



Gaussian Process-Aided Function Comparison Using Noisy Scattered Data

Abhinav Prakash, Rui Tuo & Yu Ding

To cite this article: Abhinav Prakash, Rui Tuo & Yu Ding (2021): Gaussian Process-Aided Function Comparison Using Noisy Scattered Data, Technometrics, DOI: [10.1080/00401706.2021.1905073](https://doi.org/10.1080/00401706.2021.1905073)

To link to this article: <https://doi.org/10.1080/00401706.2021.1905073>



View supplementary material [↗](#)



Published online: 30 Apr 2021.



Submit your article to this journal [↗](#)



Article views: 105



View related articles [↗](#)



View Crossmark data [↗](#)



Gaussian Process-Aided Function Comparison Using Noisy Scattered Data

Abhinav Prakash, Rui Tuo, and Yu Ding

Department of Industrial and Systems Engineering, Texas A&M University, College Station, TX

ABSTRACT

This work proposes a nonparametric method to compare the underlying mean functions given two noisy datasets. The motivation for the work stems from an application of comparing wind turbine power curves. Comparing wind turbine data presents new problems, namely the need to identify the regions of difference in the input space and to quantify the extent of difference that is statistically significant. Our proposed method, referred to as funGP, estimates the underlying functions for different data samples using Gaussian process models. We build a confidence band using the probability law of the estimated function differences under the null hypothesis. Then, the confidence band is used for the hypothesis test as well as for identifying the regions of difference. This identification of difference regions is a distinct feature, as existing methods tend to conduct an overall hypothesis test stating whether two functions are different. Understanding the difference regions can lead to further practical insights and help devise better control and maintenance strategies for wind turbines. The merit of funGP is demonstrated by using three simulation studies and four real wind turbine datasets.

ARTICLE HISTORY

Received January 2020
Accepted March 2021

KEYWORDS

Difference region identification; Functional hypothesis test; Wind power curves

1. Introduction

Comparing information from two datasets is an important topic in statistics. Various methods exist to compare datasets arising from univariate and multivariate distributions, for example, two sample t -test (Fisher 1925) and Hotelling's T^2 test (Hotelling 1931), respectively. The literature is not just limited to comparing finite-dimensional objects, but also extends to functions. In this article, we focus on nonparametric methods that compare functions.

Our work is motivated by an application in the wind energy sector, where the goal is to compare two power curves. The power curve of a wind turbine is a function that maps wind speed and other environmental inputs to wind power output. Power curves are used to characterize the performance of wind turbines (IEC 2005). Hence, comparing power curves plays a critical role in assessing and benchmarking turbine performance, devising maintenance plans, and justifying expensive overhauls or retrofits (Hwangbo, Johnson, and Ding 2017; Ding 2019). Some important aspects for comparing power curves are to understand where the power curves differ (in the input space) and how much is the difference. It is not only important to check whether two power curves differ, but more crucial to identify the regions of difference and quantify the difference for guiding economically justifiable maintenance and repairs. In addition to the need for identifying the regions of difference, the datasets arising from wind turbines entails two other features: the input conditions (e.g., wind speed or wind direction) for the observations cannot be controlled, and as a result, the input points for any two datasets are not the same, and there lacks replicates for any input point. Taken altogether, our research

objective is to develop a nonparametric function comparison method that meets the following three requirements:

- The method can identify the regions between two functions that are statistically different and quantify the difference,
- The input data points associated with the two functions are not necessarily the same.
- There are no replications in the data points.

The problem of testing the equality of two nonparametric functions has been studied extensively in the literature. One early work is Hall and Hart (1990). They defined a test statistic for the problem using the smoothed (estimated) function values and obtained a distribution of their test statistic using the bootstrap method. King, Hart, and Wehrly (1991) also studied the same problem using smoothing techniques and proposed an exact distribution for their test statistic under the normality assumption for the errors. Delgado (1993) proposed another test statistic using the marked empirical process. Fan and Lin (1998) worked on reducing the dimension of the problem using discrete Fourier transforms so that standard multivariate techniques can be used to test the hypothesis. These works assume that the two datasets under comparison have identical input points. This assumption is relaxed in Kulasekera (1995), Kulasekera and Wang (1997), Munk and Dette (1998), and Neumeyer and Dette (2003), which proposed tests that are valid under different input points among the datasets. The literature mentioned hitherto are *global* tests, providing a binary answer on whether the functions are statistically the same or not. They do not provide any insights on the regions of the input space where the functions are different, or which function has higher or lower function values.

Cox and Lee (2008) addressed this problem of identifying difference regions using a pointwise testing procedure based on the Westfall-Young randomization technique (Westfall and Young 1993). However, Cox and Lee's method does not meet our aforementioned requirements, because (a) Cox and Lee (2008) required replications of data points, as Cox and Lee's method is based on permutation, which requires data replication. (b) Their method produces pointwise p -values rather than a coherent confidence band for functional differences. It is not straightforward to convert the pointwise p -values into functional differences. (c) Cox and Lee (2008) developed their method for the cases when the input points for the two functions are at the same locations.

In this work, we propose a new nonparametric function comparison method that satisfies the three requirements posed above. We first use a Gaussian process (GP) regression model to recover the functions from the noisy datasets. Given a prescribed Type I error, we then build a confidence band on the difference between the functions throughout the input space under the null hypothesis. If the actual difference between the functions computed using the data are beyond the confidence band, then we reject the null hypothesis. We call the method *function comparison using Gaussian Process* or *funGP*.

When the null hypothesis is rejected, funGP identifies the regions of difference in the input space and quantifies the estimated difference using the established confidence band. Although we assume the functions as realizations of GPs, we demonstrate that the method works well for deterministic functions also. We apply our method to real wind turbine datasets and compare the results with some existing work for turbine performance characterization. That GP regression works for a large class of functions makes the proposed method applicable to many problems.

We organize the rest of the article as follows. Section 2 provides the details of the proposed method. Section 3 presents the simulation examples and comparison studies with two existing functional tests. We apply the funGP method to wind turbine datasets in Section 4. We conclude the work with some discussions in Section 5.

2. The funGP Method

In this section, we describe the mathematical formulation and the implementation details of the proposed funGP method.

2.1. Problem Formulation

Let us consider two datasets, $\{\mathcal{D}_i \mid i = 1, 2\}$, with n_1 and n_2 data points, respectively. Assume that \mathcal{D}_1 can be denoted by an ordered pair $\{\mathbf{X}^{(1)}, \mathbf{y}^{(1)}\}$, where $\mathbf{X}^{(1)}$ is an $n_1 \times d$ matrix with each row corresponding to input variable values for one data point and $\mathbf{y}^{(1)}$ is a vector of length n_1 with each component as response for one data point. Similarly, \mathcal{D}_2 can be denoted as $\{\mathbf{X}^{(2)}, \mathbf{y}^{(2)}\}$. Specifically,

$$\mathbf{y}^{(1)} = \begin{bmatrix} y_{11} \\ y_{12} \\ \vdots \\ y_{1n_1} \end{bmatrix}, \quad \mathbf{X}^{(1)} = \begin{bmatrix} -\mathbf{x}_{11}^\top - \\ -\mathbf{x}_{12}^\top - \\ \vdots \\ -\mathbf{x}_{1n_1}^\top - \end{bmatrix},$$

$$\mathbf{y}^{(2)} = \begin{bmatrix} y_{21} \\ y_{22} \\ \vdots \\ y_{2n_2} \end{bmatrix}, \quad \mathbf{X}^{(2)} = \begin{bmatrix} -\mathbf{x}_{21}^\top - \\ -\mathbf{x}_{22}^\top - \\ \vdots \\ -\mathbf{x}_{2n_2}^\top - \end{bmatrix}.$$

We also assume that these datasets come from underlying models given by

$$y_{ij} = f_i(\mathbf{x}_{ij}) + \epsilon_{ij}, \quad i = 1, 2, \quad j = 1, \dots, n_i, \quad (1)$$

where $f_1(\cdot)$ and $f_2(\cdot)$ are two smooth continuous functions with the same compact domain $\mathcal{X} \subset \mathbb{R}^d$ and $\epsilon_{ij} \stackrel{iid}{\sim} \mathcal{N}(0, \sigma_\epsilon^2)$ with a constant variance $\sigma_\epsilon^2 < \infty$.

The goal is to test the following null and alternative hypotheses. The null hypothesis is that the functions are identical, whereas the alternative hypothesis is that the functions differ for at least one $\mathbf{x} \in \mathcal{X}$. Under the null hypothesis, H_0

$$f_1(\mathbf{x}) = f_2(\mathbf{x}) \quad \text{for all } \mathbf{x} \in \mathcal{X}.$$

And, under the alternative hypothesis, H_1

$$\text{there exists } \mathbf{x} \in \mathcal{X} \quad \text{such that} \quad f_1(\mathbf{x}) \neq f_2(\mathbf{x}).$$

A rigorous frequentist testing of the null hypothesis H_0 usually relies on a test statistic whose distribution is (approximately) independent of the underlying function $f := f_1 = f_2$ under H_0 . One would consider using the estimator of $f_1 - f_2$ to build a test statistic. Specifically, we in this work invoke a GP framework for calculating the distribution in the presence of an unknown f . This assumption allows for calculating the distribution of an intuitive estimator of $f_1 - f_2$. Of course, doing this requires us to replace the null hypothesis H_0 . Details will be presented in the next subsection.

In addition to the above discussion, we stress that our application requires the test statistic to be a functional statistic, as we are interested in identifying the region of input space where the functions are different. Specifically, when H_0 is rejected, we need to identify the set $\mathcal{S} = \{\mathbf{x} : f_1(\mathbf{x}) \neq f_2(\mathbf{x})\}$. Most of the existing methods reviewed in Section 1, such as Munk and Dette (1998), use a univariate statistic to test the hypothesis and cannot identify the region of difference. Under certain conditions, such as the input points of the two datasets are identical and replicated response, a functional test is available; see Cox and Lee (2008). To the best of our knowledge, no statistics have been proposed in the literature under the general conditions as in the current context.

2.2. Hypothesis Testing With a GP Prior

The general idea for a hypothesis testing is to find a test statistic and subsequently have a decision rule to either accept or reject the null hypothesis based on the value of the test statistic. In the application described, we are not only interested in the binary answer that whether the two functions are different, but also want to understand where the difference lies in the input space. This requires us to obtain a test statistic at the input points for which we do not have any data. We also assume that the input points for the two datasets are not the same. Thus, we would have to assume some structure in the functions (such as the

functions are smooth and continuous) in order to recover the functions and estimate the noise in the model. Here, we adopt a Bayesian idea that imposes a prior structure on the functions. Specifically, we use a GP prior with zero mean and a covariance function given by $k(\mathbf{x}, \mathbf{x}')$. The zero mean assumption is for mathematical simplicity, and we can assume a different mean function, if necessary.

Despite the use of the GP prior, we still follow a frequentist hypothesis testing framework, by considering a new null hypothesis H_0^{GP} , still stating $f_1(\mathbf{x}) = f_2(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X}$, but incorporating the following prior information:

$$\begin{aligned} y_{ij} &= f(x_{ij}) + \epsilon_{ij}, \\ f &\sim \mathcal{GP}(0, k(\mathbf{x}, \mathbf{x}')), \end{aligned}$$

where $\epsilon_{ij} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma_\epsilon^2)$. In Sections 2.2 and 2.3, we assume that $k(\cdot, \cdot)$ and σ_ϵ^2 are known.

So far, our goal can be described as testing H_0^{GP} against the alternative hypothesis H_1 . We will propose a test method, so that its Type I error under H_0^{GP} has a probability controlled by a prespecified significance level α . Note that the Type I error under H_0^{GP} is

$$\int_{\mathcal{X}} \mathbb{P}(H_0^{\text{GP}} \text{ is rejected} | f_1 = f_2 = f) d\mathbb{P}_{\text{GP}},$$

where \mathbb{P}_{GP} denotes the probability measure of the GP prior. It is worth noting that the Type I above is not identical to the Type I error under the original null hypothesis H_0 . However, we expect that the proposed method can serve as an approximate method for the fixed-function testing problems, and we will verify this expectation via numerical studies in Section 3.

The main idea of our test is as follows. First, we can reconstruct f based on the datasets \mathcal{D}_1 and \mathcal{D}_2 separately, and denote the reconstructed functions as \hat{f}_1 and \hat{f}_2 , respectively. Under H_0^{GP} , \hat{f}_1 and \hat{f}_2 should be close. Thus, we can test H_0^{GP} by computing the difference between \hat{f}_1 and \hat{f}_2 .

To reconstruct f , we start off by defining a cross-covariance matrix $\mathbf{K}_{\mathbf{X}, \mathbf{X}'}$ between a pair of input variable matrix \mathbf{X} and \mathbf{X}' , and a covariance vector $\mathbf{r}(\mathbf{x})$ between the input data \mathbf{X} and any point \mathbf{x} as follows:

$$\begin{aligned} \mathbf{K}_{\mathbf{X}, \mathbf{X}'} &= \begin{bmatrix} k(\mathbf{x}_1, \mathbf{x}'_1) & k(\mathbf{x}_1, \mathbf{x}'_2) & \dots & k(\mathbf{x}_1, \mathbf{x}'_n) \\ k(\mathbf{x}_2, \mathbf{x}'_1) & k(\mathbf{x}_2, \mathbf{x}'_2) & \dots & k(\mathbf{x}_2, \mathbf{x}'_n) \\ \vdots & \vdots & \ddots & \vdots \\ k(\mathbf{x}_m, \mathbf{x}'_1) & k(\mathbf{x}_m, \mathbf{x}'_2) & \dots & k(\mathbf{x}_m, \mathbf{x}'_n) \end{bmatrix}, \\ \mathbf{r}(\mathbf{x}) &= \begin{bmatrix} k(\mathbf{x}_1, \mathbf{x}) \\ k(\mathbf{x}_2, \mathbf{x}) \\ \vdots \\ k(\mathbf{x}_m, \mathbf{x}) \end{bmatrix}, \end{aligned} \quad (2)$$

where $\mathbf{x}_1 \dots \mathbf{x}_m$ are the vectors in the rows of the matrix \mathbf{X} , and $\mathbf{x}'_1 \dots \mathbf{x}'_n$ are the vectors in the rows of the matrix \mathbf{X}' . When $\mathbf{X} = \mathbf{X}'$, $\mathbf{K}_{\mathbf{X}, \mathbf{X}}$ is then a symmetric covariance matrix. The standard GP prediction theory suggests (Rasmussen and Williams 2006)

$$\hat{f}_1(\mathbf{x}) = \mathbf{r}_1(\mathbf{x})^\top [\mathbf{K}_{\mathbf{X}^{(1)}, \mathbf{X}^{(1)}} + \sigma_\epsilon^2 \mathbf{I}_{n_1}]^{-1} \mathbf{y}^{(1)}, \quad (3)$$

$$\hat{f}_2(\mathbf{x}) = \mathbf{r}_2(\mathbf{x})^\top [\mathbf{K}_{\mathbf{X}^{(2)}, \mathbf{X}^{(2)}} + \sigma_\epsilon^2 \mathbf{I}_{n_2}]^{-1} \mathbf{y}^{(2)}, \quad (4)$$

where $\mathbf{r}_1(\mathbf{x})$ is the covariance vector between $\mathbf{X}^{(1)}$ and any point \mathbf{x} , $\mathbf{K}_{\mathbf{X}^{(1)}, \mathbf{X}^{(1)}}$ is the covariance matrix for $\mathbf{X}^{(1)}$, and \mathbf{I}_{n_1} is the identity matrix of proper size— $n_1 \times n_1$ in this case. The notations in Equation (4) are likewise defined.

It is worth noting that although $\hat{f}_1(\mathbf{x})$ and $\hat{f}_2(\mathbf{x})$ are posterior means from a Bayesian perspective, here we take a frequentist point of view and regard them merely as statistics, that is, functions of the data. To test the null hypothesis H_0^{GP} , we use the statistic $G(\mathbf{x}) := \hat{f}_1(\mathbf{x}) - \hat{f}_2(\mathbf{x})$. Clearly, given $(\mathbf{X}^{(1)}, \mathbf{X}^{(2)})^\top$, the randomness of $G(\mathbf{x})$ comes solely from the data $(\mathbf{y}^{(1)}, \mathbf{y}^{(2)})^\top$, which follows a zero-mean multivariate normal distribution under H_0^{GP} . Therefore, under H_0^{GP} and given the input data, $G(\mathbf{x})$ is a centered GP and we write

$$G(\cdot) | \mathbf{X}^{(1)}, \mathbf{X}^{(2)}, H_0^{\text{GP}} \sim \mathcal{GP}(0, c(\cdot, \cdot)). \quad (5)$$

The expression for covariance function $c(\cdot, \cdot)$ is given as follows (see supplementary material of Section S.1 for details):

$$\begin{aligned} c(\mathbf{x}, \mathbf{x}') &= \mathbf{r}_2(\mathbf{x})^\top [\mathbf{K}_{\mathbf{X}^{(2)}, \mathbf{X}^{(2)}} + \sigma_\epsilon^2 \mathbf{I}_{n_2}]^{-1} \mathbf{r}_2(\mathbf{x}') \\ &\quad + \mathbf{r}_1(\mathbf{x})^\top [\mathbf{K}_{\mathbf{X}^{(1)}, \mathbf{X}^{(1)}} + \sigma_\epsilon^2 \mathbf{I}_{n_1}]^{-1} \mathbf{r}_1(\mathbf{x}') \\ &\quad - 2\mathbf{r}_2(\mathbf{x})^\top [\mathbf{K}_{\mathbf{X}^{(2)}, \mathbf{X}^{(1)}} \\ &\quad + \sigma_\epsilon^2 \mathbf{I}_{n_2}]^{-1} \mathbf{K}_{\mathbf{X}^{(2)}, \mathbf{X}^{(1)}} [\mathbf{K}_{\mathbf{X}^{(1)}, \mathbf{X}^{(1)}} \\ &\quad + \sigma_\epsilon^2 \mathbf{I}_{n_1}]^{-1} \mathbf{r}_1(\mathbf{x}'). \end{aligned}$$

Therefore, to test H_0^{GP} , it suffices to find a $1 - \alpha$ probability band of $G(\cdot)$ under H_0^{GP} , that is, a pair of functions $l(\mathbf{x})$ and $u(\mathbf{x})$ such that

$$\mathbb{P}_{G \sim \mathcal{GP}(0, c(\cdot, \cdot))} (l(\mathbf{x}) \leq G(\mathbf{x}) \leq u(\mathbf{x}) \text{ for all } \mathbf{x}) \geq 1 - \alpha.$$

With a slight abuse of terminology, we shall call the band between $l(\mathbf{x})$ and $u(\mathbf{x})$ a $1 - \alpha$ confidence band for $\mathcal{GP}(0, c(\cdot, \cdot))$. It is worth noting that this band is related to the distribution of $G(\mathbf{x})$ only under the null hypothesis H_0^{GP} . The test then proceeds by checking whether $G(\cdot)$ remains within the confidence band. If there exists an \mathbf{x} for which $G(\mathbf{x})$ is outside the band, then we reject the null hypothesis. Clearly, such a testing method ensures a $1 - \alpha$ Type I error under H_0^{GP} . The question now is how to efficiently build an effective confidence band for $\mathcal{GP}(0, c(\cdot, \cdot))$.

2.3. Building the Confidence Band

To build a $1 - \alpha$ confidence band for $\mathcal{GP}(0, c(\cdot, \cdot))$, the main idea is to sample from a set with a coverage probability of $1 - \alpha$. For notational simplicity, we suppose $G(\mathbf{x}) \sim \mathcal{GP}(0, c(\cdot, \cdot))$ in this subsection, that is, the null hypothesis H_0^{GP} is true. Since $G(\mathbf{x})$ is an infinite-dimensional object, it is more convenient to work with a finite-dimensional representation of it. To this end, we employ the Karhunen Loève (KL) expansion on $G(\mathbf{x})$. The KL expansion for the zero mean GP $G(\mathbf{x})$ is given as follows: $G(\mathbf{x}) = \sum_{k=1}^{\infty} \sqrt{\lambda_k} \phi_k(\mathbf{x}) z_k$, where $\{z_k\}_{k=1}^{\infty}$ are uncorrelated standard normal random variables, $\{\phi_k(\cdot)\}_{k=1}^{\infty}$ are the basis functions, $\{\lambda_k\}_{k=1}^{\infty}$ are the eigenvalues. In practice, this infinite sum is truncated for two reasons: (i) If the process is smooth, the eigenvalues would decay rapidly, (ii) To make the computation tractable. Under the assumption that the underlying functions under comparison are smooth, we discard all the eigenvalues smaller than a certain threshold. For practical purposes, we find

a threshold value of $10^{-6} \times \lambda_{\max}$ to work well, where λ_{\max} is the largest eigenvalue. Thus, the KL expansion decomposes the process into independent components and also reduces the dimension of the problem by finding a sparse representation of the process using its eigenfunction basis.

Let the truncation number computed from the aforementioned rule be m , that is, only the m largest eigenvalues are significantly “large”. Then, we write the truncated KL expansion for $G(\mathbf{x})$ as follows:

$$G(\mathbf{x}) \approx \sum_{k=1}^m \sqrt{\lambda_k} \phi_k(\mathbf{x}) z_k. \quad (6)$$

In Equation (6), the randomness is introduced by z_k 's. Hence, in order to build a $1 - \alpha$ confidence band on $G(\mathbf{x})$, we build the same level confidence band for the joint distribution of $z_k \mid k = 1, \dots, m$, which can be constructed as follows. We know that for an m -dimensional uncorrelated standard normal vector, a confidence region \mathcal{R} , with probability $\mathbb{P}(\mathcal{R}) = 1 - \alpha$, can be described using a hypersphere of radius r . This radius can be expressed as $r = \|\mathbf{z}\|$, where $\|\mathbf{z}\| = \|z_1^2 + z_2^2 + \dots + z_m^2\|$ is the ℓ^2 norm. The sum of squares of m uncorrelated standard normals follows a chi-square distribution with m degrees of freedom, that is, $r^2 = z_1^2 + z_2^2 + \dots + z_m^2 \sim \chi_m^2$. Hence, r is computed by inverting the CDF of a chi-square distribution in the following way:

$$r = \sqrt{F_m^{-1}(1 - \alpha)}, \quad (7)$$

where $F_m^{-1}(\cdot)$ is the inverse CDF of χ_m^2 . Once we have the radius r , we sample \mathbf{z} from the region with a coverage probability of $1 - \alpha$ using the following rule:

- sample z_i from $\mathcal{N}(0, 1) \mid i \in \{1 \dots m\}$,
- if $\sum_{i=1}^m z_i^2 \leq r^2$, accept $\mathbf{z} = (z_1, z_2, \dots, z_m)^\top$; else, reject it.

The samples of \mathbf{z} obtained as above can be easily converted to samples from $G(\mathbf{x})$ that are from $1 - \alpha$ confidence set using Equation (6). In order to test the hypothesis, we would need to compare the actual difference in the predictive means $g(\mathbf{x})$ with the confidence band at all the points $\mathbf{x} \in \mathcal{X}$. This is practically intractable, as for any continuous function, there are infinitely many points in the domain. Hence, we discretize the domain using a finite-sized evenly spaced test grid to approximate \mathcal{X} . Let \mathbf{X}_{test} be a $n_{\text{test}} \times d$ matrix with each row corresponding to one grid point $\mathbf{x}_{t_j} \mid j = 1, \dots, n_{\text{test}}$. We compare the function difference with the confidence band on these grid points. Testing on this regular grid is a reasonable approximation to testing for all $\mathbf{x} \in \mathcal{X}$ because of our underlying assumption that the functions are continuous and smooth. Let $\mathbf{C}_{\mathbf{X}_{\text{test}}, \mathbf{X}_{\text{test}}}$ be the covariance matrix generated using the covariance function $c(\mathbf{x}, \mathbf{x}')$ using all the points in \mathbf{X}_{test} in a similar way as $\mathbf{K}_{\mathbf{X}, \mathbf{X}'}$ is defined in Equation (2). Let $\mathbf{\Lambda}$ be an $m \times m$ diagonal matrix with m largest eigenvalues of $\mathbf{C}_{\mathbf{X}_{\text{test}}, \mathbf{X}_{\text{test}}}$ and let \mathbf{U} be an $n_{\text{test}} \times m$ matrix whose columns are the eigenvectors corresponding to the m largest eigenvalues of $\mathbf{C}_{\mathbf{X}_{\text{test}}, \mathbf{X}_{\text{test}}}$. Then, the KL expansion at all the points in \mathbf{X}_{test} , denoted by a random vector \mathbf{G} such that its j th component $(\mathbf{G})_j = G(\mathbf{x}_{t_j})$, can be expressed using the matrix

notation as follows (see Supplementary Material Section S.2 for details):

$$\mathbf{G} = \mathbf{U} \mathbf{\Lambda}^{\frac{1}{2}} \mathbf{z}. \quad (8)$$

In order to construct the confidence band, we sample a large number (say 1000) of \mathbf{z} from its confidence set, then the values of the confidence band at all points in \mathbf{X}_{test} is given by the vectors

$$\begin{aligned} \mathbf{ub} &= \max_{\mathbf{z}} \mathbf{U} \mathbf{\Lambda}^{\frac{1}{2}} \mathbf{z}, \\ \mathbf{lb} &= -\mathbf{ub}, \end{aligned} \quad (9)$$

where \mathbf{ub} is the vector of upper bounds and \mathbf{lb} is the vector of lower bounds for the confidence band. We accept the null hypothesis H_0^{GP} at the confidence level of $1 - \alpha$ if the value of $g(\mathbf{x}_{t_j})$ are within the band, that is,

$$(\mathbf{lb})_j \leq g(\mathbf{x}_{t_j}) \leq (\mathbf{ub})_j \quad \text{for all } j = 1 \dots n_{\text{test}},$$

where the notation $(\mathbf{a})_j$ is the j th component of a vector \mathbf{a} . Similarly, we reject the null hypothesis H_0^{GP} at the $1 - \alpha$ confidence level if there is at least one violation, that is,

there exists $j \in \{1 \dots n_{\text{test}}\}$ such that $g(\mathbf{x}_{t_j}) \notin [(\mathbf{lb})_j, (\mathbf{ub})_j]$.

It is also worth noting that we are getting an approximate band because of using a truncated KL expansion. This may result in losing some confidence on the test. In other words, the probability of the confidence band would be less than $1 - \alpha$. A possible compensation can be made by setting a slightly higher confidence level than the nominal level.

The points for which the null hypothesis is rejected would form a discrete grid on the difference region(s) and the absolute value of the statistically significant difference at these points would be given as: $\delta(\mathbf{x}_{t_j}) = |g(\mathbf{x}_{t_j})| - (\mathbf{ub})_j$. Needless to say that the difference at the points where the null hypothesis is not rejected would be considered zero.

2.4. Estimating the Hyperparameters

Until now, we have assumed the values of the hyperparameters of the covariance matrix and the nugget σ_ϵ are known. Next, we describe the method we use to estimate these hyperparameters.

Let us assume that $\boldsymbol{\theta}$ is the vector containing all the hyperparameters of the covariance function and the nugget σ_ϵ . We estimate these hyperparameters by merging the two datasets as $\mathbf{y} = (\mathbf{y}^{(1)}, \mathbf{y}^{(2)})$ and $\mathbf{X} = (\mathbf{X}^{(1)}, \mathbf{X}^{(2)})$, and jointly maximizing the likelihood as follows:

$$\hat{\boldsymbol{\theta}} = \arg \max \mathcal{L}(\boldsymbol{\theta}; \mathcal{D}_1, \mathcal{D}_2), \quad (10)$$

where $\mathcal{L}(\boldsymbol{\theta}; \mathcal{D}_1, \mathcal{D}_2) = (2\pi)^{-(n_1+n_2/2)} |\mathbf{K}_{\mathbf{X}, \mathbf{X}} + \sigma_\epsilon^2 \mathbf{I}|^{-1/2} \exp(-0.5(\mathbf{y}^\top [\mathbf{K}_{\mathbf{X}, \mathbf{X}} + \sigma_\epsilon^2 \mathbf{I}]^{-1} \mathbf{y}))$. We formally state the funGP algorithm in Section 2.5.

2.5. funGP Algorithm

Algorithm 1: funGP: function comparison using GP

Input: $\mathcal{D}_1 = \{\mathbf{X}^{(1)}, \mathbf{y}^{(1)}\}$, $\mathcal{D}_2 = \{\mathbf{X}^{(2)}, \mathbf{y}^{(2)}\}$, \mathbf{X}_{test} , α

Procedure:

- 1: Choose a covariance function.
- 2: Estimate the hyperparameters for the covariance function and the nugget, σ_ϵ , by optimizing the likelihood function given in Equation (10).
- 3: Compute the predictive mean functions \hat{f}_1 using \mathcal{D}_1 , and \hat{f}_2 using \mathcal{D}_2 using Equations (3) and (4).
- 4: Compute the covariance matrix $\mathbf{C}_{\mathbf{X}_{\text{test}}, \mathbf{X}_{\text{test}}}$ using the covariance function in Equation (6) for the points in \mathbf{X}_{test} .
- 5: Compute the difference between predictive means for the points in \mathbf{X}_{test} ,
 $g(\mathbf{x}_{t_j}) = \hat{f}_2(\mathbf{x}_{t_j}) - \hat{f}_1(\mathbf{x}_{t_j}) \mid j = 1, \dots, n_{\text{test}}$.
- 6: Do the eigen decomposition of $\mathbf{C}_{\mathbf{X}_{\text{test}}, \mathbf{X}_{\text{test}}}$ and store the m largest eigenvalues following the truncation rule in Section 2.3 in a diagonal matrix $\mathbf{\Lambda}$ and the corresponding eigenvectors in a matrix \mathbf{U} .
- 7: Compute the radius, r , of a standard normal vector of dimension m with a coverage probability of $1 - \alpha$ using Equation (7).
- 8: Sample a large number (say 1000) of standard normal vector \mathbf{z} such that $\|\mathbf{z}\| \leq r$.
- 9: Compute the vector of upper bounds, \mathbf{ub} , and lower bounds, \mathbf{lb} , for all the test points using Equation (9).

Output:

If $(\mathbf{lb})_j \leq g(\mathbf{x}_{t_j}) \leq (\mathbf{ub})_j \quad \forall j = 1, \dots, n_{\text{test}}$, functions are same at $1 - \alpha$ confidence level.

Else, functions are different at $1 - \alpha$ confidence level.

3. Simulation Study

In this section, we present three simulation studies for the funGP method to estimate the Type I and the Type II errors and compare it with two methods from the existing literature. We estimate the Type II error for some small perturbations. In order to quantify the difference between a function and its perturbation, we use an L^2 -distance percentage defined as follows:

$$L^2 \text{ dist \%} = \frac{\|f - g\|_{L^2}}{\|f\|_{L^2}} \times 100\%,$$

where f is the underlying function, and g is its perturbation. After fixing the nominal level of H_0^{GP} to $\alpha = 0.05$, we conduct 1000 runs for each simulation example to estimate the Type I/Type II errors. We also examine the effect of the sample size, the number of points in the test grid, and the truncation number in the KL expansion on the Type I and Type II errors using different experiments.

3.1. Functions Used in the Simulations

The first simulation study is based on functions sampled from a known GP. Our method also assumes the functions to be GP samples, so this study represents a case when there is no model misspecification, that is, the GP part in H_0^{GP} is indeed

true. Whereas the other two simulation studies are based on some parametric functions available in the literature; we use GP as a surrogate for the true function. Thus, for these two fixed functions, the estimated Type I/Type II errors are for the original hypothesis test H_0 , even though we are controlling the Type I error only under H_0^{GP} . Hence, the last two simulations evaluate the efficacy of the funGP method under a potential model misspecification. In order to generate the datasets, we randomly sample two sets of points from the input domain of the functions. We then generate response by adding some iid Gaussian noise to the function values at the sampled input points. For conducting all the simulation studies, we use a constant mean and a squared exponential covariance function for the GP modeling.

The following are the specifications of the simulated functions. We consider a one-dimensional input $x \in [0, 1]$ for the first simulation study, the GP sample. The model can be described as: $y = f(x) + \epsilon$; $f(x) \sim \mathcal{GP}(0, k(x, x'))$; $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$. The covariance function $k(x, x')$ is squared exponential with the following form: $k(x, x') = \sigma_f^2 \exp(-0.5[(x - x')/\theta]^2)$. The hyperparameters for the covariance function, $k(x, x')$, are set to $\sigma_f = 5$, and $\theta = 0.2$. The standard deviation of the noise, σ_ϵ , is set to 0.5. For each simulation run, a different sample is generated from the given GP model, and the estimated Type I error is the percentage of runs for which the null hypothesis is rejected. For estimating the Type II error, we create a perturbation $g(x)$ in the following way:

$$g(x) = \begin{cases} f(x) + \frac{1}{3} \sin\left(\pi \left(\frac{x-0.2}{0.8-0.2}\right)\right), & x \in [0.2, 0.8], \\ f(x), & \text{otherwise.} \end{cases}$$

The functions $f(x)$ and $g(x)$ sampled for one simulation run (left panel) along with two noisy datasets generated from it (right panel) are shown in Figure 1. One can see that the difference between the functions is small and gets masked visually in the noisy data.

For the other two studies, we use two parametric functions available in the literature: piston simulation function (Kenett and Zacks 1998) and borehole simulation function (Harper and Gupta 1983). We use these functions with two dimensional input by fixing the rest of their input variables to certain values. More details about these functions, including the function plots, are provided as supplementary material S.3 to maintain the flow for the readers and save space.

3.2. Results

Table 1 shows the estimated Type I and Type II errors for all the simulation studies, along with the L^2 distance between the function f and its perturbation, g . The results in Table 1 are based on the following specifications. The sample size for one-dimensional function (GP sample) and two-dimensional functions (piston and borehole) are 500 and 1000, respectively, randomly sampled from their respective input domain. The test grid is 500 evenly spaced points in the domain for the GP sample and 50×50 evenly spaced grid for the piston and the borehole functions. We follow the same truncation rule for the truncation number m as described in Section 2.3. The L^2 distance varies between 3% and 5 %.

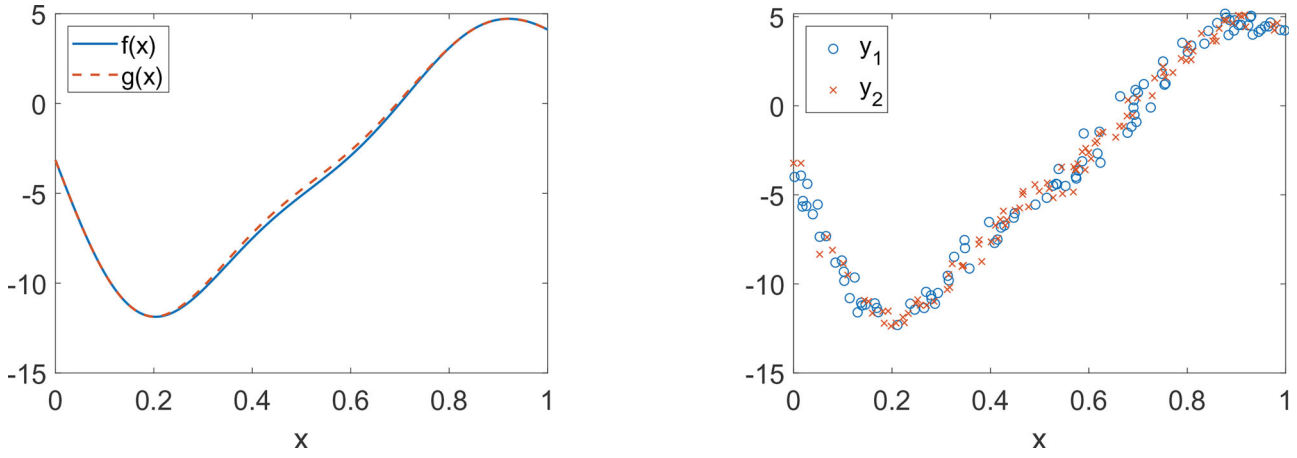


Figure 1. Plots for GP sample. Left panel: $f(x)$ and its perturbation, $g(x)$; Right panel: Noisy realizations from $f(x)$ and $g(x)$.

Table 1. Estimated Type I and Type II errors for the simulated functions.

Function	Type I error	Type II error	L^2 dist %
GP sample	0.049	0.031	4.7
Piston	0.041	0.008	3.8
Borehole	0.065	0.022	3.4

For the first case study, when the true functions are GP samples, the estimated Type I error is very close to the nominal level of 0.05 (5%). This result is very much expected, as both we are controlling the Type I error under H_0^{GP} and the GP assumption is indeed true. In the other two simulation studies, the estimated Type I error is not as close to the nominal value as the first simulation study. This can be attributed to the fact that the estimated Type I error is for H_0 , and we are controlling the Type I error for H_0^{GP} . The form of the mean and covariance function required to sample these functions from a GP is not known and we use approximations in these studies. The agreement between the estimated Type I error and the nominal value would depend on how well the GP approximates the function. If it is difficult to approximate a function using a known parametric covariance function, then we can either come up with more sophisticated mean and covariance functions, or we can increase the confidence level of the test to a value greater than the desired level to account for model uncertainty. For the given sample size, we are satisfied that the method can identify the difference in the underlying functions even with small perturbations. We would, next, conduct experiments to see how the method performs under different sample sizes, test grid sizes and truncation numbers.

3.3. Further Experiments

We repeat the three simulation experiments carried out previously under different sample sizes while keeping the test grid and the truncation rule fixed. The sample sizes are set at four levels: 100, 200, 500, and 1000. Table 2 presents the results of these experiments. The table clearly shows a reduction in Type II error as the sample size increases while keeping the Type I error stable, which is consistent with our understanding of statistical hypothesis tests. The numerical results indicate that in order

Table 2. Estimated Type I and Type II errors under different sample sizes.

Function	Estimate	Sample size for each dataset			
		100	200	500	1000
GP sample	Type I error	0.032	0.040	0.049	0.049
	Type II error	0.721	0.450	0.031	0.001
piston	Type I error	0.031	0.032	0.027	0.041
	Type II error	0.845	0.639	0.201	0.008
borehole	Type I error	0.058	0.058	0.070	0.065
	Type II error	0.757	0.460	0.091	0.022

Table 3. Estimated Type I and Type II errors under different test grid sizes.

Function	Estimate	Number of test points			
		100	400	900	2500
GP sample	Type I error	0.039	0.038	0.059	0.049
	Type II error	0.033	0.033	0.034	0.022
piston	Type I error	0.025	0.032	0.044	0.041
	Type II error	0.020	0.014	0.014	0.008
borehole	Type I error	0.035	0.050	0.046	0.065
	Type II error	0.063	0.027	0.024	0.022

to render sufficient detection power, a large enough sample is needed for detecting small difference between two functions.

The test grid size experiment is carried out while keeping the sample size fixed at the same value as used for the main result in Section 3.2. We use a test grid of size 100, 400, 900, and 2,500 for each simulated function so that it corresponds to 10×10 , 20×20 , 30×30 , and 50×50 test grid, respectively, for two-dimensional functions (piston and borehole). Table 3 presents the results of this experiment. There is no significant effect of the test grid size on the Type I and Type II errors. This is expected as our test relies on the truncated KL expansion and the number of eigenvalues (m) remains constant for different test grid sizes because we use the same truncation rule.

We contemplate how the results may change when we use any arbitrary truncation number instead of using the aforementioned rule for calculating the truncation number. Table 4 displays the result of using different truncation numbers on the Type I and Type II errors. We note that the hypothesis test remains a level- α test as long as the truncation number is

Table 4. Estimated Type I and Type II errors under different truncation numbers.

Function	Estimate	Truncation number		
		10	50	100
GP sample	Type I error	0.051	0.034	0.029
	Type II error	0.023	0.043	0.048
piston	Type I error	0.036	0.013	0.024
	Type II error	0.016	0.020	0.026
borehole	Type I error	0.260	0.046	0.039
	Type II error	0.014	0.020	0.018

larger than a certain threshold, of which the specific value would depend on the function under study. For the first two cases, a truncation number of 10 or greater appears sufficient, whereas for the third case, a truncation number may need to be as large as 50. When one chooses a smaller truncation number than the problem demands, then one cuts off a significant portion of the $1 - \alpha$ confidence band, resulting in a high Type I error. Thus, we suggest using the recommended truncation rule, which adapts the truncation number according to the problem.

3.4. Comparison With Other Methods

We compare our method with two other methods available in the literature. The first comparison is with Munk and Dette (1998), which is a global test that works for datasets without requiring common input points and replicates. Although it is a global test, we can still use this method to compare with the funGP method in terms of the Type I and Type II errors. This method builds its test statistic based on the L^2 -distance between the functions. The method is developed for functions with one-dimensional input. For this reason, we use this method only for the first simulation study, the GP sample.

Table 5 presents the results for the comparison. We note that our method is significantly more powerful than Munk and Dette's method. Munk and Dette (1998) provided an expression for approximating the power of their test, given the L^2 -distance between the functions, the sample size, and the noise level; see Equation (17) in Munk and Dette (1998). The approximate power computed using that expression is 0.435, which is consistent with the empirically estimated Type II error in Table 5 (power = $1 - \text{Type II error}$).

We also compare our method with Cox and Lee (2008), which identifies the difference region in terms of p -values. Cox and Lee (2008) is based on a permutation test and requires the datasets to have replicates and the same input points. Since the datasets simulated for the funGP method do not have replicates and do not share the same input points, we simulate different sets of samples with replicates keeping the input points the same for the two functions. We apply both funGP and Cox and Lee's methods to these newly generated datasets to estimate

Table 5. Comparison between the funGP and Munk and Dette (1998) methods for the GP sample simulation study.

Function	Method	Type I error	Type II error
GP sample	funGP	0.049	0.031
	Munk & Dette	0.117	0.570

Table 6. Comparison between funGP and Cox and Lee methods.

Function	Method	Type I error	Type II error
GP sample	funGP	0.039	0.042
	Cox & Lee	0.023	0.160
piston	funGP	0.027	0.011
	Cox & Lee	0.017	0.071
borehole	funGP	0.042	0.098
	Cox & Lee	0.020	0.017

the Type I and Type II errors. We still use 1000 runs for the simulation. We use 50 input points with 10 replications each for 1-dimensional case (GP sample function) and 100 input points with 10 replications each for two-dimensional cases (piston and borehole functions). The sample sizes are chosen such that the total number of the samples is equal to that of the main simulation study, that is, 500 for one-dimensional case and 1000 for two-dimensional case. In each case, the nominal level of the test is set to $\alpha = 0.05$.

The results for this comparison are presented in Table 6. The proposed funGP method performs better than Cox and Lee in two out of three cases—the GP sample and the piston cases, and worse for the borehole case, in terms of the Type II error. We would like to articulate that we advocate the merit of our method as identifying the difference region and quantifying the difference for datasets arising from a broader setting, namely without the same input points and replicates, and not purely in terms of its power in a binary decision. Yet, our method performs comparably, and sometimes even better, than other methods such as Cox and Lee (2008) and Munk and Dette (1998).

4. Application

In this section, we apply the funGP method to a wind energy problem. A common technique to characterize the performance of a wind turbine is through the use of its power curve (Ding 2019, chap. 5 and 6). A univariate power curve is a functional curve with wind speed as the input and the generated wind power as the output. But researchers realize that the wind power output is affected by other inputs more than just the wind speed. Consequently, multivariate power curves have been developed; see, for instance Ding (2019, chap. 5) or Lee et al. (2015).

A nominal wind power curve is shown in Figure 2. The turbine does not produce power below the cut-in wind speed V_{ci} . Above the cut-in speed, the power gradually rises till the rated power and then capped at that level till the cut-out wind speed V_{co} , at which the turbine operation is stopped in order to protect its components against damage. The pitch control is one of the main mechanisms to regulate a wind turbine's power output (Senjyu et al. 2006); in Figure 2, we mark the wind speed region where the pitch control is active.

The power curve (univariate or multivariate) is generally learned through data; please see Chapter 5 of Ding (2019) for various methods. If one wants to compare the performance of two turbines or the same turbine over multiple time periods, then they can do so by comparing the learned power curves. This raises a question that whether the difference in the learned

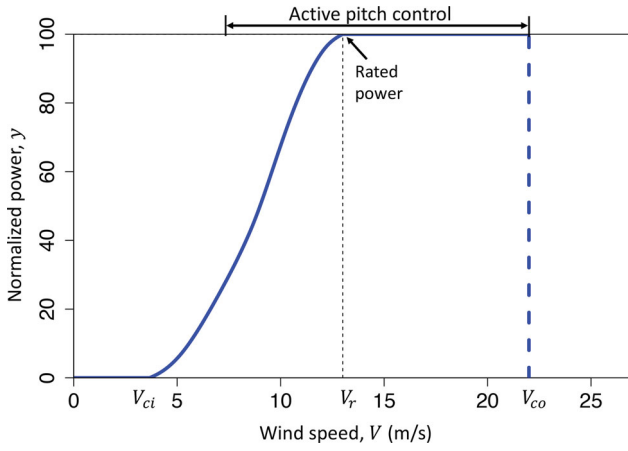


Figure 2. A nominal wind power curve. V_{ci} : the cut-in wind speed; V_r : the rated wind speed; V_{co} : the cut-out wind speed.

curves is due to the randomness in the samples, or the difference is genuine in turbine performance beyond random fluctuation. Our proposed method can, hence, be employed to answer this question.

We apply our method to the four datasets as used by Hwangbo, Johnson, and Ding (2017), which also constitutes a large portion of (Ding 2019, chap. 6), and we download the four datasets from the book website of Ding (2019) (<https://aml.engr.tamu.edu/book-dswe/dswe-datasets>, *Inland and Offshore Wind Farm Dataset2*). Each dataset corresponds to a different turbine. The four turbines are labeled as WT1, WT2, WT3, and WT4. The datasets WT1 and WT2 are from onshore turbines and have the following five input variables: wind speed (V), wind direction (D), air density (ρ), turbulence intensity (I), and wind shear (S). The other two datasets (WT3 and WT4) correspond to offshore wind turbines with the input variable S replaced with humidity (H), with the rest of the variables the same as that of the onshore turbines. Each of the four datasets comprises four years of data. We conduct a year to year comparison for each turbine, as done in Hwangbo, Johnson, and Ding (2017). For this reason, each turbine's dataset is divided into four annual datasets.

The marginal distributions of the covariates are different for each year, thus before computing their metric, Hwangbo, Johnson, and Ding (2017) applied a method called covariate matching to the annual datasets. Covariate matching tries to match the marginal distributions of all the available environmental variables among the annual datasets by selecting the proper data subsets. Covariate matching is applied here in order to enable a fair comparison in turbine performance by ensuring that the distributions of the environmental variables are similar. We follow the same strategy with the same specifications as given in Hwangbo, Johnson, and Ding (2017). After the covariate matching, Hwangbo, Johnson, and Ding (2017) used only the wind speed as the input variable to estimate the power curve. We also proceed in a similar way. In other words, we have wind speed as the input and wind power as the output. We input these datasets to our funGP algorithm and do a pairwise comparison between the annual datasets for each turbine using the following specification. We select 1000 evenly spaced points from the range of the input variable (wind speed) as the test grid and compare the power curves for any two annual datasets for

Table 7. Percentage of test points with statistically significant difference between annual datasets.

Turbine	Year 1 & 2	Year 1 & 3	Year 1 & 4	Year 2 & 3	Year 2 & 4	Year 3 & 4
WT1	49.5	58.1	53.6	13.9	0	0
WT2	40.6	41.3	41.3	0	0	0
WT3	85.6	81.4	73.1	55.4	72.7	41.9
WT4	74.9	60.8	64.3	44.4	69.6	2.7

a given turbine on the defined test grid. A typical wind turbine operates at wind speeds between 5 m/s to 15 m/s for most of the time. Thus, we select this range to test the difference. Hwangbo, Johnson, and Ding (2017) developed a 90 % confidence interval for their performance metric using the bootstrap method. For comparison, we also build a 90 % confidence band on the difference of the power curves.

The outputs from our method are the pointwise difference in the power curves and the 90 % confidence band on the difference for the power curves to be the same. In Table 7, we report the percentage of points, out of the 1000 test points, where the difference between two given yearly datasets is statistically significant. Whenever the percentage is greater than zero, we claim that the difference between corresponding two curves is statistically significant.

Speaking of the current industry practice for turbine performance comparison in the wind energy sector, the most popular method is to compare their peak power coefficient estimated from the data (IEC 2005). The power coefficient, C_p , of a turbine is computed by using the following formula:

$$C_p = \frac{2y}{\rho AV^3},$$

where y is the wind power output and A is the sweeping area of the turbine blades. Here, C_p is not a constant but rather a function of wind speed and a few other factors. The exact formula linking C_p to other physical variables does not exist. So it is empirically estimated. Using a functional C_p is not easy, and because of that, practitioners simply choose the peak value on the C_p -versus-wind-speed curve to represent the performance of a turbine. The power coefficient has a theoretical upper bound, known as the Betz limit, which is 0.593 (Ding 2019) but the practical C_p is generally smaller than 0.5. It is obvious that this C_p metric is just a point metric of an otherwise functional difference.

Hwangbo, Johnson, and Ding (2017) suggested another technique to compare the performance of wind turbines using the concepts of production economics. They devise a performance metric called productive efficiency which takes into account the overall power curve and not just the peak performance. But their final output is again a point metric of the functional difference, much like the power coefficient. Hwangbo, Johnson, and Ding (2017)'s study found the productive efficiency metric has a good similarity with the power coefficient metric, although not exactly the same. Using the four datasets mentioned above, the performance quantifications using the two metrics registered a correlation of 0.75 (Hwangbo, Johnson, and Ding 2017). Other than being a point metric, both the power coefficient and the productive efficiency methods do not quantify the estimation uncertainty on their own—one can go through an expensive

bootstrap approach to get a confidence interval on the performance metrics. The funGP method, on the other hand, can lead to any level of confidence bands on the difference of the performance.

We compare our results with the metrics, peak power coefficient and productive efficiency, obtained by Hwangbo, Johnson, and Ding (2017, Table II). We illustrate the comparison in a chart (see Figure 3) using vertical and horizontal lines with the following criteria:

- If the two metrics used by Hwangbo, Johnson, and Ding (2017) agreed with each other (i.e., they both say the two annual periods are different or they both say the same), and they also agree with our result, then we use vertical lines to demonstrate that.
- If the two metrics do not agree with each other, but one of them agree with our result, then we still use vertical lines.
- However, if the two metrics agree with each other, but they do not agree with our method, then we use horizontal lines to show that.

In other words, the vertical lines imply an agreement between our method and at least one of the two metrics, whereas the horizontal lines mean a disagreement between the two metrics and the funGP method. We observe that when the difference between two power curves is statistically significant, the confidence intervals of the peak power coefficient or the productive efficiency for the same two curves tend not to overlap, leading naturally to the overwhelming agreement pattern observed in Figure 3.

There is one comparison outcome for which using funGP and either metric in Hwangbo, Johnson, and Ding (2017) disagree: WT1 for Year 2 versus Year 3. Taking a closer look reveals that the percentages of test points where the two curves are different, as reported in Table 7, is 13.9%. The percentage is much smaller than the percentage values in other cases for which two curves are declared different. When we look at the power coefficient and productive efficiency values in Hwangbo, Johnson, and Ding (2017, Table II), they are as such:

- WT 1's power coefficient. Year 2: 0.388 with the 90% confidence intervals as [0.386, 0.392], and Year 3: 0.393 with the 90% confidence intervals as [0.390, 0.397].

- WT1's productive efficiency. Year 2: 0.969 with the 90% confidence intervals as [0.966, 0.973], and Year 3: 0.972 with the 90% confidence intervals as [0.969, 0.975].

Apparently, for the power coefficient and productive efficiency metrics, their 90% confidence intervals are only marginally overlapping, not really contradicting with the small regions of difference detected by using the funGP method. It is not unreasonable to consider that the funGP method is more sensitive to the difference between the two curves.

The funGP method provides a quantification of the regions of difference. Better yet, funGP can be used to compute the difference in the power curves at any point in the domain of the curve, and thus, gives a more detailed picture of the difference between any two curves, so that the practitioners can see where the difference lies and thus make an informed decision regarding whether the difference region matters or not. Figure 4 shows this difference vs. wind speed plot for all the annual datasets for the first turbine (WT1).

As described in Section 1, knowing the regions of difference is helpful in deciding the maintenance plan for the turbine. For instance, if the difference occur in low power range, one may not necessarily need to go for expensive maintenance as doing so is unlikely to result in large change in the power output. Another important implication of knowing the difference regions is to decide the pitch control configuration of the turbine. As Creaby, Li, and Seem (2009) explained, wind turbine's aerodynamic characteristics change with time because of surface wear, dirt and other factors. Therefore, knowing the region of difference can help adjust the control laws to optimize the pitch control for different regions of operations, in order to maximize the power output. The funGP method is better suited in this application as a more powerful and informative testing and comparison method.

5. Discussions

This work presents a new nonparametric method that compares functions, referred to as the funGP method. Unlike many methods in the literature, the novelty of funGP lies in its ability to identify the regions of difference in the input space of the functions and quantify this difference, rather than simply

Turbine	Year 1 & 2	Year 1 & 3	Year 1 & 4	Year 2 & 3	Year 2 & 4	Year 3 & 4
WT1	Vertical lines	Vertical lines	Vertical lines	Horizontal lines	Vertical lines	Vertical lines
WT2	Vertical lines	Vertical lines	Vertical lines	Vertical lines	Vertical lines	Vertical lines
WT3	Vertical lines	Vertical lines	Vertical lines	Vertical lines	Vertical lines	Vertical lines
WT4	Vertical lines	Vertical lines	Vertical lines	Vertical lines	Vertical lines	Vertical lines

Figure 3. Comparison chart for the results obtained using funGP method to that of the peak power coefficient and the productive efficiency method. Vertical lines imply that the results agree. Horizontal lines imply that the results differ.

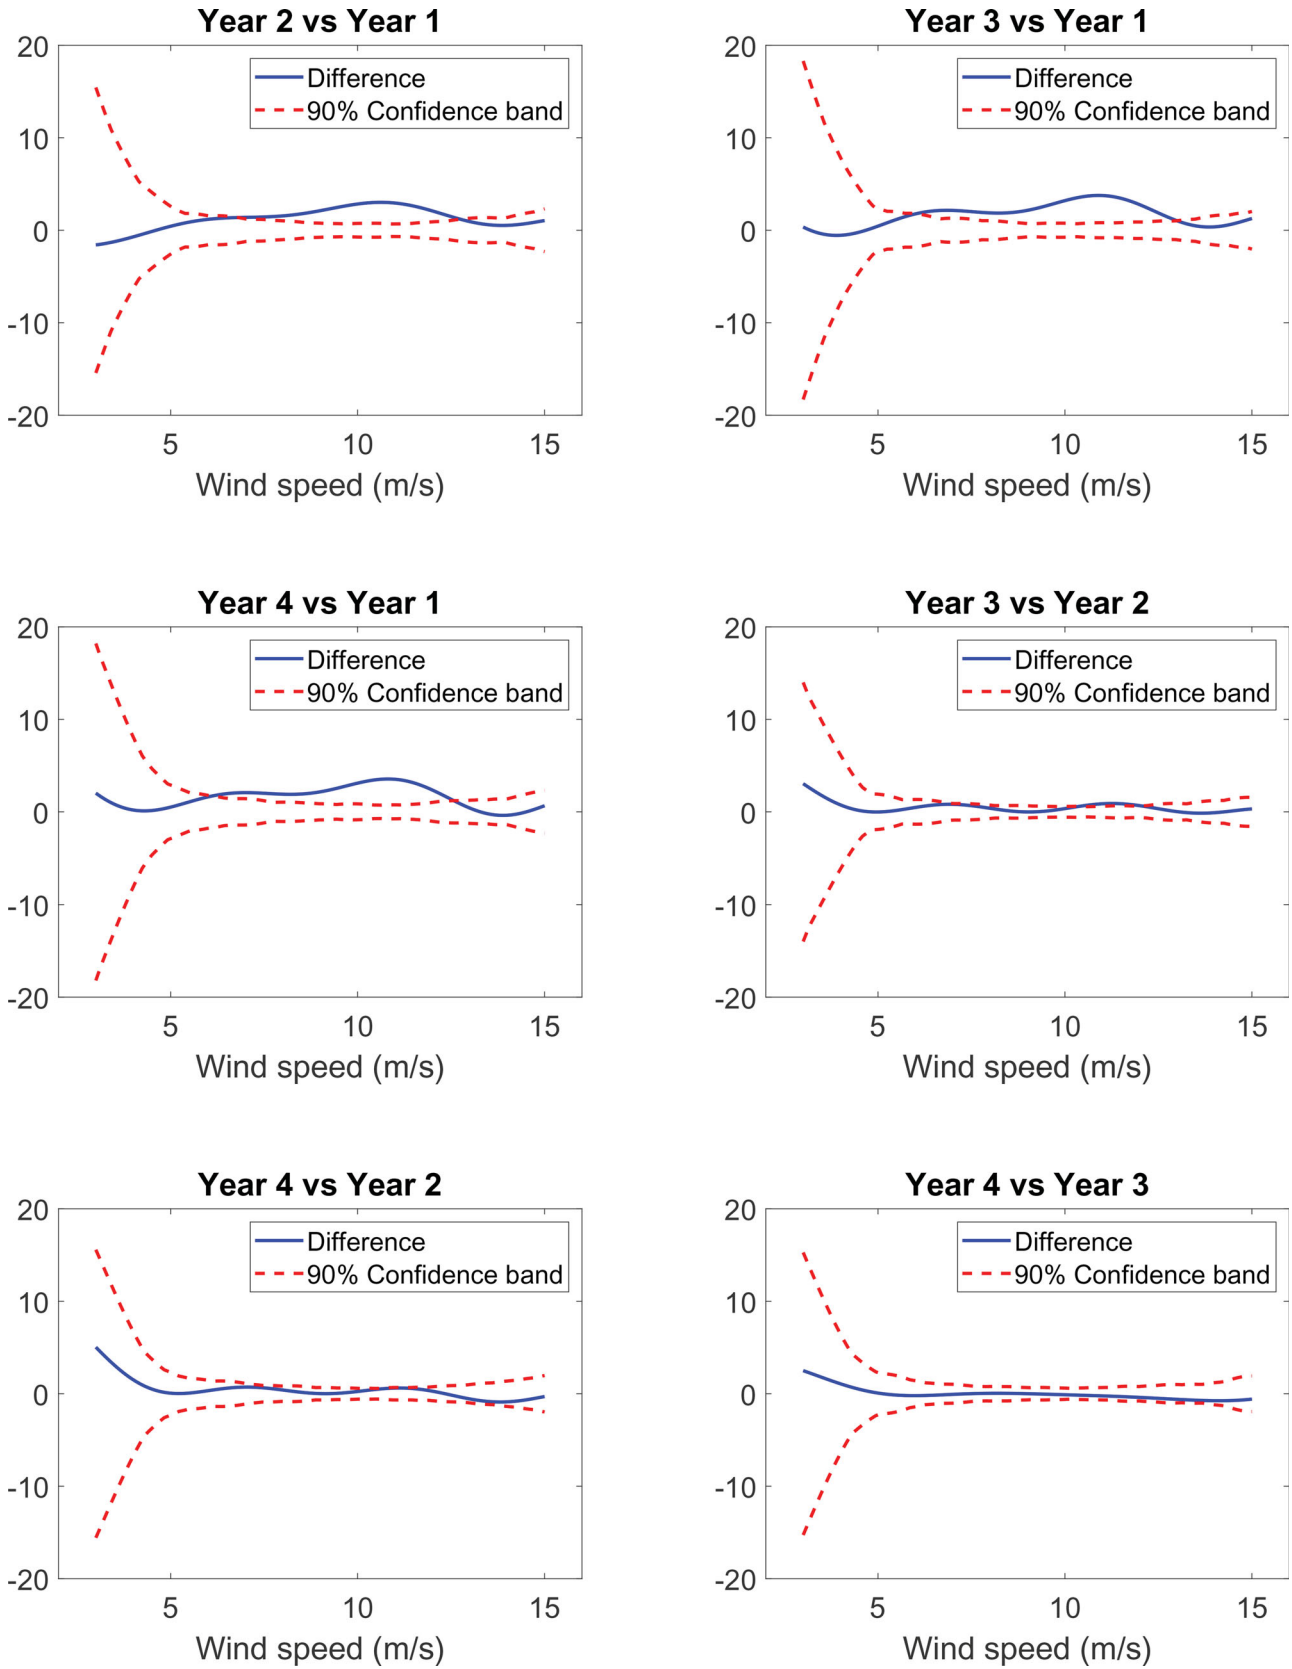


Figure 4. Difference in curves vs wind speed for WT1.

returning a binary answer on whether the difference exists or not. This ability makes the funGP method a truly functional test.

From an application point of view, particularly in engineering, comparing processes often mean subsequent decision making. For instance, comparing wind power curves guides

the maintenance strategy. Under these circumstances, a binary answer for function comparison can easily run to its limit, as it may not be of much help in driving the decision-making process. Understanding a fuller picture of function difference through region identification and subsequent quantification, on the other hand, could lead to better engineering and economic decisions. We hope that our work paves the way and shifts the focus of function comparison research toward more informative function tests, which would have broader applications and impact in the engineering fields.

In the work, we use evenly spaced input points to conduct the comparison of the curves. In the higher dimensions, the numbers of grid points can grow rapidly and may become computationally burdensome. One worthy future research direction that would directly advance this work is to devise an adaptive grid—based on the characteristics of the function under study—to quickly identify and quantify the differences while reducing the computational time.

Supplementary Material

Supplementary Material: The PDF file contains: (i) Derivation of the covariance function $c(\mathbf{x}, \mathbf{x}')$, (ii) Brief description of the Karhunen-Loève expansion of a Gaussian process, and (iii) Details on borehole and piston functions.

Computer Code: The computer code to reproduce all the results in this article are available on GitHub at <https://github.com/TAMU-AML/funGP-Paper>. A generic R function for applying the funGP algorithm to any dataset is available in DSWE package in R available through CRAN at <https://CRAN.R-project.org/package=DSWE>.

Acknowledgments

The authors thank the Editor, the Associate Editor, and the Reviewers for providing valuable feedback. Their comments has led to a significant improvement in the article. The authors would also like to acknowledge the role of Texas A&M's high-performance research computing (HPRC), which enabled the authors to efficiently run their experiments.

Funding

Prakash and Ding's research is partially supported by NSF grant IIS-1741173. Tuo's research is supported by NSF DMS-1914636. Ding and Tuo's research is also supported by NSF grant CCF-1934904.

References

- Cox, D., and Lee, J. S. (2008), "Pointwise Testing With Functional Data Using the Westfall-Young Randomization Method," *Biometrika*, 95, 621–634. [2,7]
- Creaby, J., Li, Y., and Seem, J. E. (2009), "Maximizing Wind Turbine Energy Capture Using Multivariable Extremum Seeking Control," *Wind Engineering*, 33, 361–387. [9]
- Delgado, M. A. (1993), "Testing the Equality of Nonparametric Regression Curves," *Statistics & Probability Letters*, 17, 199–204. [1]
- Ding, Y. (2019), *Data Science for Wind Energy*, Boca Raton, FL: Chapman & Hall/CRC Press. [1,7,8]
- Fan, J., and Lin, S.-K. (1998), "Test of Significance When Data Are Curves," *Journal of the American Statistical Association*, 93, 1007–1021. [1]
- Fisher, R. A. (1925), "Application of 'Student's' Distribution," *Metron*, 5, 90–104. [1]
- Hall, P., and Hart, J. D. (1990), "Bootstrap Test for Difference Between Means in Nonparametric Regression," *Journal of the American Statistical Association*, 85, 1039–1049. [1]
- Harper, W., and Gupta, S. (1983), *Sensitivity/Uncertainty Analysis of a Borehole Scenario Comparing Latin Hypercube Sampling and Deterministic Sensitivity Approaches*. BMI/ONWI-516, Columbus, OH: Office of Nuclear Waste Isolation, Battelle Memorial Institute. [5]
- Hotelling, H. (1931), "The Generalization of Student's Ratio," *The Annals of Mathematical Statistics*, 2, 360–378. [1]
- Hwangbo, H., Johnson, A., and Ding, Y. (2017), "A Production Economics Analysis for Quantifying the Efficiency of Wind Turbines," *Wind Energy*, 20, 1501–1513. [1,8,9]
- IEC. (2005). *Wind Turbines-Part 12-1: Power Performance Measurements of Electricity Producing Wind Turbines*, Geneva, Switzerland: International Electrotechnical Commission 61400-12-1 Ed. 1. [1,8]
- Kenett, R. S., and Zacks, S. (1998), *Modern Industrial Statistics: The Design and Control of Quality and Reliability*, Pacific Grove, CA: Duxbury Press. [5]
- King, E., Hart, J. D., and Wehrly, T. E. (1991), "Testing the Equality of Two Regression Curves Using Linear Smoothers," *Statistics & Probability Letters*, 12, 239–247. [1]
- Kulasekera, K. B. (1995), "Comparison of Regression Curves Using Quasi-residuals," *Journal of the American Statistical Association*, 90, 1085–1093. [1]
- Kulasekera, K. B. and Wang, J. (1997), "Smoothing Parameter Selection for Power Optimality in Testing of Regression Curves," *Journal of the American Statistical Association*, 92, 500–511. [1]
- Lee, G., Ding, Y., Genton, M. G., and Xie, L. (2015), "Power Curve Estimation With Multivariate Environmental Factors for Inland and Offshore Wind Farms," *Journal of the American Statistical Association*, 110, 56–67. [7]
- Munk, A., and Dette, H. (1998), "Nonparametric Comparison of Several Regression Functions: Exact and Asymptotic Theory," *The Annals of Statistics*, 26, 2339–2368. [1,2,7]
- Neumeyer, N., and Dette, H. (2003), "Nonparametric Comparison of Regression Curves: An Empirical Process Approach," *The Annals of Statistics*, 31, 880–920. [1]
- Rasmussen, C. E., and Williams, C. K. I. (2006), *Gaussian Processes for Machine Learning*, Cambridge, MA: The MIT Press. [3]
- Senjyu, T., Sakamoto, R., Urasaki, N., Funabashi, T., Fujita, H., and Sekine, H. (2006), "Output Power Leveling of Wind Turbine Generator for All Operating Regions by Pitch Angle Control," *IEEE Transactions on Energy conversion*, 21, 467–475. [7]
- Westfall, P. H., and Young, S. S. (1993), *Resampling-Based Multiple Testing: Examples and Methods for p-Value Adjustment*. New York: Wiley. [2]