

<sup>2</sup>Vattenfall Research and Development, Edinburgh, UK

### Corresponding author:

J.H. Mclean, Dynamics Research Group, Department of Mechanical Engineering, University of Sheffield, Mappin Building, Mappin Street, Sheffield S1 3JD, UK.

Email: jmclean1@sheffield.ac.uk

financial return predictions, power output predictions or grid management.

An example of a typical power curve is shown in Figure 1. The power curve is bounded on two sides in the power dimension. The lower bound of the power curve occurs at zero power, where the wind does not contain enough energy to move the turbine. The upper bound is the maximum power that the turbine is designed to safely extract from the wind, also known as the rated power. If wind speed increases while the turbine is operating at rated power, it will regulate itself to maintain rated power before shutting down should speeds get too high.

As seen in Figure 1, the data that compose a power curve are inherently stochastic. The authors argue that the full richness of information in this noisy data cannot be truly represented by a deterministic approach, and should be modelled probabilistically for maximum utility. A probabilistic approach (specifically a Bayesian approach) provides value because it allows for the uncertainty of the predicted power outputs to be found. This additional information is useful because it can be used to understand and model uncertainty propagation on financial returns and structural health of the turbine, allowing for robust long term risk planning, as shown in Nielson et al.<sup>3</sup> and Hughes et al.<sup>4</sup> The quantification of uncertainty also means that predictions are made with a quantified level of confidence, bolstering the decision-making processes in relation to the turbine.

It is clear then that robust probabilistic models are beneficial to the industry, and provide tangible benefits to operators. Whilst probabilistic models are highly useful, it is important that the predictive distributions returned are physically meaningful if they are to be relied on for operational decision making. It is common for probabilistic models to assume (implicitly or explicitly) that the noise is Gaussian, and can be modelled as such even at the bounded limits of the power curve. This leads to predicted probability mass outside the realm of what is physically possible for the turbine to achieve, that is, the models may suggest that there is some probability of producing more than a turbine's rated power. While the predictive density could simply be truncated this may not be statistically robust or representative of the true behaviour; it is clear that the distribution below rated power is also not Gaussian. A robust method of modelling the uncertainty such that it remains physically meaningful is therefore desirable.

This paper seeks to address this issue in probabilistic power curve modelling by investigating two bounded Gaussian Process (GP) models. The GP was chosen as a framework because it offers excellent performance and innately quantifies uncertainty for its estimates. On top of this, it is an intuitive framework, with a wealth

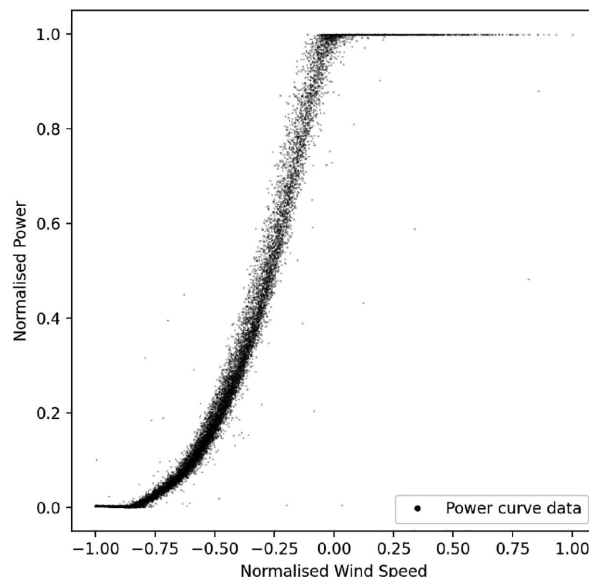


Figure 1. Example of a power curve.

of literature and associated work to draw from. As will be seen in the literature review, choice of probabilistic models is quite limited, leaving the GP as an obvious choice.

The bounded models will be assessed against each other and a standard GP, which was generated as a performance baseline for comparison. The first bounding method uses a warping function to transform the data and output of the GP in order to satisfy the boundary conditions. This is coupled with a heteroscedastic<sup>1</sup> GP, allowing the input dependent variance in the data to be captured as shown in Rogers et al.<sup>5</sup>

The second bounding method will replace the standard Gaussian likelihood in a GP with a Beta likelihood, again to satisfy boundary conditions. The Beta likelihood GP is approximated through quadrature and variational inference of a latent heteroscedastic GP. This method will be referred to as the Heteroscedastic Beta Process (HBP). The HBP method is beneficial because it models heteroscedasticity and the non-Gaussian noise whilst allowing for the use of fast inference techniques, rather than relying on significantly slower approaches to approximate a non-Gaussian posterior, such as Markov Chain Monte Carlo (MCMC).<sup>6</sup>

In this paper two novel probabilistic methodologies are introduced for modelling wind turbine power curves. When compared to existing approaches in the literature (to be discussed presently) these probabilistic methods differ in that they provide uncertainty quantification which obeys known physical *a priori* knowledge of the turbine within a consistent Bayesian

framework. The methodologies presented here are applicable to any given wind turbine, wind farm, or collection of similar turbines should there be sufficient data. It is important to note that power curves will differ between turbine models and locations, and as such care should be taken when using the methods presented to compare radically different turbines.

This paper is structured in the following manner. The section ‘Related work’ will provide a brief overview of the existing literature in wind turbine SHM and power curve modelling. The section ‘GP regression’ and its various subsections will cover the fundamental theory used to construct the bounded GP models. The section ‘Case study’ illustrates how the data was collected and used, outlines the metrics used to assess the success of the models and presents the performance results of the models. The subsections of section ‘Case study’ provide a discussion on the results of their respective models. Finally, concluding remarks and a discussion on possible avenues are made in their respective sections.

### Related work

The application of SHM in wind turbines is a well-established field, and comprehensive reviews have been published previously.<sup>7–9</sup> A wide variety of techniques exist and have been proposed for monitoring turbines, ranging from traditional Non-Destructive Evaluation (NDE) approaches to entirely data-based methods. NDE approaches are often limited to accessible parts of the turbine and require specialised sensors and equipment but can be highly effective at identifying faults in the turbine. Examples of NDE approaches for wind turbines include (but are not limited to) acoustic emissions<sup>10</sup> and ultrasonic methods<sup>11</sup> and thermal imaging.<sup>12</sup> In the remote environments wind turbines often find themselves in, the periodic nature of NDE methods might be considered less attractive than the online and continuous monitoring offered by SHM methodologies.

As mentioned in the introduction, most modern turbines come installed with a SCADA system. Using SCADA data for monitoring is a popular approach within the field, as it does not require additional equipment or expertise and allows for an online approach. This can be seen in Canizo et al.,<sup>13</sup> where a big data approach was used. Zaher et al.<sup>14</sup> apply a neural net to SCADA data to detect faults in the gearbox and generator. Yang et al.<sup>15</sup> use SCADA data to detect blade and drive train faults. Finally, power curves using SCADA data are also commonly used data to identify faults in turbines.<sup>1,2,16</sup>

Power curve modelling (whether related to SHM or not) is also an established area of research. A

comprehensive literature review of power curve models was published by Lydia et al.<sup>17</sup> Power curve models can be conceptually divided in numerous ways, but the most relevant division for this paper is that between deterministic and probabilistic models.

Historically, power curve models have been simple polynomial models. Such models are, by default, deterministic but can be highly accurate. A wide variety of polynomial power curve models exist in the literature; ranging from a basic piece-wise linear model<sup>18</sup> to more complicated and commonly used models such as a cubic polynomial.<sup>19</sup> High order polynomials have also been tested, such as the ninth order polynomial used in Raj et al.<sup>20</sup> However, higher order polynomial models can be prone to overfitting to the training data as discussed in Bishop and Nasrabadi.<sup>21</sup>

Modern approaches have steered away from classic polynomial models, and a variety of non-polynomial deterministic models have been utilised in the literature. Many of these power curve models are applications of classic machine learning algorithms; such as K Nearest Neighbour (K-NN) and random forest models,<sup>22,23</sup> support vector machines,<sup>24</sup> clustering centre fuzzy logic modelling<sup>25</sup> and Copula modelling.<sup>26</sup> A nonlinear regression model was used in Marčiukaitis et al.<sup>27</sup> This is one of the few works in the deterministic power curve literature space which explicitly obeys physical limitations, however it is still a deterministic model and only returns confidence intervals. Neural nets are another popular approach to power curve models. Models can be fairly simple power curve models<sup>28,29</sup> or can be more complicated multi-input models such as in Pelletier et al.<sup>30</sup>

Probabilistic power curve models are comparatively rare, and often these models simply make no mention if the models or their returned distributions obey physical limitations. Probabilistic models vary widely in their approaches and chosen algorithms; from a probabilistic analytical solution<sup>31</sup> to algorithms such as Monte Carlo and Fuzzy Clustering methods<sup>32</sup> to a neural net used to estimate quantiles.<sup>33</sup> Many probabilistic models assume Gaussianity in the power curve data, leading to physically impossible uncertainty densities. This can be seen in heteroscedastic models such as Rogers et al.,<sup>5</sup> where Gaussianity was assumed and, although the heteroscedastic noise models reduced the issue, predictions still led to physically impossible distributions. The assumption of Gaussian noise is also made in Jin and Tian,<sup>31</sup> where it is only used to model the region between the bounds, lending credence to the idea that the Gaussian distribution cannot effectively model power output noise at the bounds of the power curve. Several physically meaningful probabilistic models exist in literature. The use of the Weibull distribution is common here, as seen in Ge et al.<sup>34</sup> and Yun,

Hur<sup>35</sup> The Weibull distribution is useful in this case because it is bound on its lower end, and as such can be used to reflect the bounded nature of the power curve in one direction. Finally, the GP has also been used for probabilistic power curve modelling in literature. Papatheou et al.<sup>1</sup> use a GP for outlier detection in power curves for SHM purposes. Rogers et al.<sup>5</sup> use a heteroscedastic GP to model the power curve, and Pandit et al.<sup>36,37</sup> also use a GP for anomaly detection in power curves.

Power curve modelling and SHM are active fields of research. As such, this brief review has not been able to cover all work completed in these fields. Instead it has sought to highlight the relevant state-of-the-art and illustrate that the research topic of this work is an important area that has seen some work but is comparatively underdeveloped.

## GP regression

The GP<sup>38</sup> is a probability distribution over the possible functions that fit a given set of points. The GP is designed to model regression problems of the form  $y=f(x)+\varepsilon$ , where  $\varepsilon\sim\mathcal{N}(0,\sigma_n^2)$  is used to capture noise that may be present on the true function evaluations.  $\mathbf{x}$  is some vector of inputs and  $\mathbf{y}$  the target data with additive Gaussian noise  $\varepsilon$  which has a variance  $\sigma_n^2$ . The GP is mathematically defined in Equation (1), where  $\mu(\mathbf{x})$  is the mean function and  $k(\mathbf{x},\mathbf{x}')$  is the covariance function (otherwise known as the kernel).

$$f\sim\mathcal{GP}(\mu(\mathbf{x}),k(\mathbf{x},\mathbf{x}')) \quad (1)$$

GPs are a type of Bayesian inference, meaning that a defined prior is updated with observed data to produce a posterior distribution. The prior for the GP is the form of  $\mu(\mathbf{x})$  and  $k(\mathbf{x},\mathbf{x}')$ . The posterior can be computed through the use of Bayes theorem, by conditioning the test data  $\mathbf{x}_*$  on the training data  $\mathbf{x}$  through the joint Gaussian distribution (as given in Equation (2)). In Equation (2) the notation  $K_{\mathbf{x}*\mathbf{x}}$  is used to write the covariance matrix between the test variables  $f(\mathbf{x}')$  and training variables  $f(\mathbf{x})$ . Similar subscript notation is used for denoting covariance matrices between sets of training variable and test variables, for example,  $K_{\mathbf{x}\mathbf{x}}$  for the covariance of the training data with itself.

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{y}_* \end{bmatrix} \sim \mathcal{N} \left( \begin{bmatrix} m(\mathbf{x}) \\ m(\mathbf{x}_*) \end{bmatrix}, \begin{bmatrix} K_{\mathbf{x}\mathbf{x}} + \sigma_n^2 \mathbb{I} & K_{\mathbf{x}\mathbf{x}_*} \\ K_{\mathbf{x}_*\mathbf{x}} & K_{\mathbf{x}_*\mathbf{x}_*} + \sigma_n^2 \mathbb{I} \end{bmatrix} \right) \quad (2)$$

Assessing this distribution will yield the posterior predictive distribution of the GP, which shows that it is possible to predict new outputs  $\mathbf{y}_*$  subject to new inputs  $\mathbf{x}_*$  given the training inputs  $\mathbf{x}$  and their respective outputs  $\mathbf{y}$ . The posterior predictive is given in Equation

(3), where  $\mathbb{E}[\mathbf{y}_*]$  and  $\mathbb{V}[\mathbf{y}_*]$  are the mean and variance of the predicted distribution respectively.

$$\begin{aligned} p(\mathbf{y}_*|\mathbf{x}_*,\mathbf{y},\mathbf{x}) &\sim \mathcal{N}(\mathbb{E}[\mathbf{y}_*],\mathbb{V}[\mathbf{y}_*]) \\ \mathbb{E}(\mathbf{y}_*) &= m(\mathbf{x}_*) + K_{\mathbf{x}_*\mathbf{x}}(K_{\mathbf{x}\mathbf{x}} + \sigma_n^2\mathbb{I})^{-1}(\mathbf{y} - m(\mathbf{x})) \\ \mathbb{V}(\mathbf{y}_*) &= K_{\mathbf{x}_*\mathbf{x}_*} - K_{\mathbf{x}_*\mathbf{x}}(K_{\mathbf{x}\mathbf{x}} + \sigma_n^2\mathbb{I})^{-1}K_{\mathbf{x}\mathbf{x}_*} + \sigma_n^2\mathbb{I} \end{aligned} \quad (3)$$

Equation (1) shows that the GP is fully defined by its mean and covariance functions. The mean is commonly fixed to zero in literature<sup>39</sup> and in practice (as is the case in this paper). The covariance function specifies the family of distributions that may have generated the observed data. Each kernel enforces different types of behaviour on the corresponding function draws, with the ability to combine kernels allowing for a number of different types of structure to be embedded into the GP prior. A popular kernel for GPs is the squared exponential kernel, given in Equation (4), where  $\sigma^2$  is the signal variance and  $l$  is the lengthscale of the kernel; these two values are the *hyperparameters* of the kernel.

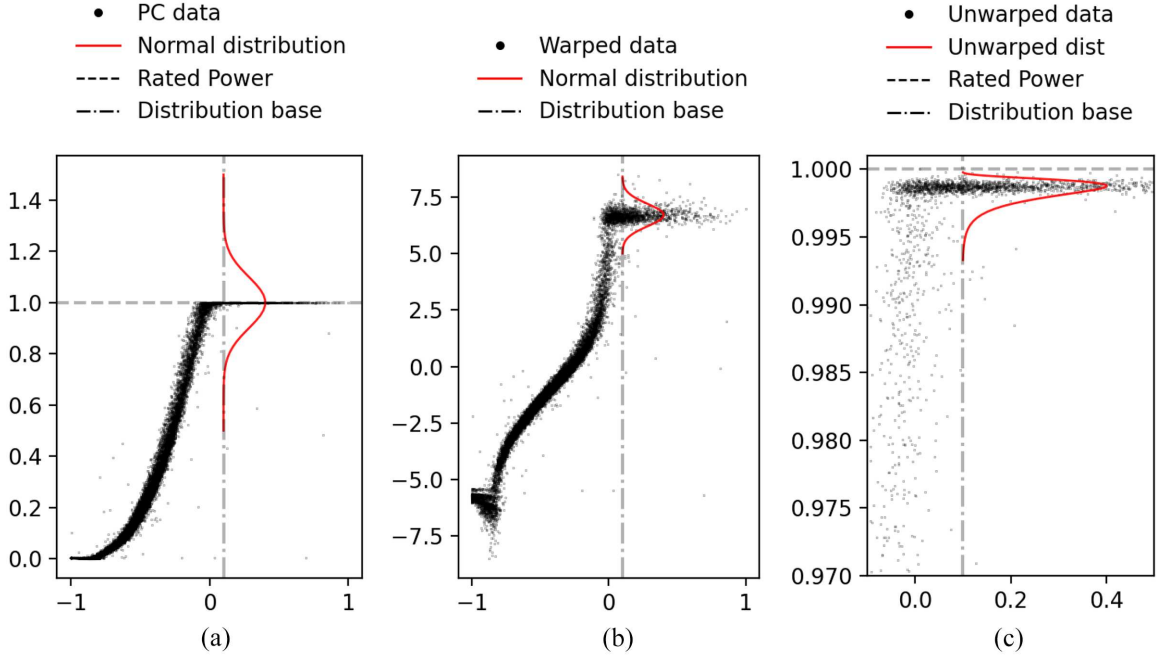
$$k(\mathbf{x},\mathbf{x}') = \sigma^2 \exp\left(-\frac{(\mathbf{x}-\mathbf{x}')^2}{2l^2}\right) \quad (4)$$

Signal variance and length scales are examples of hyperparameters; these control the shape of the family of functions defined by the kernel (and by extension the GP). They are not known *a priori*, and are learnt by minimising the negative log of the marginal likelihood of the model (Equation (5)), which solves a Type-II maximum likelihood problem such that  $\theta = \underset{\theta}{\operatorname{argmin}}\{-\log p(\mathbf{y}|\mathbf{x},\theta)\}$ .

$$\begin{aligned} \log p(\mathbf{y}|\mathbf{x},\theta) &= -\frac{1}{2}\mathbf{y}^T [K + \sigma_n^2\mathbb{I}]_y^{-1}\mathbf{y} \\ &\quad -\frac{1}{2}\log[K + \sigma_n^2\mathbb{I}] - \frac{n}{2}\log 2\pi \end{aligned} \quad (5)$$

## Heteroscedastic noise

Heteroscedastic noise is noise that varies as a function of a function's input. The power curve data in Figure 1 is an excellent example of this; it is clear that the data is much more noisy in the centre of the power curve than towards its limits – the input (the wind speed) is correlated to the amount of noise, that is, level of uncertainty. The standard (homoscedastic) GP assumes that noise and hence uncertainty is constant across the full input range, and as such cannot model the way the noise changes with wind speed. In order to model the heteroscedastic noise and return a more useful probabilistic power curve, a heteroscedastic GP model can be used, as was done in Rogers et al.<sup>5</sup>



**Figure 2.** Illustration of the warping process – the normal distributions are for illustrative purposes only. (a) Gaussian uncertainty in bounded space, (b) Gaussian uncertainty in unbounded space, and (c) unwarped distribution in bounded space.

The particular heteroscedastic model adopted in this work (distinct from Rogers et al.<sup>5</sup>) uses two latent GPs to learn the location (Equation (6)) and scale (Equation (7)) of the distributions required, first shown in Lázaro-Gredilla et al.<sup>40</sup> This is done by connecting a Multi-output kernel to a heteroscedastic likelihood. This maps the latent GPs onto a single function. The heteroscedastic GP model can be described as:

$$f_1(\mathbf{x}) \sim \mathcal{GP}(\mu(\mathbf{x}), k_1(\mathbf{x}, \mathbf{x}')) \quad (6)$$

$$f_2(\mathbf{x}) \sim \mathcal{GP}(\mu(\mathbf{x}), k_2(\mathbf{x}, \mathbf{x}')) \quad (7)$$

$$m(\mathbf{x}) = f_1(\mathbf{x}) \quad (8)$$

$$s(\mathbf{x}) = g(f_2(\mathbf{x})) \quad (9)$$

$$y(\mathbf{x}) \sim \mathcal{N}(m(\mathbf{x}), s(\mathbf{x})) \quad (10)$$

The transform  $g(\cdot)$  in Equation (9) in this case is an exponential transform, in order to keep the function positive, since it models the variance of the noise which must be greater than zero.

### Warping functions

The first bounded model investigated in this paper is the warping function. A visual illustration of how the warping function works is shown in Figure 2. Figure 2(a) is an illustrative example of how a simple normal distribution modelling uncertainty at rated power has

a large portion of probability mass outside of what is physically possible.

The warped bounding method starts with transforming the power data using the logit function,<sup>41</sup> shown in Equation (11). Conceptually, the warped GP moves the data from the bounded space to an unbounded one by means of a nonlinear transformation. The warped power curve is shown in Figure 2(b). The logit function moves the power curve from a bounded space ( $p \in (0, 1)$ ) in the power dimension to an unbounded space ( $p \in (-\infty, \infty)$ ) in the power dimension. The GP is then applied to the warped power curve, as in this unbounded space the bounds on the power curve no longer apply, so the assumption of Gaussianity holds.

$$\text{logit}(p) = \ln\left(\frac{p}{1-p}\right) \quad \text{for } p \in (0, 1) \quad (11)$$

After fitting, the data and the returned uncertainty distribution are unwarped using the inverse of the logit function, resulting in the bounded data and distribution shown in Figure 2(c). An important note on this method is that it is difficult to evaluate the unwarped uncertainty, as no closed form solution exists for it; the density shown in Figure 2(c) is presented for visualisation but is not exact. Therefore, it is simpler to work with and evaluate the model directly in the warped space. To illustrate how the uncertainty looks in the

unwarped space for the entire power curve, a confidence bound can be created in the warped space at  $\pm 2\sigma$  and unwarped with the predictive mean.

### Non-Gaussian likelihoods

A typical GP model is formulated using a Gaussian likelihood. This assumes that the data is drawn from a normal distribution, that is,  $f(x)$  has noise  $\varepsilon$  added and  $\varepsilon \sim \mathcal{N}(0, \sigma_n^2)$ . This results in a normally distributed posterior function, as shown in Figure 2(a). Without any modifications, this is not a suitable model for probabilistic power curve modelling. Instead of warping the output space, the likelihood can be changed to a non-Gaussian likelihood – one that better represents the distribution of the observed data for the problem at hand.

The likelihood chosen here is the Beta likelihood, imposing the assumption that the data is drawn from a Beta distribution. The Beta distribution is a continuous probability distribution in the interval  $[0,1]$ . It is defined in Equation (12) where  $\alpha$  and  $\beta$  are the shape parameters of the distribution. The probability density function is given in Equation (13), where  $B(\alpha, \beta) = \Gamma(\alpha)\Gamma(\beta)/\Gamma(\alpha + \beta)$  and  $\Gamma$  is the Gamma function. The use of the Beta likelihood is a bounding method because the GP with a Beta likelihood would imply that data were generated from Beta distributions, which are inherently bound between  $[0,1]$ .

$$y \sim \text{Beta}(\alpha, \beta) \quad (12)$$

$$p(y) = y^{\alpha-1} \frac{(1-y)^{\beta-1}}{B(\alpha, \beta)} \quad (13)$$

Using a non-Gaussian likelihood in the GP (like the Beta likelihood) results in an analytically intractable posterior.<sup>42</sup> The posterior must instead be approximated, numerous approximations exist for this problem. A common method (used in the popular GP packages GPflow<sup>43</sup> and GPML<sup>44</sup>) to approximate the posterior is MCMC<sup>6</sup> sampling. However, this can require a large amount of samples to approximate the posterior well, and can be very slow.

Instead of using MCMC, this paper approximates the Beta likelihood using variational inference and quadrature of a heteroscedastic GP. The model proceeds by introducing two latent GPs to calculate  $\alpha$  and  $\beta$  as functions of the input, instead of the mean and variance as used in the normal heteroscedastic GP as described previously. This is expressed mathematically below, where  $t(\cdot)$  represents an exponential transform to keep the functions over  $\alpha$  and  $\beta$  positive.

$$[f_1, f_2] = \text{GP}(m(x), k(x, x')) \quad (14)$$

$$g_1 = t(f_1) \quad (15)$$

$$g_2 = t(f_2) \quad (16)$$

$$y \sim \text{Beta}(g_1, g_2) \quad (17)$$

A flow diagram of this process is shown in Figure 3. Figure 3 illustrates how  $f_1$  and  $f_2$  are calculated from the power curve data using the heteroscedastic GP, both  $f_1$  and  $f_2$  are distributed normally for each input. These are then transformed exponentially to create  $g_1$  and  $g_2$  which are approximations of  $\alpha$  and  $\beta$ . Using total variance and mean, the distributions of  $g_1$  and  $g_2$  are used to find a final  $\alpha$  and  $\beta$  value for each wind speed input. These are subsequently used to create the final Beta distributions at each input shown in the final panel.

This approach is useful because it allows for fast posterior approximation with established methods, in this case the Stochastic Variational Inference approximation.<sup>45</sup> This is further explained in the following section.

A further benefit to this approach is that the heteroscedasticity of the model is not lost, as  $\alpha$  and  $\beta$  vary with inputs and as such capture the varying uncertainty with inputs as one would expect, alongside capturing the progression of the mean function.

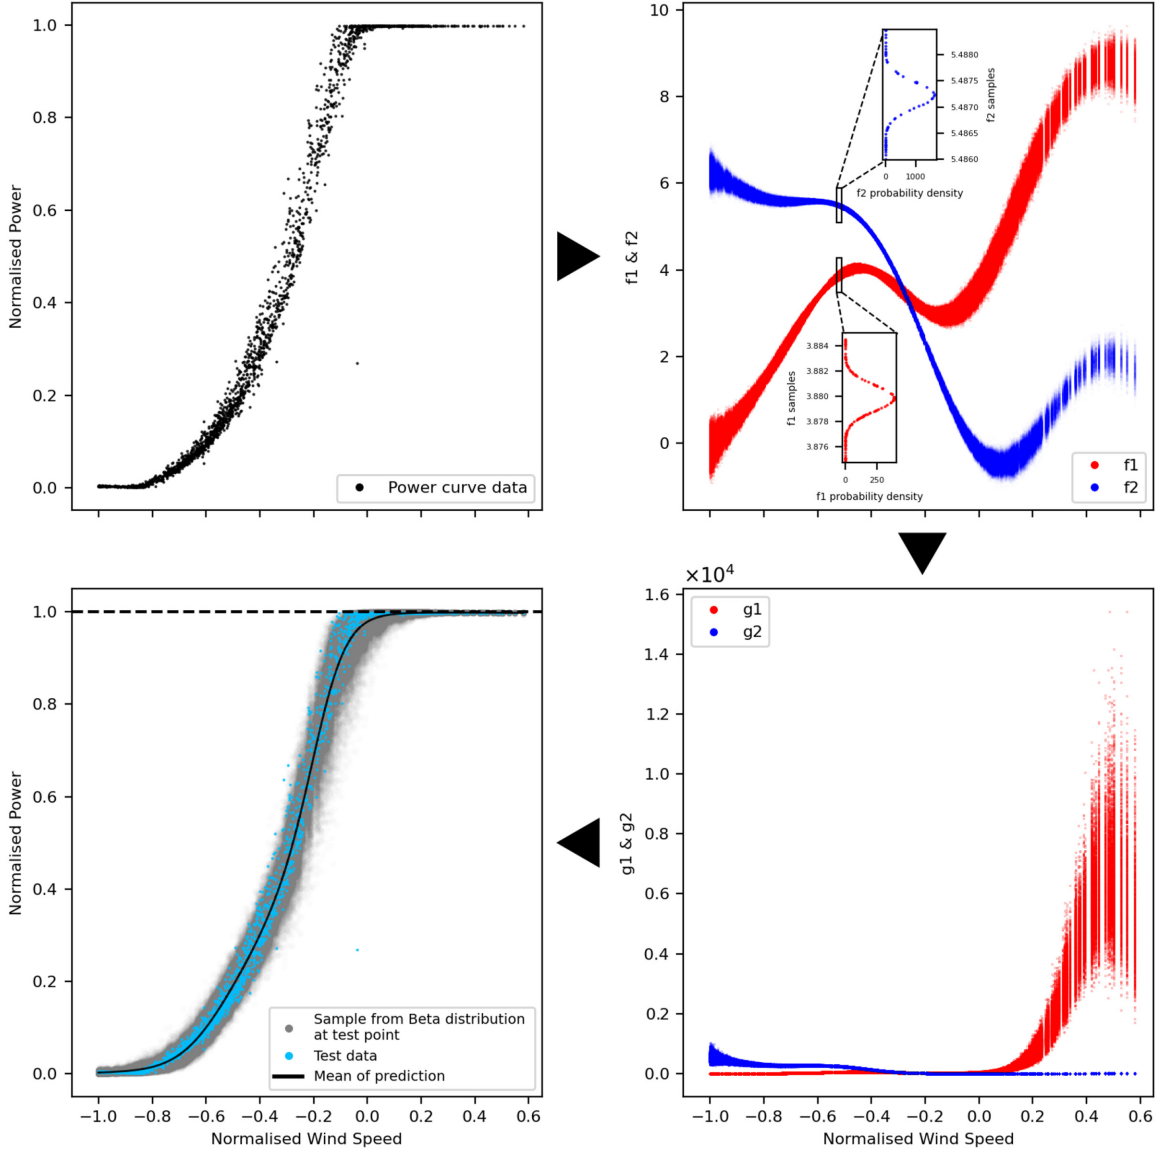
### Approximations and quadrature solutions

When considering the large heteroscedastic models used here, two problems arise. The first is that solutions for heteroscedastic and non-Gaussian models do not, in general, exist in closed form, and must be approximated to reach a solution. The second is that GPs are computationally inefficient when used on large datasets. This is due to their reliance on a matrix inversion when calculating the posterior and the marginal likelihood, which has time complexity  $\mathcal{O}(N^3)$  and therefore scales poorly with increasing numbers of training data points. The SCADA dataset being used in this paper is quite large so calculating the entire posterior is impractical and would require an excessive amount of computational resource and time.

This paper uses a stochastic variational inference approach<sup>45</sup> in order to address these problems. Similarly to other sparse variational inference methods, the method applied in this work uses a set of inducing points  $\bar{u}$  to summarise the model. The approximated posterior is given in Equation (18).

$$q(f, u) = p(f, u)\phi(u) \quad (18)$$

The approximate posterior is found by maximising the bound on the marginal likelihood. This bound is known as the Evidence Lower Bound (ELBO) and is used instead of the marginal likelihood to reduce



**Figure 3.** Flow chart of the Heteroscedastic Beta Process.

computational cost. The global stochastic variational approach<sup>45</sup> differs from other variational inference bounds in that it allows for the inducing points to be global points representing the entire GP despite being calculated locally – thus allowing faster computation.

A formulation of the ELBO of a heteroscedastic GP<sup>46</sup> is shown in Equation (19), where  $\int q(f_i)q(g_i) \log p(y_i|f_i, g_i)df_idg_i$  is the likelihood term. The terms  $q(f_i)$  and  $q(g_i)$  are the marginal variational distributions of the mean and variance latent GPs respectively, and  $\log p(y_i|f_i, g_i)$  is the conditional probability term. The model can then be optimised from this bound using stochastic gradient descent.

$$\log p(\mathbf{y}) \geq \sum_{i=1}^n \int q(f_i)q(g_i) \log p(y_i|f_i, g_i)df_idg_i - \text{KL}(q(u_f)||p(u_f)) - \text{KL}(q(u_g)||p(u_g)) \tag{19}$$

The likelihood term is clearly an  $n$  term integration over the marginal variational distributions and the conditional probability. As long as these terms are all Gaussian (such as in the heteroscedastic GP model), this integration has an analytical solution. However, the conditional term  $\log p(y_i|f_i, g_i)$  is not Gaussian in the HBP model, and as such the likelihood term does



not have an analytical solution. It is therefore approximated with quadrature, as shown in Saul et al.<sup>46</sup>

As previously mentioned, the standard GP approach scales  $\mathcal{O}(N^3)$ , that is, doubling the number of data points in the training set causes an eight-fold increase in computational load. The use of sparse variational inference allows the GP to be computed with time complexity  $\mathcal{O}(NM^2)$  instead of  $\mathcal{O}(N^3)$ , where  $M$  is the number of inducing points. Clearly, when  $M \ll N$  the reduction in burden is significant; since  $M$  is a user choice it can be used to control the level of computational load balanced against the expressive power of the GP. Since the underlying function of the power curve is relatively smooth, it can be well represented with a small  $M$ . This approach solves the computational complexity problem, as it now requires significantly less time to train the models. As mentioned earlier, the approximated posterior is found by maximising the ELBO. This means that the heteroscedastic methods can be used, as they can be approximated efficiently.

## Case study

The data used in this paper comes from the SCADA system of an operating wind turbine. The SCADA data is the 10 min average of power output and wind speed (as taken on the nacelle) comprised of 50,000 readings. The measured values of wind and power have been obscured by scaling for confidentiality reasons.<sup>2</sup>

The dataset was split into three equal parts: a training dataset, a testing dataset and a validation set. The training data is used to optimise the hyperparameters of the GP. The predictive performance of the trained models are evaluated on the test data, which they have not seen before. The validation set is put aside should further validation of predictive performance be required.

After the initial scaling to anonymity, the power output data was again normalised between 0 and 1 in order to allow for the Beta likelihood and logit warping approaches to work. Neither the original scaling or the subsequent normalisation impacted the relationship between power output and wind speed as both were linear transforms. The training data was then cleaned of severe outliers, for example known curtailments.

Two performance metrics will be used in order to compare the success of the various models in modelling the power curve. The first is Normalised Mean Squared Error (NMSE), shown in Equation (20), which is used to assess mean predictive performance (or more concretely, it measures how close to the mean of the predictive distribution the test value was). The NMSE can be thought of (and functions similarly to) percentage

**Table 1.** Case study results.

Model	NMSE	Joint log likelihood
Standard GP	<b>0.24</b>	32,312
WH GP	5.4	-12,469
HBP	0.53	<b>68,866</b>

GP: Gaussian Process; WH GP: Warped Heteroscedastic GP; HBP: Heteroscedastic Beta Process.

error. It is an intuitive method of assessing point predictions, and is important to consider when only the point predictions matter, or if the model needs to be compared to a deterministic model.

$$\text{NMSE} = \frac{100}{N\sigma^2} \sqrt{(y - \hat{y})^T (y - \hat{y})} \quad (20)$$

The second metric being used is the joint log predictive likelihood which is computed as the sum of the predictive log likelihood at every test value under the posterior. This is an especially important metric for this paper as it is a measurement of how good the uncertainty prediction of a specific model was at capturing the test data and not solely the performance of the mean estimate. The larger the joint log likelihood, the better the model has captured the true distribution of the test data and the more successful the model. Interpretation of the absolute values of the log likelihood is difficult; however, comparatively models which better capture the distribution of the data will have higher log likelihoods.

To facilitate the discussion of results, the NMSE and joint log likelihood of each model is shown in Table 1, where WH GP is the Warped Heteroscedastic GP. The results of each model will be further discussed in their respective following subsections. A small figure illustrating how the mean predictions differ from each other can be seen in Figure 4.

All models use the same kernel configuration; a Matérn 3/2 kernel combined with a Linear kernel to capture the gross trend of the data. The choice of a Matérn kernel over other kernels is supported in the literature by Stein.<sup>47</sup> However, the power curve is a strong and simple signal, and any generic suitably flexible kernel (e.g. squared exponential, Matern or rational quadratic) will model the problem well and comparably.

It will be argued by the authors in the remainder of this section that the HBP best captures the physical behaviour and uncertainty associated with the power curve of an operating wind turbine. All the models presented in this paper were written in Python and run using GPflow.<sup>43</sup>

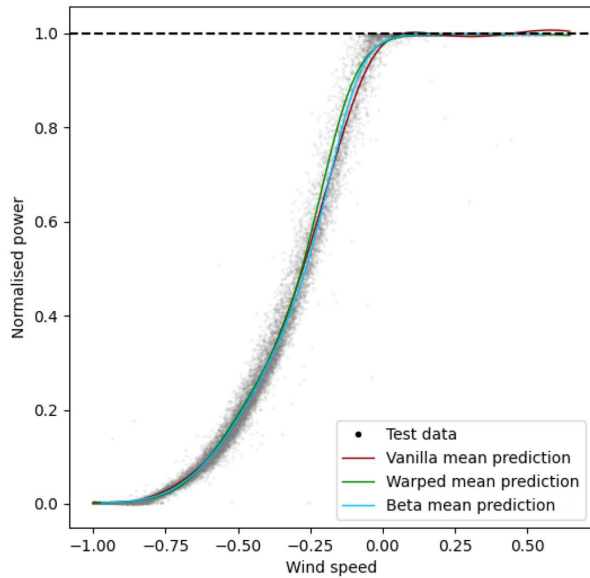


Figure 4. Mean comparison of models.

Simple GP regression

A standard (homoscedastic sparse) GP model was trained over the power curve to act as a benchmark to the two bounded models being tested. The standard GP used an additive composite kernel comprised of a Matérn 3/2 kernel and a Linear kernel. This combination was used to capture the upward trend of the data.

In Figure 5(a) the results of the standard GP model can be seen. The model seemingly performs well in its mean predictions with an NMSE of 0.24. It also shows a log likelihood of 32,312, less than half of that seen for the HBP, indicating worse performance in the uncertainty quantification. Figure 5(b) shows a magnified portion of the same model at rated power in order to examine how successful the model was at quantifying uncertainty at rated power. The standard GP model captured the shape of the power curve very well on average across the range of wind speeds, and the NMSE reflects this. The uncertainty prediction is comparatively poor; the uncertainty distribution is underestimated away from the bounds and overestimated close to them. This can be seen in that some data is not covered by the distribution in the middle of the power curve, and that the distributions at rated and zero power overestimate the variance.

While the GP has an excellent NMSE and good joint log likelihood, it is clear from visual examination of Figure 5(b) that the predicted mean and uncertainty do not obey physical limits and overestimate possible values of the rated power. Perhaps the point of most concern here, is that despite the excellent NMSE score, there are still predictions above the rated power output. This output is then not a particularly useful model if carried forward into further analysis, as the errors in the uncertainty and mean estimation could mask damage in the turbine or suggest that the financial performance of the wind could exceed what is physically possible.

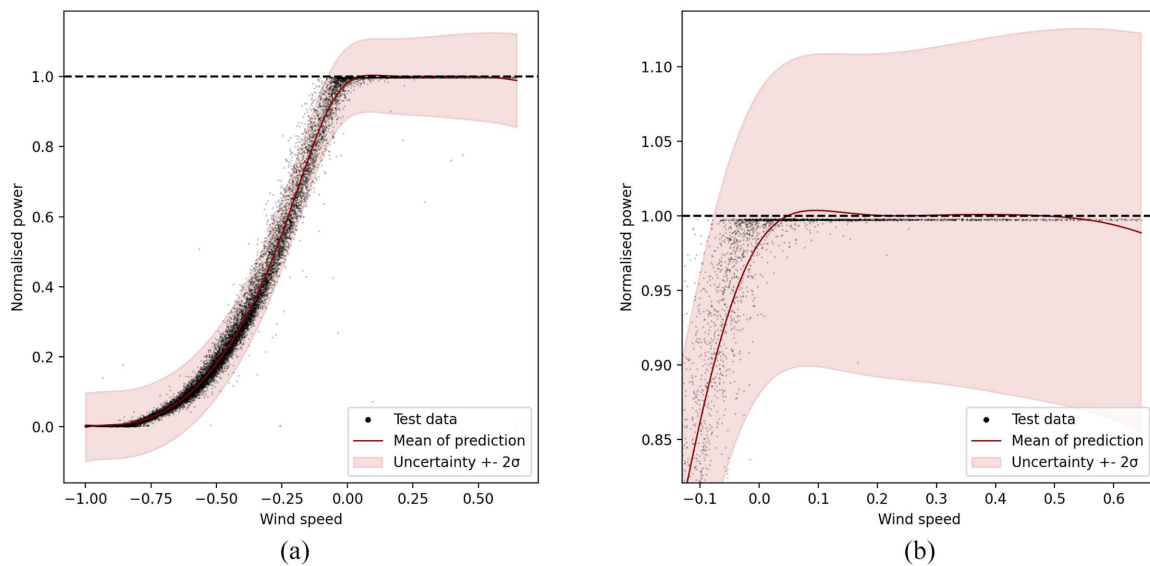
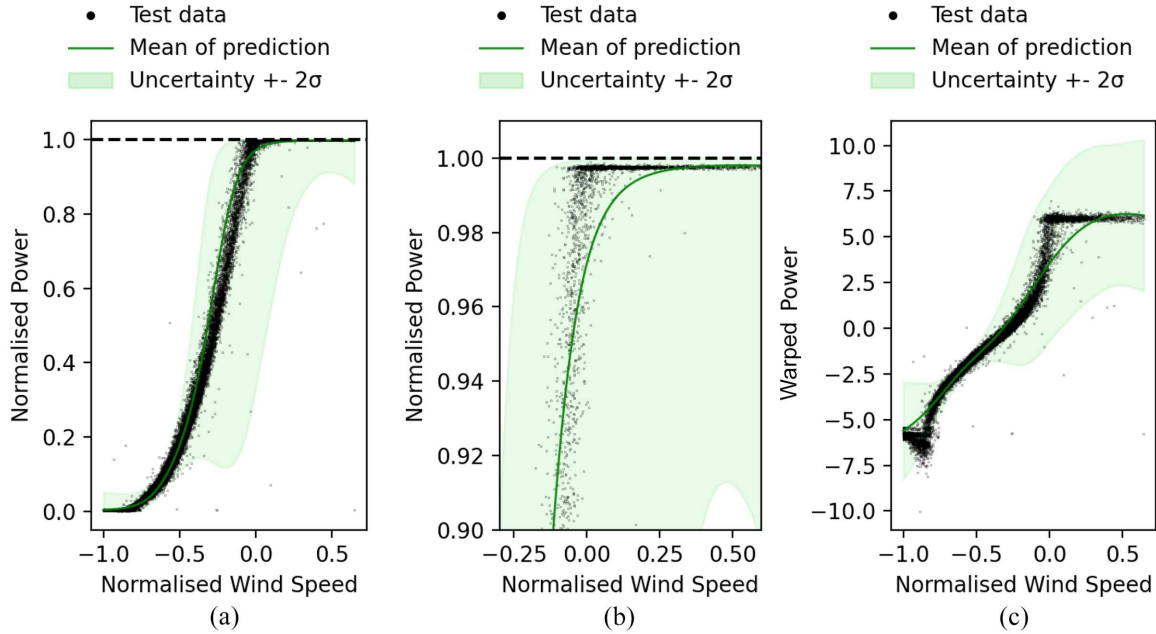


Figure 5. Standard GP result. (a) Power curve model of standard GP and (b) rated power model performance of standard GP. GP: Gaussian process.



**Figure 6.** Heteroscedastic GP results. (a) Power curve model of warped Heteroscedastic GP in the original data space, (b) magnification of rater power, and (c) Heteroscedastic GP in the warped space. GP: Gaussian process.

Finally, it is clear that the GP struggles to correctly predict the transition to rated power. This could be because continuous function models such as the GP can struggle with abrupt transitions such as this one. There are ways of addressing this shortcoming of GPS, such as using Bayesian committees to model the transition to rated power.<sup>5</sup>

### Heteroscedastic warped GP

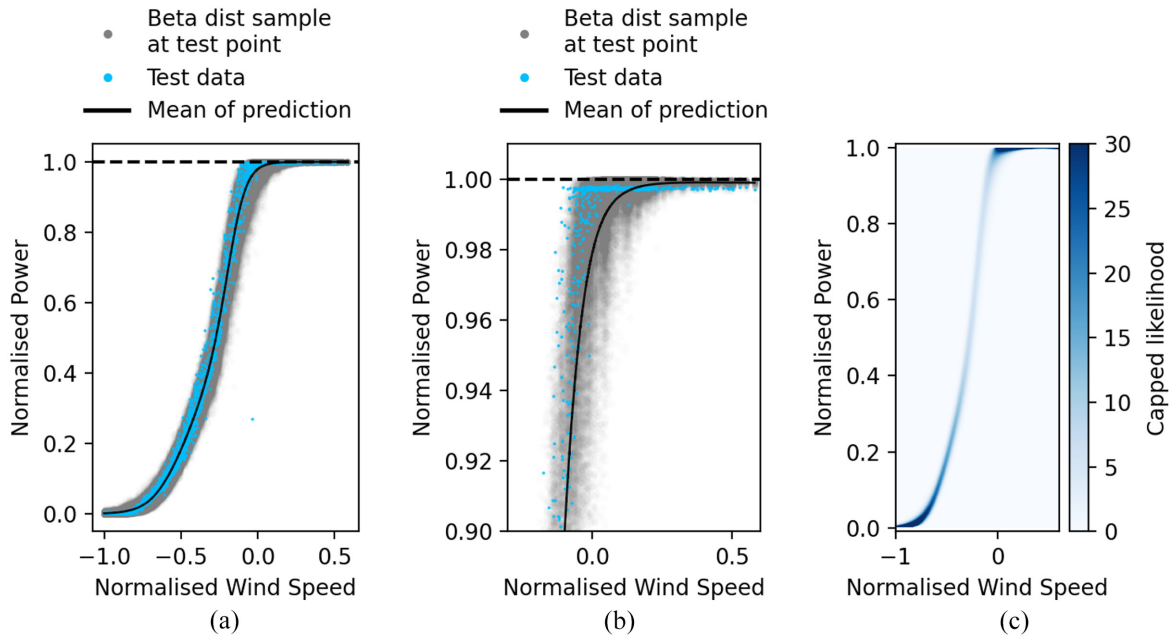
The prediction of the power curve with the warped GP is shown in Figure 6(a), and a magnified image of the model at rated power is shown in Figure 6(b). The performance metrics for this model were comparatively very poor, with an NMSE of 5.4 and a log likelihood of  $-12,469$ . The warping approach has failed to correctly capture either the mean behaviour or the uncertainty in the power curve.

The performance metrics clearly show that the warped GP has performed poorly. In order to understand why, the fitted GP must be examined in the warped space (shown in Figure 6(c)), as this is where the GP interacts with the data being modelled. It is clear from Figure 6(c) that the GP has failed to map the warped function correctly. This is likely because

warping the power curve has resulted in a data shape that is too noisy and discontinuous for the GP with a Gaussian likelihood to model well.

The inherent noise of the power curve in the power axis was exacerbated during the warping process, resulting in large vertical noise distributions that the GP cannot capture except by increasing the variance. This would confuse the training of the GP and is why such large uncertainty distributions are returned in Figure 6(a) and (c). The discontinuity of the warped shape is also a problem because the GP is a continuous function and struggles to model abrupt changes as discussed in the previous section and in Cornford et al.<sup>48</sup> This results in the GP smoothing over the abrupt changes in the model as seen in Figure 6(c) space. It appears that this challenge in modelling the discontinuous transition is amplified when projecting the data through the warping function.

Despite being an intuitive approach, the heteroscedastic warped GP is clearly not that useful in assessing either the predictive mean or the uncertainty. The issues with this model lie in the warped space, where the data shape becomes too noisy and discontinuous. These problems could be addressed in several ways such as with further cleaning in the warped space, with



**Figure 7.** HBP results. (a) HBP power curve, (b) magnification of rated power, and (c) capped heat map of the predictive uncertainty of the HBP.

HBP: Heteroscedastic Beta Process.

a Bayesian committee<sup>5</sup> or an alternative warping function (although it is not clear what that warping function may be).

### Beta likelihood

The results from the HBP model can be seen in Figure 7(a). A magnification of the model at rated power is shown in Figure 7(b). Figure 7(c) is a heat map of the predictive uncertainty over the test data in order to illustrate that the uncertainty is recoverable as a closed form solution. The heat map is capped at 30 due to the likelihood being so high towards the bounds.

The NMSE for this model is 0.53, while the log likelihood is 68,866. These are excellent results. The log likelihood is the highest of all the models tested, and shows that the bounded distribution best captures the power curve data. The NMSE is higher than for the standard GP, however the difference (0.53 vs 0.24) is small enough that these results can be considered comparable. It is worth noting again the previous discussion regarding the failure of the standard GP model to obey the physical bounds of the system in both the mean and the variance despite its lower NMSE.

Choosing between the standard GP and HBP will depend on the desired use case. When uncertainty matters, the HBP is clearly a superior model. If only the mean prediction matters then the standard GP could be considered the better choice; however it is worth

considering that the lower NMSE of the standard GP is misleading; the NMSE is simply an average of performance over the curve, and will not penalise physically implausible predictions. As seen in Figure 5(b) the predictive mean of the standard GP does exceed rated power, and as such is physically implausible. Blindly using the unbounded model could have serious consequences should its predictions be taken forward without further handling – such as over-predicting financial returns or masking turbine damage.

As seen in Figure 7(a) the model has captured the shape of the power curve well. Similarly to the other models, this model also struggles to accurately fit the transition to rated power. This is marginally less pronounced as the use of Beta distributions allow the mean to be ‘squashed’ into the transition better than a normal distribution would allow.

It is clear from the performance metrics that the HBP model is successful. It presents significant utility in that a full bounded predictive uncertainty distribution is returned. This uncertainty distribution also captures the data better than the standard GP and warped model. This has significant benefits as an improved power curve model for applications outside of SHM and as an improved starting point for turbine anomaly detection within SHM. The NMSE is comparable to the standard GP but offers the extra security that the mean prediction is bounded, removing the potential for error that an unbounded model carries and

providing the basis for increased operator confidence in the model.

## Conclusions

The novelty of this work was to propose and investigate two power curve models with physically meaningful predictions of mean and uncertainty. As discussed in ‘Related work’, this is a comparatively underdeveloped area of power curve modelling. Predictions of the mean and uncertainty that obey the physical laws imposed on the system being modelled are vital to understand in order to be able to extract meaningful information from these models.

Of the two methods investigated, the authors have argued that HBP model is the most successful. When compared to the traditional GP, the HBP model performs favourably. The joint likelihood is significantly higher than the standard GP, showing that the bounded uncertainty region better captures the test points than its unbounded counterpart. This is an especially important result when the use of predictive uncertainty in further calculations is considered. The difference in NMSE between the standard GP and HBP is small enough to be functionally comparable, and ignores the fact that some of the predictions returned by the standard GP are not physically plausible.

Predictions from the standard GP will not necessarily exceed the bounds with its mean predictions; however the model cannot provide confidence that its mean predictions are guaranteed to be physically plausible. The HBP is therefore superior in establishing operator trust and reliance on the power curve model, even when compared to a theoretically perfect standard GP.

In conclusion, the HBP model is an important step in better quantifying the uncertainty of power curve models. It was shown that bounding the GP to respect physical limitations significantly improves uncertainty quantification, a result with interesting ramifications for SHM and other industrial use moving forward.

## Further work

While the proposed methods for bounding the GP met with varying levels of success, both models open up avenues for further work. It was clear that the heteroscedastic GP did not perform adequately in the warped space. Several methods to remedy this exists, and although outside the scope of this current work, may be considered in the future, including using Bayesian committees to tackle the truncated nature of the

warped data shape or experimenting with alternative warping functions.

The GP with Beta likelihood was successful, and while the fast approximation of the likelihood with the heteroscedastic method worked well, the random nature of the sampling and unification of the various distributions into an average distribution at each test point is not the most statistically robust method. Uniting the Beta distributions of each sample mathematically rather than using the mean and total variance would be more robust. Finally, it may be worth considering again the use of committee machines to better handle the transition to rated power seen in the power curve.


## Declaration of conflicting interests


The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

## Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: The authors gratefully acknowledge the support of the Engineering Physical Sciences Research Council (EPSRC) through grant numbers EP/S001565/1 and EP/W002140/1. For the purpose of open access, the authors have applied a Creative Commons Attribution (CC-BY) licence to any Author Accepted Manuscript version arising.

## ORCID iDs

J.H. Mclean  <https://orcid.org/0000-0003-0594-7634>

A. E Maguire  <https://orcid.org/0000-0002-1004-1979>

## Notes

1. A heteroscedastic GP is a GP designed to model heteroscedastic noise, otherwise known as input-dependent noise. This will be expanded upon in the subsection ‘Heteroscedastic noise’.
2. Due to the confidentiality of the data, it is not possible to release the data used in this case study, however the code used to generate the results is made available at <https://github.com/JacquesMclean/Physically-meaningful-Gaussian-Processes>

## References

1. Papatheou E, Dervilis N, Maguire AE, et al. A performance monitoring approach for the novel Lillgrund offshore wind farm. *IEEE Trans Ind Electron* 2015; 62(10): 6636–6644.
2. Gonzalez E, Stephen B, Infield D, et al. Using high-frequency SCADA data for wind turbine performance

- monitoring: a sensitivity study. *Renewable Energy* 2019; 131: 841–853.
3. Nielsen JS, Tcherniak D and Ulriksen MD. A case study on risk-based maintenance of wind turbine blades with structural health monitoring. *Struct Infrastruct Eng* 2021; 17(3): 302–318.
  4. Hughes AJ, Barthorpe RJ, Dervilis N, et al. A probabilistic risk-based decision framework for structural health monitoring. *Mech Syst Signal Process* 2021; 150: 107339.
  5. Rogers TJ, Gardner P, Dervilis N, et al. Probabilistic modelling of wind turbine power curves with application of heteroscedastic Gaussian process regression. *Renewable Energy* 2020; 148: 1124–1136.
  6. Gamerman D and Lopes HF. *Markov chain Monte Carlo: stochastic simulation for Bayesian inference*. New York, NY: CRC Press, 2006.
  7. Martinez-Luengo M, Kolios A and Wang L. Structural health monitoring of offshore wind turbines: a review through the statistical pattern recognition paradigm. *Renewable Sustainable Energy Rev* 2016; 64: 91–105.
  8. Márquez FPG, Tobias AM, Pérez JMP, et al. Condition monitoring of wind turbines: techniques and methods. *Renewable Energy* 2012; 46: 169–178.
  9. Ciang CC, Lee JR and Bang HJ. Structural health monitoring for a wind turbine system: a review of damage detection methods. *Meas Sci Technol* 2008; 19(12): 122001.
  10. Dutton A, Blanch M, Vionis P, et al. Acoustic emission condition monitoring of wind turbine rotor blades: laboratory certification testing to large scale in-service deployment. In: *European Wind Energy Conference-EWEC*, Madrid, 2003. Citeseer.
  11. Raišutis R, Jasiūnienė E and Žukauskas E. Ultrasonic NDT of wind turbine blades using guided waves. *Ultrasonics/Ultrasound* 2008; 63(1): 7–11.
  12. Dutton A. Thermoelastic stress measurement and acoustic emission monitoring in wind turbine blade testing. In: *European Wind Energy Conference London*, 2004. pp. 22–25.
  13. Canizo M, Onieva E, Conde A, et al. Real-time predictive maintenance for wind turbines using big data frameworks. In: *2017 IEEE international conference on prognostics and health management (ICPHM)*, Dallas, 2017, pp. 70–77. IEEE.
  14. Zaher A, McArthur S, Infield D, et al. Online wind turbine fault detection through automated SCADA data analysis. *Wind Energy* 2009; 12(6): 574–593.
  15. Yang W, Court R and Jiang J. Wind turbine condition monitoring by the approach of SCADA data analysis. *Renewable Energy* 2013; 53: 365–376.
  16. Uluyol O, Parthasarathy G, Foslien W, et al. Power curve analytic for wind turbine performance monitoring and prognostics. In: *Annual conference of the PHM society*, Montreal, 2011, vol. 3.
  17. Lydia M, Kumar SS, Selvakumar AI, et al. A comprehensive review on wind turbine power curve modeling techniques. *Renewable Sustainable Energy Rev* 2014; 30: 452–460.
  18. Khalfallah MG and Koliub AM. Suggestions for improving wind turbines power curves. *Desalination* 2007; 209(1–3): 221–229.
  19. Saint-Drenan YM, Besseau R, Jansen M, et al. A parametric model for wind turbine power curves incorporating environmental conditions. *Renewable Energy* 2020; 157: 754–768.
  20. Raj MM, Alexander M and Lydia M. Modeling of wind turbine power curve. In: *ISGT2011-India*, India, 2011, pp. 144–148. New York: IEEE.
  21. Bishop CM and Nasrabadi NM. *Pattern recognition and machine learning*, 2006, vol. 4. Springer.
  22. Kusiak A, Zheng H and Song Z. On-line monitoring of power curves. *Renewable Energy* 2009; 34(6): 1487–1493.
  23. Schlechtingen M, Santos IF and Achiche S. Using data-mining approaches for wind turbine power curve monitoring: a comparative study. *IEEE Trans Sustainable Energy* 2013; 4(3): 671–679.
  24. Ouyang T, Kusiak A and He Y. Modeling wind-turbine power curve: a data partitioning and mining approach. *Renewable Energy* 2017; 102: 1–8.
  25. Üstüntaş T and Şahin AD. Wind turbine power curve estimation based on cluster center fuzzy logic modeling. *J Wind Eng Ind Aerodyn* 2008; 96(5): 611–620.
  26. Gill S, Stephen B and Galloway S. Wind turbine condition assessment through power curve copula modeling. *IEEE Trans Sustainable Energy* 2011; 3(1): 94–101.
  27. Marčiukaitis M, Žutautaitė I, Martišauskas L, et al. Non-linear regression model for wind turbine power curve. *Renewable Energy* 2017; 113: 732–741.
  28. Li S, Wunsch DC, O’Hair EA, et al. Using neural networks to estimate wind turbine power generation. *IEEE Trans Energy Convers* 2001; 16(3): 276–282.
  29. Panahi D, Deilami S and Masoum MA. Evaluation of parametric and non-parametric methods for power curve modelling of wind turbines. In: *2015 9th international conference on electrical and electronics engineering (ELECO)*, Bursa, 2011, pp. 996–1000. New York: IEEE.
  30. Pelletier F, Masson C and Tahan A. Wind turbine power curve modelling using artificial neural network. *Renewable Energy* 2016; 89: 207–214.
  31. Jin T and Tian Z. Uncertainty analysis for wind energy production with dynamic power curves. In: *2010 IEEE 11th international conference on probabilistic methods applied to power systems*, Singapore, 2010, pp. 745–750. IEEE.
  32. Yan J, Zhang H, Liu Y, et al. Uncertainty estimation for wind energy conversion by probabilistic wind turbine power curve modelling. *Appl Energy* 2019; 239: 1356–1370.
  33. Xu K, Yan J, Zhang H, et al. Quantile based probabilistic wind turbine power curve model. *Appl Energy* 2021; 296: 116913.
  34. Ge S, Zuo MJ and Tian Z. Wind turbine power output estimation with probabilistic power curves. In: *2020 Asia-Pacific international symposium on advanced reliability and maintenance modeling (APARM)*, pp. 1–6. IEEE.

35. Yun E and Hur J. Probabilistic estimation model of power curve to enhance power output forecasting of wind generating resources. *Energy* 2021; 223: 120000.
36. Pandit RK and Infield D. SCADA-based wind turbine anomaly detection using Gaussian process models for wind turbine condition monitoring purposes. *IET Renewable Power Gener* 2018; 12(11): 1249–1255.
37. Pandit RK and Infield D. Comparative analysis of Gaussian process power curve models based on different stationary covariance functions for the purpose of improving model accuracy. *Renewable Energy* 2019; 140: 190–202.
38. Rasmussen CE. Gaussian processes in machine learning. In: *Summer school on machine learning*, 2004, pp. 63–71. Springer.
39. Williams CK and Rasmussen CE. *Gaussian processes for machine learning*, vol. 2. Cambridge, MA: MIT Press, 2006.
40. Lázaro-Gredilla M and Titsias MK. Variational heteroscedastic Gaussian process regression. In: *International conference on machine learning*, Washington, 2011.
41. Swiler LP, Gulian M, Frankel AL, et al. A survey of constrained Gaussian process regression: approaches and implementation challenges. *J Mach Learn Model Comput* 2020; 1(2): 7.
42. Jensen BS, Nielsen JB and Larsen J. Bounded Gaussian process regression. In: *2013 IEEE international workshop on machine learning for signal processing (MLSP)*, Southampton, 2013, pp. 1–6. IEEE.
43. Matthews AGdG, van der Wilk M, Nickson T, et al. GPflow: a Gaussian process library using TensorFlow. *J Mach Learn Res* 2017; 18(40): 1–6.
44. Rasmussen CE and Nickisch H. Gaussian processes for machine learning (GPML) toolbox. *J Mach Learn Res* 2010; 11: 3011–3015.
45. Hensman J, Fusi N and Lawrence ND. Gaussian processes for big data. *arXiv preprint arXiv:13096835* 2013.
46. Saul AD, Hensman J, Vehtari A, et al. Chained Gaussian processes. In: *Artificial intelligence and statistics*, pp. 1431–1440. PMLR.
47. Stein ML. *Interpolation of spatial data: some theory for kriging*. Berlin: Springer Science & Business Media, 1999.
48. Cornford D, Nabney I and Williams C. Adding constrained discontinuities to Gaussian process models of wind fields. *Adv Neural Inf Process Syst* 1998; 11: 1.