



On switching response surface models, with applications to the structural health monitoring of bridges



K. Worden^{*}, E.J. Cross

Dynamics Research Group, Department of Mechanical Engineering, University of Sheffield, Mappin Street, Sheffield S1 3JD, United Kingdom

ARTICLE INFO

Article history:

Received 13 October 2016

Received in revised form 13 March 2017

Accepted 18 April 2017

Available online 5 May 2017

Keywords:

Structural Health Monitoring (SHM)

Environmental and operational variations

Confounding influences

Response surfaces

Switching models

Gaussian processes

ABSTRACT

Structural Health Monitoring (SHM) is the engineering discipline of diagnosing damage and estimating safe remaining life for structures and systems. Often, SHM is accomplished by detecting changes in measured quantities from the structure of interest; if there are no competing explanations for the changes, one infers that they are the result of damage. If the structure of interest is subject to changes in its environmental or operational conditions, one must understand the effects of these changes in order that one does not falsely claim that damage has occurred when changes in measured quantities are observed. This problem – the problem of *confounding influences* – is particularly pressing for civil infrastructure where the given structure is usually openly exposed to the weather and may be subject to strongly varying operational conditions. One approach to understanding confounding influences is to construct a data-based *response surface model* that can represent measurement variations as a function of environmental and operational variables. The models can then be used to remove environmental and operational variations so that change detection algorithms signal the occurrence of damage alone. The current paper is concerned with such response surface models in the case of SHM of bridges. In particular, classes of response surface models that can *switch* discontinuously between regimes are discussed.

Recently, it has been shown that Gaussian Process (GP) models are an effective means of developing response surface or *surrogate* models. However, the GP approach runs into difficulties if changes in the latent variables cause the structure of interest to abruptly switch between regimes. A good example here, which is well known in the SHM literature, is given by the Z24 Bridge in Switzerland which completely changed its dynamical behaviour when it cooled below zero degrees Celsius as the asphalt of the deck stiffened. The solution proposed here is to adopt the recently-proposed Treed Gaussian Process (TGP) model as an alternative. The approach is illustrated here on the Z24 bridge and also on data from the Tamar Bridge in the UK which shows marked switching behaviour in certain of its dynamical characteristics when its ambient wind conditions change. It is shown that treed GPs provide an effective approach to response surface modelling and that in the Tamar case, a linear model is in fact sufficient to solve the problem.

© 2017 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

^{*} Corresponding author.

E-mail address: k.worden@sheffield.ac.uk (K. Worden).

1. Introduction

In very brief terms, Structural Health Monitoring (SHM) is the engineering discipline concerned with inferring the state of health of a structure or system from measurements obtained from sensors permanently installed on the structure or within the system [1]. It is possible to exploit a very diverse range of sensor technologies in the implementation of an SHM system, but one of the more common choices is to monitor dynamical response using accelerometers. This choice leads to vibration-based SHM, and this is the main choice considered in this paper.

It is critical to note that an SHM system is much more than a sensor network. It is almost always the case that the information about the health of the structure is well hidden in the raw time series data acquired by sensing. This issue arises because small incipient damage will not usually cause a major departure from the dynamical behaviour of the healthy structure. Because of this fact, the vital ingredient in any SHM system is an inference engine which constructs low-dimensional data vectors called *features* in which the effect of damage is much more visible. An example of a damage-sensitive feature vector often used in vibration-based SHM would be a set of the natural frequencies or resonance frequencies of the structure of interest. Natural and resonance frequencies are functions of the structural stiffness and will (usually) decrease when damage – such as a fatigue crack – causes a local reduction in stiffness. Determining natural frequencies from the raw time data is one example of *feature extraction* as it is referred to in the context of pattern recognition or machine learning [1]. Once damage-sensitive features have been determined, the SHM inference engine can proceed to an analysis which provides diagnostic and prognostic information about the health of the structure.

One of the major problems associated with SHM, is that features may change as a result of mechanisms other than damage and one does not usually wish to raise an alarm as a result of these benign changes. These other influences on the features will be referred to here as *confounding influences*; they most often arise in the context of engineering as the result of changes in the environment or operating conditions of the structure of interest. For the bridges discussed in this paper, ambient temperature is an environmental variable which strongly affects the SHM features, while traffic loading is an equally important operational influence. If natural frequencies are to be used as features for SHM, it has long been known that variations in the frequencies due to temperature changes can mask variations due to damage [2]. In order to implement damage detection by detecting changes in features, one must clearly produce features that are sensitive to damage but insensitive to environmental and operational variations, or alternatively, one must project out from the features the influence of the benign variations. This process is commonly referred to in the SHM literature as *data normalisation*; various techniques can be applied and a good, fairly recent, survey of the field can be found in [3].

Among the techniques available for data normalisation, one of the simplest is a regression-based approach. This relies on the availability of measurements of the environmental or operational variables of interest. When the features for SHM are based on the dynamics of the structure – as in vibration-based SHM – the response variables almost always change on a much shorter time-scale than the confounding influences. For example, accelerations measured on a bridge will have frequencies associated with tens of Hertz, while cycles of variation associated with temperature or traffic will be on scales of hours or more. This means that time histories acquired over hours or days will show their main variation as a result of the confounding influences, with the dynamical behaviour superimposed as a form of high-frequency ‘noise’. Fitting a regression model to such data with the environmental or operational variables as the independent variables will then capture only the dependence on the confounding influences, predictions from this model can then be subtracted from subsequent data, with the remaining residual (hopefully) sensitive *only* to damage. Regression models used in this context are often called *response surface models* and can vary in sophistication from simple polynomials [4], to state-of-the-art structures derived from modern machine learning theory like artificial neural networks and support vector machines [5,6]; examples from both ends of the spectrum will be presented in this paper. Complications can arise if the confounding influences cause discontinuous changes in the features as the ambient variables change, for example if polynomial models are selected, discontinuous behaviour may force the choice of very high-order polynomials with the result that very many coefficients need to be estimated. If the response surface models have the capability to *switch* between simple (e.g. linear) submodels, the number of parameters for estimation from the data can be much smaller, such models are often referred to as *parsimonious*. Parsimonious models are always selected where possible as they require less training data for their estimation problem, and data from structures in engineering SHM problems, particularly data corresponding to damage states, can sometimes be in short supply. In the machine learning context, parsimonious models are desirable because they are less prone to overfitting [7].

As mentioned above, when nonlinear models are required, there are numerous options for the model structure. The structure chosen here is the *Gaussian Process* (GP) [8]; this represents a powerful nonparametric regression technique which has been developed considerably within the machine learning community in the last 10–15 years. Advantages of the GP approach include a natural Bayesian framework for analysis and the automatic availability of confidence intervals for model predictions. In fact, Gaussian processes have a pedigree in terms of response surface modelling and sensitivity analysis [9], and the current authors have exploited this for their previous studies on engineering problems [10,11].

The approach to data normalisation discussed above and in the remainder of the paper can be referred to as a *subtraction* strategy. Another powerful approach can be based on the idea of *projection*, whereby the subspace of the feature space containing the confounding influences is identified and the features are projected onto the orthogonal complement of the corrupted subspace. The projection approach has various merits, including the property that one does not require measurements of the latent variables driving the confounding influences. The projection approach is not discussed further

in this paper; however, the curious reader may consult the overview [12] for a discussion of its merits and demerits. Further details of one of the more promising projection methods – *cointegration* – can be found in Refs. [13,14].

This paper will illustrate the use of switching models in the context of the SHM of bridges. Illustration will be via two case studies of real bridges. In the first case, relating to the Tamar Bridge, it will be shown how engineering insight allowed a switching model based on simple polynomials to be hand-crafted. Following a discussion on the *treed Gaussian Processes* (TGPs), a class of powerful switching models, a solution to the Tamar problem is presented that requires much less *a priori* engineering insight. Finally, a case study on the Z24 bridge data is presented which shows how the TGPs allow enhanced damage detection in the presence of confounding influences.

2. Case study 1: the Tamar Bridge

2.1. Background

The Tamar Bridge (Fig. 1) in the south-west of the UK carries a major road across the River Tamar from the town of Saltash to the city of Plymouth. Details of the history and current construction of the bridge can be found in [4]. For the current paper it suffices to say that data were acquired from a vibration-based monitoring system installed by members of the Vibration Engineering Section (VES) of the Department of Civil and Structural Engineering in the University of Sheffield in 2006.¹ The system is based around a set of accelerometers on the deck and some selected cables. The monitoring system records time data at a sampling rate of 64 Hz at 10-min intervals; this data is passed to a computer which carries out an automated modal analysis in order to extract the natural frequencies of the structure. For further details of this, and other monitoring systems on the bridge, the reader can consult [4].

In the recent study [4], simple polynomial response surface models were fitted to the natural frequency data from the Tamar Bridge in order to gain insight into which environmental and operational effects were driving the feature variation. The analysis yields a type of sensitivity analysis that can be used to rank variables according to their effect on natural frequencies. The response surface models used were simple multinomials; if the operational and environmental variables are grouped in a vector $\underline{\theta} = \{\theta_1, \dots, \theta_d\}$, the models take the form,

$$f_i(\underline{\theta}) = \underline{h}(\underline{\theta})^T \underline{\beta} \quad (1)$$

where f_i is the natural frequency under study, $\underline{h}(\underline{\theta})$ is a vector of multinomial basis functions and $\underline{\beta}$ their corresponding vector of coefficients. In (almost) the simplest case, one can take $\underline{h}(\underline{\theta}) = (1, \underline{\theta}^T)^T$ and the model will be linear in the θ_i . Models of the form (1) have the advantage that simple least-squares methods allow coefficient estimation. A major advantage of linear response surface models is that (if the θ_i are standardised), the coefficients in the expansion model give an indication of the importance of the expansion variables. As an illustration, the results of fitting a linear response surface to the first natural frequency of the bridge will be given here. In this case, it was clear from the coefficient estimates in the original study [4] that the dominant effect on the natural frequency was from traffic loading; Fig. 2 shows how well the behaviour is captured using this single independent variable in the model.

It is important to manage expectations of what can be achieved using response surface models. With respect to Fig. 2, one can see that, although the low frequency behaviour is captured very well by the linear model, the higher frequency fluctuations in the natural frequency are not. In this case, the reason is as follows: although the natural frequency measurements are captured at ten-minute intervals and reflect (almost) instantaneous estimates of the natural frequencies, the traffic loadings are estimated from hourly traffic counts and then interpolated onto half-hourly estimates. The estimated traffic loading is thus highly-smoothed, compared to the natural frequency estimates which can show the effect of the true short time-scale variations in traffic at busy periods. One should also bear in mind that random fluctuations in natural frequency estimates will generally be around 1% of their mean value. In general, one must take care building regression models from data from very large civil infrastructure. For example, temperatures can vary across the structure, as can local wind speed; one should take care in pairing input and output variables.

A more careful analysis of the models was carried out by using *F*- and *T*-statistics in order to establish which coefficients were *statistically significant*. For the details of this analysis, the reader can refer to [4]; however, the main result was to indicate that there was a quite small but statistically significant effect of temperature which improved the response surface models slightly from those based on traffic alone. However, a careful consideration of long periods of Tamar data showed a small number of anomalous regions where the models completely failed to capture the behaviour regardless of whether temperature variation was added or not, an example is shown in Fig. 3.

A careful consideration of the variables showed that the times at which the linear response surface failed were associated with high wind speeds, with the wind crossing the bridge deck in the transverse direction to the span, i.e. from the north or south. Further investigation showed that two regimes were visible in the vibrational behaviour of the bridge; below 25 mph wind speed (transverse to the span direction), the vertical acceleration of the bridge showed no significant dependence on wind speed; however, above 25 mph, vertical acceleration increased linearly with wind speed (Fig. 4). For easterly or westerly winds, there was no marked dependence of the vertical acceleration on wind speed. In order to try and capture this

¹ The VES has since moved to the University of Exeter, but has maintained monitoring operations on the Tamar Bridge.



Fig. 1. The Tamar Suspension Bridge.

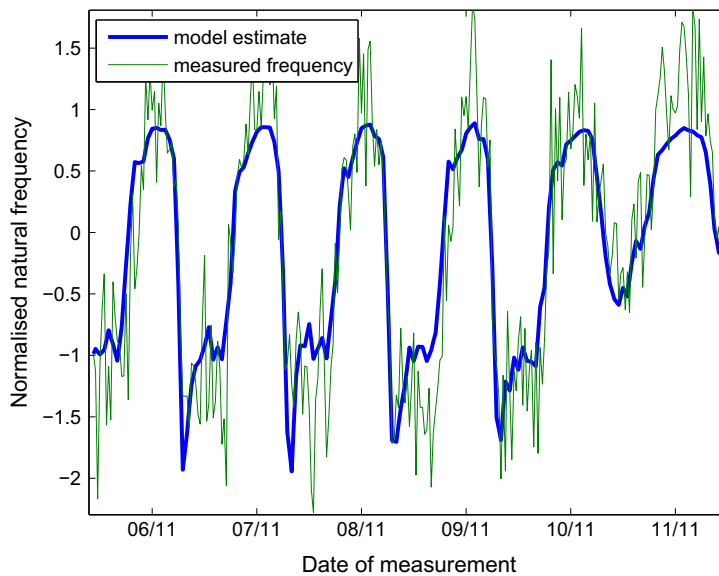


Fig. 2. Linear model of first deck modal frequency with traffic loading only. The model took the form: $\omega_1 = 0.099 - 0.79 \times (\text{traffic load})$.

behaviour, and its possible reflection in the natural frequency, it would have been possible to add wind speed as an input to the response surface model; however, this would have forced adoption of a nonlinear model in order to allow switching between regimes. In this specific case a simpler solution was available; adding the RMS vertical acceleration as an input variable incorporated the required switching behaviour while still allowing a linear response surface model. The introduction of the additional variable on the basis of engineering insight allowed a much more accurate representation of the natural frequency in the formerly anomalous region (Fig. 5). The issue really is about including the correct basis functions in models of the form (1) and this in turn involves identifying the correct variable dependencies; the objective of this paper is provide a means of finding the correct terms directly from the data rather than from engineering insight.

As an aside, Fig. 3 has another useful interpretation; the fact that the model clearly breaks down over a specific period amounts to detection of a performance anomaly. Owners and operators of bridges are interested in performance anomalies i.e. behaviour which highlights hitherto unsuspected physical effects and thus contributes to greater understanding of the bridge behaviour in general. The analysis following the breakdown of the traffic-only model exposed the interesting wind effect.




Fig. 3. Linear model of first deck modal frequency with traffic loading only; a breakdown of the model is clearly visible in the data. The model was the same as that used in Fig. 2.




Fig. 4. RMS vertical acceleration of deck as a function of wind speed, sorted in terms of wind direction.

Although adding the RMS acceleration variable was a pragmatic solution to the problem, it might be considered a little unsatisfactory as it arguably did not reflect root cause and effect. Although the natural frequencies are likely to be dependent on the response amplitude, the heightened amplitude itself is likely to be part of the aeroelastic response to the wind. A more desirable solution might be to include the variables in the model that are possible root causes of the frequency behaviour. However, capturing the switching behaviour with a polynomial response surface model would entail the use of a high-order polynomial and the estimation of many (probably orders of magnitude more) coefficients with minimal physical meaning. Fortunately, there are response surface models available that can identify and capture switching behaviour based on data alone with minimal user intervention. The issue is that these models are much more complex to formulate in the first place. In the next section a group of such models will be introduced and then will be illustrated in the remainder of the paper.

3. Switching regression models - regression trees

3.1. Regression trees

The idea of a regression tree is fairly simple to state (much of the theory and practice of such trees can be attributed to the work of Breiman and colleagues, and a good reference is [15]). The idea is to partition the independent variable space into regions over which the response behaviour is smooth, and to fit low-order regression models over each region. If the parti-




Fig. 5. Linear model of first deck modal frequency with traffic loading and RMS vertical acceleration.

tioning is carried out by hand, the resulting problem is still amenable to linear least-squares methods. The idea, however, of a regression tree in general, is that the partitions are determined from the data as part of the modelling problem; this renders the estimation problem highly nonlinear and alternatives to least-squares are needed. Breiman and co-workers established a greedy algorithm for fitting the trees that gave good (but suboptimal) solutions. If an effective partition of the data is found, linear regression models over each distinct region can give excellent results; however, in principle, *any* regression algorithm can be used once the data have been partitioned into sensible regions. Once the concept of *Classification and Regression Trees* (CART) was established, arguably the next major advance was the development of a Bayesian framework for the algorithm [16,17]. The new algorithm was based on rigorous concepts of probability theory and proved an effective departure from the greedy algorithm. In Bayesian CART, a prior probability distribution was proposed over all possible tree structures as well as all possible coefficients; this was then refined by using the data to determine which tree was supported by the greatest evidence. The original formulation is too complicated to describe here without taking this paper a long way from its illustrative objectives. In the original Chipman formulation, all the regression models within the tree were linear; this restriction was later removed by Gramacy, who replaced the linear models by more powerful *Gaussian Process (GP)* models [18]; Gramacy's work also extended the Bayesian formulation of the problem significantly.

3.2. Gaussian processes

Gaussian process (GP) regression has recently become a popular technique in machine learning, although its roots go back many years [8]. (In fact, the subject has its roots in geotechnical signal processing as early as the 1950s, where the idea was referred to as *Kriging* [19]). In essence, Gaussian processes are an extension of the multivariate Gaussian probability distribution. Unlike most forms of regression model,

$$y = f(\underline{x}) + \epsilon \quad (2)$$

(where \underline{x} is the input, y is the output, and ϵ represents the error), the GP does not return a crisp value $f(\underline{x})$ for any given \underline{x} , but returns a *Gaussian probability distribution*. The GP is thus a Gaussian distribution over functions. Among the advantages of the GP for regression purposes are its principled statistical (Bayesian) foundations and the fact that it automatically returns a confidence interval for predictions. GPs adhere to the Bayesian paradigm in the sense that a number of prior assumptions are made about the function being modelled, and then training data (samples of the features) are used to update and evaluate a posterior distribution over functions. A key assumption is that the model is a smooth function of its inputs and this allows extra information concerning the response to be gained at reduced computational cost. An extended variant of the GP algorithm was developed by O'Hagan and colleagues [9] using additional ideas from Bayesian probability, and it is this variant that will be described briefly here. The Bayesian formulation makes the incorporation into a Bayesian regression tree formulation more direct. Because the implementation of the GP algorithm (unlike its derivation) is straightforward to state, it will be given here in a little detail.

For any set of n input points $\{\underline{x}_1, \dots, \underline{x}_n\}$ (which represent the values of the environmental or operational variables for the specific problem considered here), each of dimension d , the prior beliefs about the corresponding outputs can be represented by a multivariate normal distribution, the mean of which is a least-squares regression fit through the training data,

$$E[f(\underline{x})|\underline{\beta}] = m(\underline{x}) = \underline{h}(\underline{x})^T \underline{\beta} \tag{3}$$

where $\underline{h}(\underline{x})^T$ is a specified (vector) regression function of \underline{x} , and $\underline{\beta}$ is the corresponding vector of coefficients. For simplicity, as it was in much of the early literature, the prior mean $m(\underline{x})$ can actually be set to zero. If the regression form of $m(\underline{x})$ is chosen, one can set $\underline{h}(\underline{x})^T$ to be $(1, \underline{x}^T)$, representing a linear regression, or choose basis functions appropriate to a higher-order polynomial fit or even a generalised linear model. The covariance between output points, which determines the prior variance, is given as,

$$\text{cov}[f(\underline{x}), f(\underline{x}')|\sigma^2, B] = \sigma^2 c(\underline{x}, \underline{x}') \tag{4}$$

where σ^2 is a scaling factor (sometimes called the *height parameter*) and B is a diagonal matrix of (inverse) length-scales, representing the roughness of the output with respect to the individual input parameters. The covariance function commonly adopted, and used here, is a squared exponential function of the form,

$$c(\underline{x}, \underline{x}') = \exp[-(\underline{x} - \underline{x}')^T B (\underline{x} - \underline{x}')] + \sigma_n^2 \tag{5}$$

where the so-called *nugget* σ_n^2 is a hyperparameter accounting for measurement noise. (For simplicity, the closed-form expressions for the case $\sigma_n = 0$ are given below, although all the results presented later were obtained without this restriction.) In the simplest formulations of GPs, σ^2 , B and σ_n^2 are hyperparameters which are determined using a Maximum Likelihood approach; in a full Bayesian approach, they are handled in a more sophisticated manner.

Eqs. (3)–(5) complete the prior specification of the problem; the *posterior distribution* of the outputs is then found by conditioning the prior distribution on the training data \underline{y} (the vector of output points corresponding to the input training set), and integrating out (or marginalising over) the hyperparameters σ^2 and $\underline{\beta}$. The calculation is straightforward but very time consuming, a detailed walkthrough can be found in [20]. The integrals involved are usually all Gaussian, and although the expressions are almost always very complicated, the results can be given in closed form. The result is a Student's t -process, conditional on B and the training data,

$$[f(\underline{x})|\underline{y}, B] \sim t_{n-q}(m^*(\underline{x}), \hat{\sigma}^2 c^*(\underline{x}, \underline{x}')) \tag{6}$$

where the posterior mean and covariance are,

$$m^*(\underline{x}) = \underline{h}(\underline{x})^T \hat{\underline{\beta}} + \underline{t}(\underline{x})^T A^{-1} (\underline{y} - H \hat{\underline{\beta}}) \tag{7}$$

$$c^*(\underline{x}, \underline{x}') = c(\underline{x}, \underline{x}') - \underline{t}(\underline{x})^T A^{-1} \underline{t}(\underline{x}') + (\underline{h}(\underline{x}) - \underline{t}(\underline{x}) A^{-1} H) (H^T A^{-1} H) (\underline{h}(\underline{x}') - \underline{t}(\underline{x}') A^{-1} H)^T \tag{8}$$

and some new notation has been introduced:

$$\underline{t}(\underline{x})^T = (c(\underline{x}, \underline{x}_1), \dots, c(\underline{x}, \underline{x}_n)) \tag{9}$$

$$H^T = (\underline{h}(\underline{x}_1), \dots, \underline{h}(\underline{x}_n)) \tag{10}$$

$$A = \begin{pmatrix} 1 & c(\underline{x}_1, \underline{x}_2) & \dots & c(\underline{x}_1, \underline{x}_n) \\ c(\underline{x}_2, \underline{x}_1) & 1 & & \vdots \\ \vdots & & \ddots & \\ c(\underline{x}_n, \underline{x}_1) & \dots & & 1 \end{pmatrix} \tag{11}$$

$$\hat{\underline{\beta}} = (H^T A^{-1} H)^T H^T A^{-1} \underline{y} \tag{12}$$

$$\hat{\sigma}^2 = \frac{\underline{y}^T (A^{-1} - A^{-1} H (H^T A^{-1} H) H^T A^{-1}) \underline{y}}{n - d - 3} \tag{13}$$

$$\underline{y} = (f(\underline{x}_1), \dots, f(\underline{x}_n))^T \tag{14}$$

Determination of this model is basically an exercise in machine learning and therefore its quality is critically dependent on the number and distribution of training data points in the input space, and the values of the hyperparameters. The expressions for $\hat{\underline{\beta}}$ and $\hat{\sigma}^2$ shown above are the result of marginalisation; however, it can be shown that they actually coincide with least-squares estimates. The diagonal matrix of roughness parameters B cannot generally be integrated out analytically and it is usually evaluated using maximum likelihood estimation or Markov Chain Monte Carlo (MCMC); this calculation typically represents the most computationally intensive part of the process.

For prediction purposes, Eqs. (7) and (8) can be used. The posterior mean (7) gives the prediction and the posterior variance $\hat{\sigma}^2 c^*(\underline{x}, \underline{x})$ is computed from (8). Any desired confidence intervals for prediction can then be determined from the posterior variance.

This almost completes the basic description of the models used here - *Treed Gaussian Processes* (TGPs). Although the actual implementation is too complex to explain here in any detail, the basic ingredients have been covered.² The TGP partitions the variable space in much the same way as a Bayesian CART and then fits GP regressors over each independent region. The software used for modelling in this work is the TGP package written by Gramacy in the language R [21]. A useful feature of the TGP software is the *limiting linear model* [22]. Because of the Bayesian framework adopted for the TGP, the software can essentially assess if the evidence for a full GP model outweighs that for a linear model; if the software judges that a linear model is sufficient over a given region of the feature, it switches to the simpler representation.

In the next section, the use of the treed models is demonstrated on the data from the Tamar Bridge presented in Section 2. Another case study will then be presented based on the Z24 Bridge.

4. Case study 1 revisited: the Tamar Bridge

The treed models were next applied to the problem discussed in the second section of this paper; namely the modelling of the first natural frequency of the Tamar Bridge on the data where the high wind speeds caused a breakdown in the model based on traffic loading alone. A multivariate linear model was fitted to the data with input variables: traffic, wind speed and wind angle. Wind angle was added to account for the observed fact that the RMS vertical acceleration of the bridge only increased at high wind speeds if the wind was in the N-S direction. The TGP package was used to fit a Bayesian treed linear model. The results were excellent; the most probable tree structure determined by the model automatically switched at (normalised) wind speeds above 0.803 (corresponding to 25 mph) as shown in Fig. 6 (Wind speed is labeled V2 in Fig. 6.) Some refinement of the model was obtained by allowing the model to switch on the traffic variable (variable V1 in Fig. 6) at low wind speeds. Interestingly, the wind angle was not used to determine any switching behaviour.

In terms of the data used for training, Fig. 7 shows the wind speed variable with the horizontal line indicating the automatically determined threshold for switching. The increase in speed is very marked over the previously anomalous period and this may explain why additional information from wind angle is not needed here; it is also the case that the wind direction did not change much over the short period captured in the training data. The predictions from the treed linear model are shown in Fig. 8. One can see that the model now captures the all of the low-frequency observed behaviour of the first natural frequency very well. Even with the higher-frequency component visible in the measured data, the data falls within the 95% confidence intervals for prediction as desired.

The treed GPs have proved effective in automatically providing a model which accommodates the switching behaviour due to the wind conditions. However, because no data from a damage state of the Tamar Bridge is available, it has not been possible to show that removal of the environmental variations improves diagnostic capability for SHM. By considering a different case study – the Z24 bridge – it will be shown in the next section how TGPs can enhance the ability to detect damage when confounding influences are present.

5. Case study 2: the Z24 Bridge

5.1. Background

The Z-24 Bridge, a pre-stressed concrete highway bridge in Switzerland (Fig. 9), was subject to a comprehensive monitoring campaign under the 'SIMCES project' [24]. Although ultimately demolished in the late 1990s, the bridge has since become a landmark in experimental benchmark studies for SHM. The monitoring campaign, which spanned a whole year, tracked modal parameters and included extensive measurement of the environmental factors affecting the structure, such as air temperature, soil temperature, humidity etc. The Z24 monitoring exercise was an important study in the history of SHM developments because towards the end of the monitoring campaign researchers were able to introduce a number of realistic damage scenarios to the structure. In order, these scenarios were [25]:

- Pier settlement
- Tilt of foundation followed by settlement removal
- Concrete spalling
- Landslide
- Concrete hinge failure
- Anchor head failure
- Tendons rupture.

² In the full implementation of the TGP code [21], all the hyperparameters are dealt with in a principled manner, including the roughness parameters. In fact the hyperparameters are represented by prior densities, which have their own hyper-hyperparameters. The result of this extra generality is that a much more complex algorithm is needed than the basic Bayesian approach outlined earlier in this section.

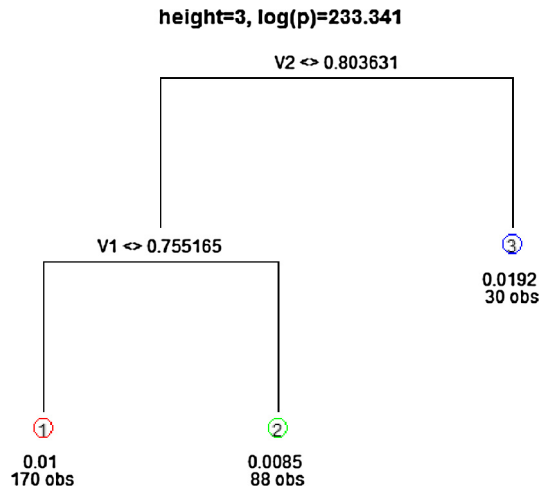


Fig. 6. Most probable (MAP) tree structure for treed linear model of Tamar first natural frequency: V1 is traffic load, V2 is wind speed, obs indicates the number of data points assigned to each leaf (model) in the decision tree. The term $\leftrightarrow a$ is a convention of the TGP software and simply indicates that the branch occurs at a and the left branch represents $< a$, while the right branch represents $> a$.




Fig. 7. Wind speed variable over the period of interest for the treed linear model.

The papers immediately produced as a result of the COST action were [26–29], although many have followed.

The SHM features of interest here are the natural frequencies of the bridge which were tracked over the period of one year and additionally over the period where the bridge was damaged according to the various scenarios. Modal properties of the bridge were extracted from acceleration data [30]. Fig. 10 shows a time history of the four natural frequencies between 0 and 12 Hz of the bridge. In total, there were 5652 records of the natural frequencies made; however, there are gaps in the records for individual frequencies where the monitoring system failed to identify a given mode. The dashed vertical line marks the start of the period where the different levels of damage (starting with pier lowering) were introduced beginning at point 4918. On inspection of Fig. 10, the natural frequencies of the bridge are by no means stationary. There are some large fluctuations in the first half of the time history before the introduction of any damage. These fluctuations occurred during periods of very cold temperatures and have been associated with an increase in stiffness of the asphalt layer on the bridge deck




Fig. 8. Bayesian treed linear model of first deck modal frequency based on traffic loading, wind speed and wind angle.




Fig. 9. Z24 bridge longitudinal section and top view [23].




Fig. 10. Time Histories of the extracted natural frequencies of the Z-24 Bridge, monitored over one year including a period when damage was introduced.

at these temperatures [30].³ The natural frequency time histories are, therefore, another good example of how damage sensitive parameters can also be very sensitive to environmental variations, in this case temperature.

As the natural frequencies in their current form would not be suitable to monitor as a damage sensitive feature, some action must be taken to remove the variable set's sensitivity to temperature. Although the regression approach discussed in this paper is one of the simpler methods conceptually, it is indicated here as, in the Z24 case, the modal properties of the bridge are nonlinearly dependent on temperature via the temperature dependence of the material and geometric properties of the structure (as an example, Fig. 11 plots how the first natural frequency changes with temperature). The bilinear form of the dependence on temperature also means that the switching models should prove useful. For the purposes of this case study, the second natural frequency will be discussed, as the first natural frequency proved rather insensitive to the damage.

5.2. Analysis

As a benchmark, the first analysis of the data simply fitted a linear regression model. However, the analysis was carried out in the Bayesian framework in order to provide confidence intervals for the model predictions [8,21]. (For this example, 99% confidence limits have been chosen, in slight contrast to the 95% limits chosen for the Tamar data.) In order to develop the models in a principled manner, the natural frequency data were divided into a training set (to establish the model) and a testing set to make sure that the model could generalise. The testing set spanned the last 2000 samples of the second natural frequency data as shown in Fig. 10, in order to encompass regions of temperature variation *and* damage; after removing the gaps from measurement failures, this set comprised 1939 points, in which the damage initiation occurs at point 1205. The training data comprised 1000 points taken alternately from the record, immediately before the test data and thus only encompassed the region of temperature variation.

Fig. 12 shows the model fit to the training data, the linear model is clearly incapable of capturing the bilinear dependence on temperature. However, an interesting feature is that the 99% confidence intervals on the predictions computed by the software do mostly encompass the data. Now, recall that the object of the exercise is to *remove* the temperature dependence in order to create a feature for damage detection. The *residual* data for the testing set could be computed by subtracting the model predictions from the data. If the structure were to remain in normal condition throughout the testing period so that temperature changes were the only source of variation, the residual should resemble a stationary white noise process with a mean of zero. Confidence intervals could be constructed from the residual values on the training data or could be adopted from the model fit and then any excursions outside the confidence intervals would potentially indicate damage. As a slightly more informative alternative to plotting residuals, the model predictions with their 99% confidence limits as compared to the measured data will be shown in this chapter; damage is indicated when the measured data deviates significantly from the predictions i.e. goes outside the model confidence intervals. For the Bayesian linear model fitted to the training data, Fig. 13 compares the measured data on the testing set with the model predictions; although the damage begins to show itself by a

³ In fact, there is not universal agreement on the cause of the nonlinearity; the most popular explanation is in terms of the stiffening of the asphalt layer, and this is the explanation assumed here.




Fig. 11. First natural frequency nonlinear dependence on temperature.

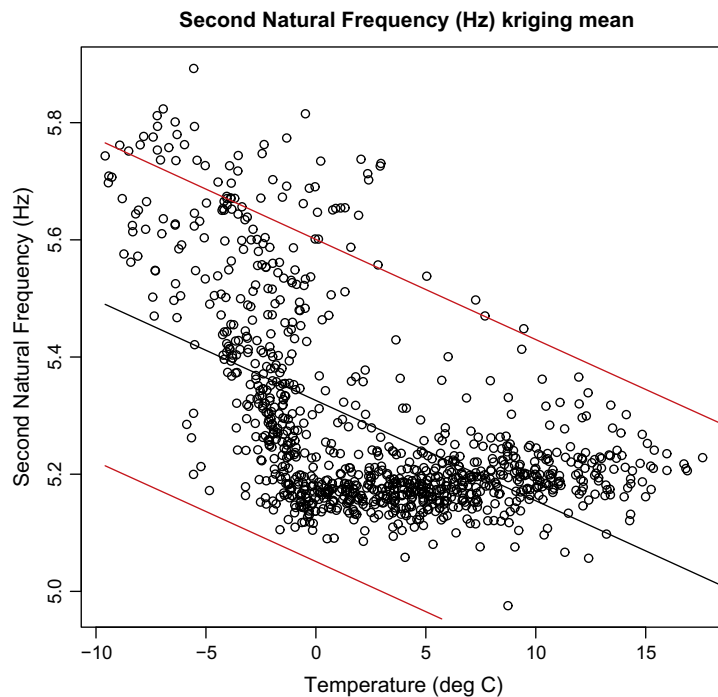


Fig. 12. Bayesian linear model of Z24 s natural frequency as a function of temperature: training data.

monotonic decline in the measured data, the model predictions have so little confidence that there are no excursions outside the 99% limits. The poor fidelity of the model in capturing the temperature variations has resulted in a residual insensitive to damage. (Within the TGP software [21], there are a number of methods of computing model predictions, including sampling from the posterior distribution of the model; however, one of the simplest methods is to use the *Kriging* means and variances as specified in Eqs. (7) and (8), and this is the approach adopted here.)

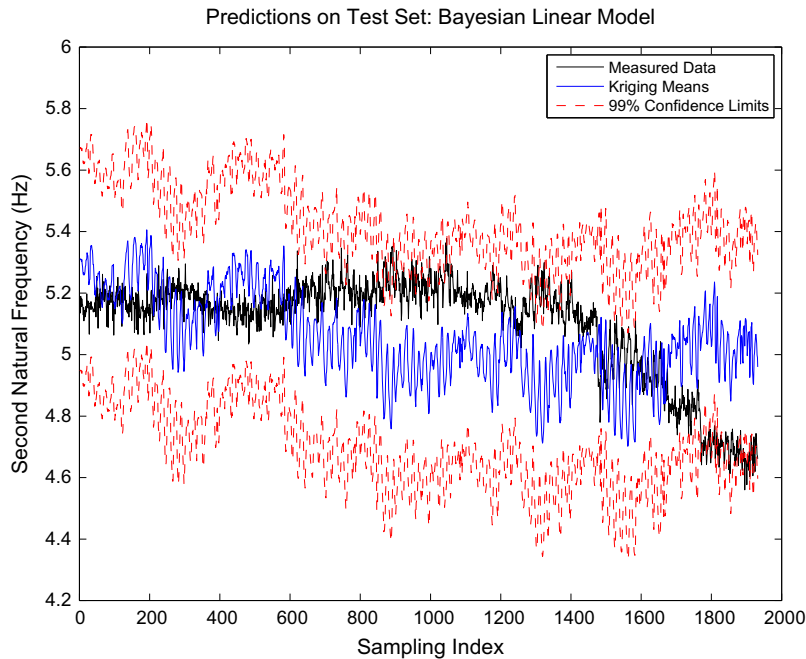


Fig. 13. Bayesian linear model predictions compared to true measurements on testing data set. The onset of damage is at point 1205.

The next model fitted was a treed linear model e.g. a Bayesian CART model. While Fig. 12 shows very clear evidence of the bilinear switching behaviour in the measured data at close to zero Celcius, the variations in the training data between zero and three degrees convince the algorithm that two switching points are needed. This is not an issue in terms of fitting a good predictive model, but it does mean that the model has perhaps not captured the true physics. As one might expect, the treed linear model fits the data much better than the simple linear one and this is reflected in the much tighter confidence bounds on the training data (Fig. 14). When the model predictions are compared to the measured data on the testing set, the results are far better than those for the linear model (consider the predictions before the onset of damage at point 1205 in Fig. 13); the much higher prediction confidence results in a very clear detection of the damage when the measured data moves outside the confidence bounds not long after initiation of damage (Fig. 15). The important point here is that the complexity of the models in the various regimes is not the issue, more important is the recognition that there *are* different regimes. Another interesting feature can be seen in Fig. 15; this is the presence of bursts of low prediction confidence at points in the early part of the testing data. The reason for this is that the points correspond to slightly higher temperatures than were present in the period over which the training data were taken; the model recognises that it is moving from interpolation to extrapolation and adjusts its confidence accordingly.

Finally, in Fig. 14, one can see that there are discontinuities in the predictive mean and the confidence intervals; this is a general feature of TGPs as independent GPs are fitted on each separate region of the input space and there is no mechanism to enforce continuity across region boundaries i.e. at switching points. Furthermore, because the confidence intervals are estimated from the Kriging means and variances, they are symmetric about the mean estimate. The presence of the outliers is arguably skewing the true distribution towards the higher natural frequency values; however, the Gaussian process cannot express this. Although one cannot see the lower confidence bound, one knows where it is because of the symmetry; however, there are no data points close to it.

For the next exercise, a Bayesian Gaussian process model was fitted. This model structure allows the nonlinear dependence on temperature to be modelled, but assumes that it can be captured by a single GP over the entire feature space. The results on the training set are shown in Fig. 16. While the results are better than the linear model, they are inferior to those from the treed linear model. The explanation for this lies with the covariance function which determines the smoothness of the predictions for the GP model. At the point of discontinuity in the data, the GP requires a covariance function changing rapidly over short time scales; away from the discontinuity the GP would like a very slowly-changing covariance function to reflect the smoothness of the model. The two objectives cannot be met with a single B hyperparameter and so the GP adopts a compromise value; this worsens the prediction confidence and generally expands the confidence interval. The effect of this on damage sensitivity can be seen in the comparison plot of Fig. 17; although the damage is detected when the test data leaves the prediction confidence intervals, it is detected at a later time than for the treed linear model.

The final model fitted to the data was the treed GP. As there is no real advantage of doing otherwise, the variant of the model allowing switching to a limiting linear model was used. The results on the training data can be seen in Fig. 18. Once

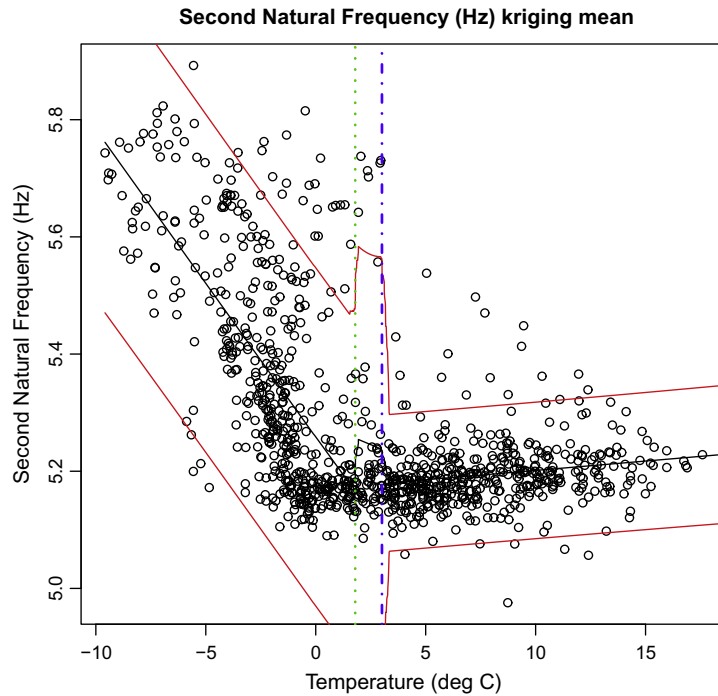


Fig. 14. Bayesian treed linear model of Z24 s natural frequency as a function of temperature: training data. Vertical dotted and dash-dot lines indicate the position of switch points.

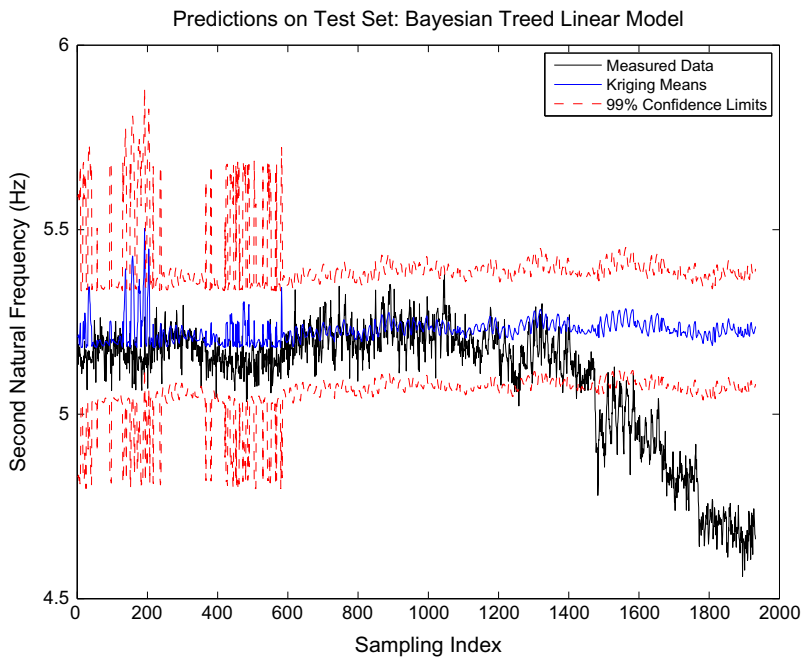


Fig. 15. Bayesian treed linear model predictions compared to true measurements on testing data set. The onset of damage is at point 1205.

again, the switching model proves superior; however, in this case there are two additional features of interest for the treed GP. In the first case, the model has recognised that only one switching point is needed. As in the treed linear case, the model switches at a higher point than freezing point; however, the nonlinear nature of the GP means that the behaviour of the data below that is captured by a single GP. Another interesting aspect of the model is that it switches to a linear model above the

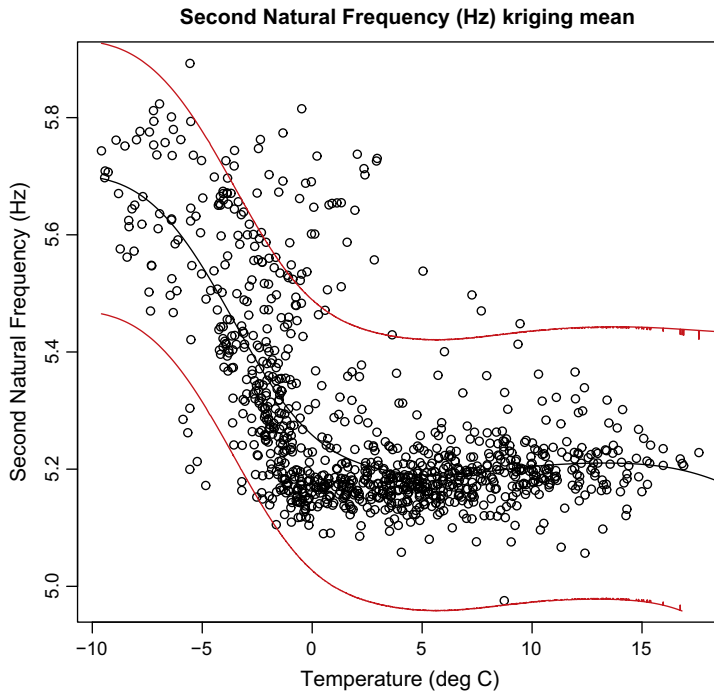


Fig. 16. Bayesian GP model of Z24 s natural frequency as a function of temperature: training data.

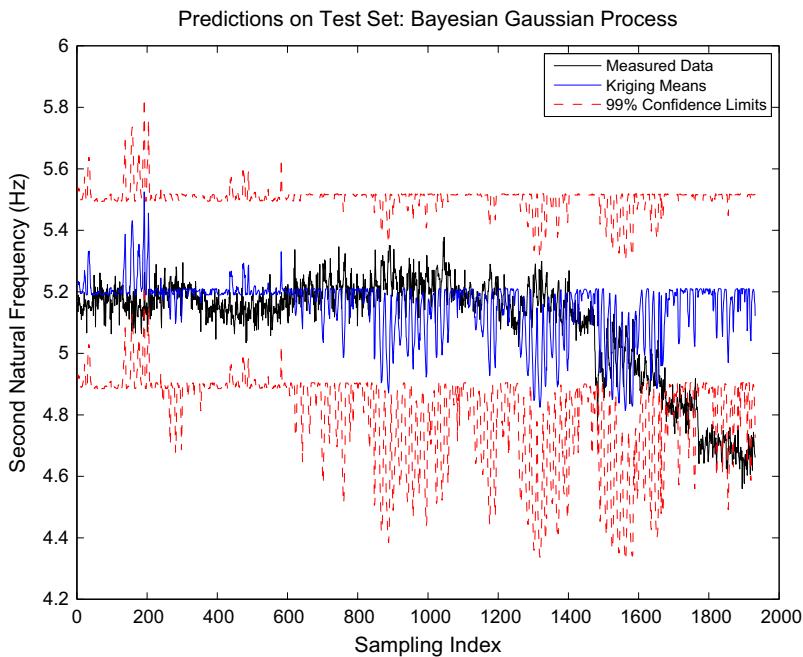


Fig. 17. Bayesian GP model predictions compared to true measurements on testing data set. The onset of damage is at point 1205.

switch point; there is nothing to be gained in terms of model evidence from adopting the more complex model. In terms of the damage detection issue, the comparison plot in Fig. 19 shows that the damage is detected as promptly as in the case of the treed linear model but is not significantly different. The treed GP essentially achieves the same result as the treed linear model, but with a simpler partition on the feature space.

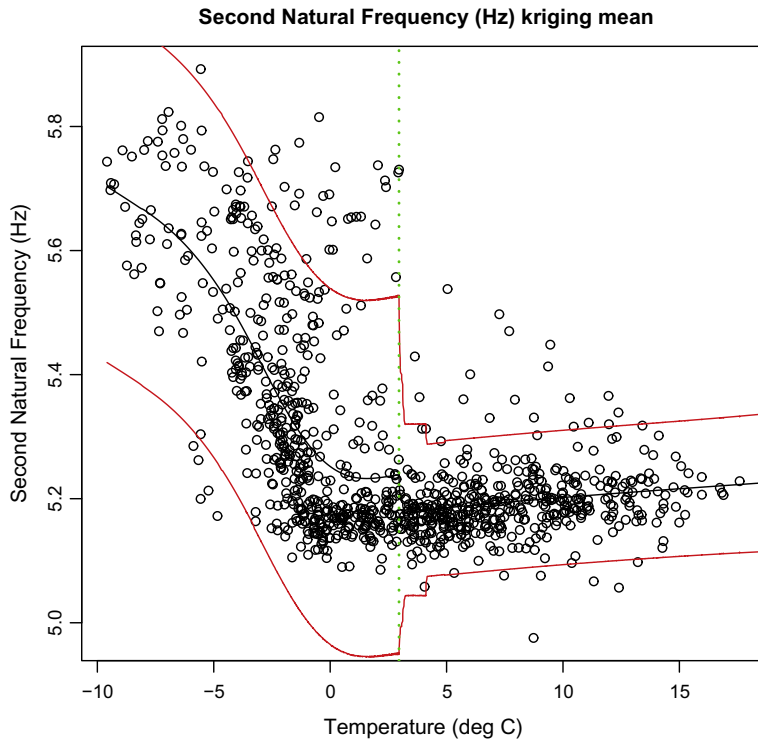


Fig. 18. Bayesian treed GP model of Z24 s natural frequency as a function of temperature: training data. The vertical dotted line indicates the position of the switch point.

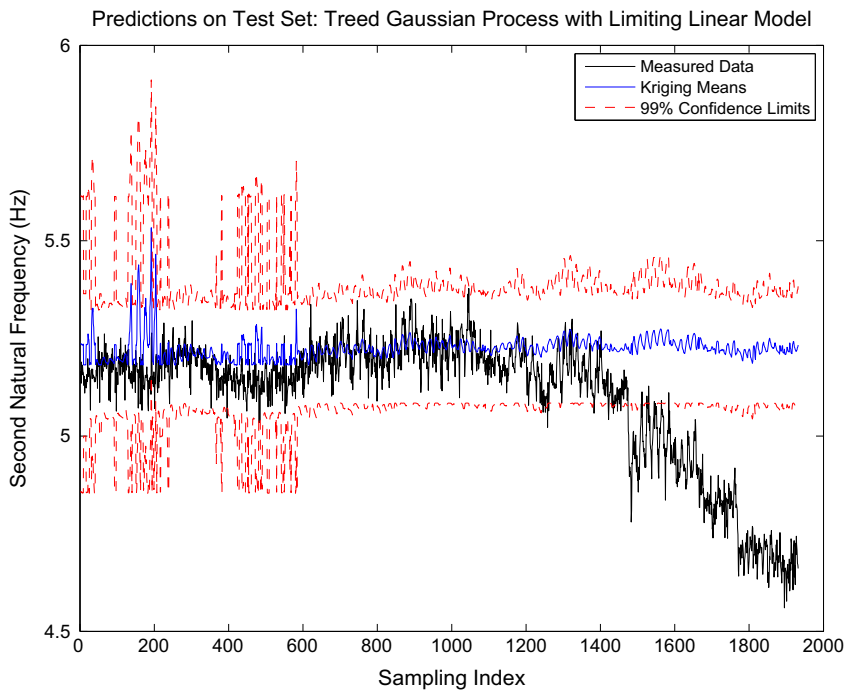


Fig. 19. Bayesian treed GP model predictions compared to true measurements on testing data set. The onset of damage is at point 1205.

6. Conclusions

There is no cause here for lengthy conclusions; the objectives of this paper have simply been to present a powerful and versatile class of response surface models with the capability to switch discontinuously between different regimes and to show the value of the models in the context of vibration data-based SHM. However, there has been the opportunity to touch on some interesting points along the way. The models presented were initially applied here as black-box models – the intention was simply to learn a predictive model from the data; however, the fact that the models can automatically learn the switching points between regimes arguably lifts them into the domain of grey-box models when a physical interpretation can be assigned to the regime boundaries. In the case of the Z24 Bridge, the physics is usually associated with changes in the stiffness of the asphalt at temperatures below zero Celcius, although alternative explanations have been offered, as mentioned before. In the case of the Tamar Bridge, the cause of the switching is arguably more obscure; however, it is possible that it is the result of unsteady aerodynamics and as this is a subject of some interest to bridge researchers, identifying and modelling the switching behaviour may well be of value. The models are also arguably superior to other schemes in terms of parsimony; if simple models are allowed over the identified regimes, this means fewer ‘coefficients’ need to be identified; this in turn reduces the demands on acquiring data and allows models to generalise away from their training data better. Finally, the authors would argue that, if responses are switching discontinuously between regimes, it may be harmful to smooth out the effects by fitting overall polynomial models for example; at the very least, one is not giving proper respect to the physics shown.

Acknowledgements

The support of the UK Engineering and Physical Sciences Research Council (EPSRC) through grant reference numbers EP/J016942/1 and EP/K003836/2 is gratefully acknowledged. The authors would like to thank their colleague Dr. Will Becker of the Joint Research Centre, Ispra for drawing their attention to the TGP software package and for a number of valuable discussions on the nature of treed and non-treed Gaussian processes. Thanks are also due to Robert Gramacy for providing the community with the TGP package in the first place. In terms of the bridge data discussed here, thanks are due to Professor James Brownjohn and Dr. Ki-Young Koo, formerly of the Department of Civil and Structural Engineering at the University of Sheffield, for providing access to the Tamar Bridge data and for valuable guidance on the SHM of civil infrastructure.

References

- [1] C. Farrar, K. Worden, *Structural Health Monitoring: A Machine Learning Perspective*, John Wiley and Sons, 2013.
- [2] C. Farrar, W. Baker, T. Bell, K. Cone, T. Darling, T. Duffey, Damage characterization and damage detection in the I-40 bridge over the Rio Grande, Tech. Rep., Los Alamos National Laboratories Report, LA-12767-MS, 2005.
- [3] H. Sohn, Effects of environmental and operational variability on structural health monitoring, *Philos. Trans. Roy. Soc. – Ser. A* 365 (2007) 539–561.
- [4] E. Cross, K.-Y. Koo, J. Brownjohn, K. Worden, Long-term monitoring and data analysis of the Tamar Bridge, *Mech. Syst. Signal Process.* 35 (2013) 16–34.
- [5] Y. Ni, X. Hua, K. Fan, J. Ko, Correlating modal properties with temperature using long-term monitoring data and support vector machine technique, *Eng. Struct.* 27 (2005) 1762–1773.
- [6] Y. Ni, H. Zhou, J. Ko, Generalization capability of neural network models for temperature-frequency correlation using monitoring data, *J. Struct. Eng.* 135 (2009) 1290–1300.
- [7] C. Bishop, *Pattern Recognition and Machine Learning*, Springer, 2013.
- [8] C. Rasmussen, C. Williams, *Gaussian Processes for Machine Learning*, The MIT Press, 2006.
- [9] M. Kennedy, C. Anderson, S. Conti, A. O’Hagan, Case studies in Gaussian process modelling of computer codes, *Reliab. Eng. Syst. Saf.* 91 (2006) 1301–1309.
- [10] W. Becker, J. Rowson, J. Oakley, A. Yoxall, G. Manson, K. Worden, Bayesian sensitivity analysis of a model of the aortic valve, *J. Biomech.* 44 (2011) 1499–1506.
- [11] D. Batterbee, N. Sims, W. Becker, K. Worden, J. Rowson, Computational model of an infant brain subjected to periodic motion: simplified modelling and bayesian sensitivity analysis, *Proc. Inst. Mech. Eng., Part H, J. Eng. Med.* 225 (2011) 1036–1049.
- [12] K. Worden, T. Baldacchino, J. Rowson, E. Cross, Some recent developments in SHM based on nonstationary time series analysis, *Proc. IEEE* 104 (2016) 1589–1603.
- [13] E. Cross, K. Worden, Q. Chen, Cointegration: a novel approach for the removal of environmental trends in structural health monitoring data, *Proc. Roy. Soc. – Ser. A* 467 (2011) 2712–2732.
- [14] E. Cross, G. Manson, K. Worden, S. Pierce, Features for damage detection with insensitivity to environmental and operational variations, *Proc. Roy. Soc. – Ser. A* 468 (2012) 4098–4122.
- [15] L. Breiman, J. Friedman, C. Stone, R. Olshen, *Classification and Regression Trees*, Chapman and Hall/CRC, 1984.
- [16] H. Chipman, E. George, R. McCulloch, Bayesian cart model search, *J. Am. Stat. Assoc.* 93 (1998) 935–948.
- [17] H. Chipman, E. George, R. McCulloch, Bayesian treed models, *Mach. Learn.* 48 (2002) 299–320.
- [18] R. Gramacy, *Bayesian Treed Gaussian Process Models*, PhD Thesis, University of California, 2005.
- [19] G. Matheron, *The intrinsic random functions and their applications*, *Adv. Appl. Probab.* 5 (1973) 439–468.
- [20] W. Becker, *Uncertainty Propagation Through Large Nonlinear Models*, PhD Thesis, Department of Mechanical Engineering, University of Sheffield, 2011.
- [21] R. Gramacy, TGP: an R package for Bayesian nonstationary, semiparametric nonlinear regression and design by treed Gaussian process models, *J. Stat. Software*, vol. 19.
- [22] R. Gramacy, H. Lee, Gaussian processes and limiting linear models, *Comput. Stat. Data Anal.* 53 (2008) 123–136.
- [23] C. Kramer, C.D. Smet, G.D. Roeck, Z24 Bridge damage tests, in: *Proceedings of 17th International Modal Analysis Conference (IMAC)*, Kissimmee, Florida, 1999.
- [24] G.D. Roeck, The state-of-the-art of damage detection by vibration monitoring: the SIMCES experience, *J. Struct. Control* 10 (2003) 127–134.
- [25] A.-M. Yan, G. Kerschen, P.D. Boe, J.-C. Golinval, Structural damage diagnosis under varying environmental conditions. Part II: local PCA for non-linear cases, *Mech. Syst. Signal Process.* 19 (2005) 865–880.
- [26] J. Maeck, G.D. Roeck, Damage assessment using vibration analysis on the Z24 Bridge, *Mech. Syst. Signal Process.* 17 (2003) 133–142.

- [27] L. Mevel, M. Gourset, M. Basseville, Stochastic subspace-based structural identification and damage detection and localisation - application to the Z24 Bridge benchmark, *Mech. Syst. Signal Process.* 17 (2003) 143–151.
- [28] L. Garibaldi, S. Marchesiello, E. Bonisoli, Identification and up-dating over the Z24 benchmark, *Mech. Syst. Signal Process.* 17 (2003) 152–161.
- [29] J. Kullaa, Damage detection of the Z24 Bridge using control charts, *Mech. Syst. Signal Process.* 17 (2003) 163–170.
- [30] B. Peeters, G.D. Roeck, One-year monitoring of the Z24-Bridge: environmental effects versus damage events, *Earthq. Eng. Struct. Dynam.* 30 (2001) 149–171.