



This is a repository copy of *Physics-informed machine learning for structural health monitoring*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/188200/>

Version: Accepted Version

Book Section:

Cross, E.J., Gibson, S.J., Jones, M.R. et al. (3 more authors) (2022) Physics-informed machine learning for structural health monitoring. In: Cury, A., Ribeiro, D., Ubertini, F. and Todd, M.D., (eds.) Structural Health Monitoring Based on Data Science Techniques. Structural Integrity (21). Springer Cham , pp. 347-367. ISBN 9783030817152

https://doi.org/10.1007/978-3-030-81716-9_17

This is a post-peer-review, pre-copyedit version of a chapter published in Structural Health Monitoring Based on Data Science Techniques. The final authenticated version is available online at: https://doi.org/10.1007/978-3-030-81716-9_17

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

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



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

Physics-informed machine learning for Structural Health Monitoring

E. J. Cross, S. J. Gibson, M. R. Jones, D. J. Pitchforth, S. Zhang and T.J. Rogers

Abstract The use of machine learning in Structural Health Monitoring is becoming more common, as many of the inherent tasks (such as regression and classification) in developing condition-based assessment fall naturally into its remit. This chapter introduces the concept of physics-informed machine learning, where one adapts ML algorithms to account for the physical insight an engineer will often have of the structure they are attempting to model or assess. The chapter will demonstrate how grey-box models, that combine simple physics-based models with data-driven ones, can improve predictive capability in an SHM setting. A particular strength of the approach demonstrated here is the capacity of the models to generalise, with enhanced predictive capability in different regimes. This is a key issue when life-time assessment is a requirement, or when monitoring data do not span the operational conditions a structure will undergo.

The chapter will provide an overview of physics-informed ML, introducing a number of new approaches for grey-box modelling in a Bayesian setting. The main ML tool discussed will be Gaussian process regression, we will demonstrate how physical assumptions/models can be incorporated through constraints, through the mean function and kernel design, and finally in a state-space setting. A range of SHM applications will be demonstrated, from loads monitoring tasks for off-shore and aerospace structures, through to performance monitoring for long-span bridges.

Key words: Physics-informed machine learning, grey-box modelling, Gaussian-process regression.

E. J. Cross, S. J. Gibson, M. R. Jones, D. J. Pitchforth, S. Zhang, T.J. Rogers
Dynamics Research Group, Department of Mechanical Engineering, University of Sheffield, Mapin St, Sheffield, S1 3JD UK. e-mail: e.j.cross@sheffield.ac.uk

1 Introduction

As performance and monitoring data from our structures become more abundant, it is natural for researchers to turn to methods from the machine learning community to help with analysis and construction of diagnostic/prognostic algorithms. Indeed, within the SHM research field, use of neural networks, support vector machines and Gaussian processes for regression and classification problems has become common place [1]. These methods bring the opportunity to learn complex relationships directly from data, without a requirement of in-depth knowledge of the system. As an example from the authors' own work, in [2] we employed a Gaussian process (GP) regression to predict strain on a landing gear from measured accelerations across the aircraft. Use of a suitably trained GP circumvents the need to build complex physics-based models of the gear for fatigue life calculations. This kind of model is often referred to as a 'black-box' model to reflect the fact the data drives the structure of the model rather than knowledge of the physics at work.

At the other end of the spectrum the term 'white-box' model can be used to describe a model purely constructed from knowledge of physics, (e.g differential equations and finite element models). Physics-based modelling and updating were common early themes in the structural health monitoring research field [3]. However, for large or critical engineering structures that operate in (often extreme) dynamic environments, such as wind turbines, aircraft, gas turbines, etc, predictive modelling from a white-box perspective presents particularly difficult challenges. Loading is often unknown and unmeasured, and dynamic behaviour during operation needs to be fully captured by a computational model, but is sensitive to small changes in (or disturbances to) the structure. Validation and updating of large complex models bring their own challenges and remain active research areas [4–7].

Due to the availability of monitoring data, the inherent challenges of the physics-based approach and the promise of machine learning methods, it is fair to say that the data-driven approach to SHM has become dominant in the research field. A significant issue with the use of any machine learning method in an engineering application, however, is the availability of suitable data with which to train the algorithm. As the model learns from the data, it is only able to accurately predict behaviour present in the data on which it was trained. As an example, Figure 1 shows a black-box model trained to predict the bending strain on an aircraft wing during different manoeuvres to inform an in-service fatigue assessment. This data set comprises of 84 flights, five of which are used for model training. The trained model is able to generalise well with a very low prediction error for the majority of the flights - Figure 1a shows a typical strain prediction for a flight not included in the training set (normalised mean-squared error¹ (nMSE)= 0.29% across the whole flight). However, for the flight shown in Figure 1b, the model is unable to predict the strain as accurately (nMSE 4.20% across the whole flight) - this flight was atypical in

¹

$$\text{nMSE} = \frac{100}{n\sigma_y^2} \sum (y_i - f_i)^2 \quad (1)$$

where y_i and f_i are the measurements and predictions respectively, $i = 1 \dots n$.

terms of operating conditions - it was a low altitude sortie over ground, characterised by the turbulent response one can see in the figure. These conditions are different from those included in the training set and the model is unable to generalise and predict the strain as well in this case.

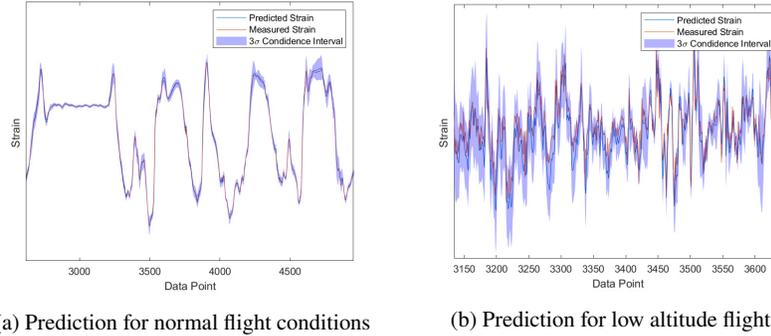


Fig. 1: Example of difficulty predicting behaviour outside of normal operation conditions - strain prediction on an aircraft wing during manoeuvres - see [8, 9] for more details.

In general, but especially because of the inherent flexibility in many of the machine learning models commonly used, extrapolation should not be attempted in this setting. For an SHM application, this will generally mean that training data are required from all possible operating conditions that the structure will see. For many applications this is currently infeasible, although as data collection becomes more commonplace, the situation will improve somewhat. Where a supervised approach is needed, this problem is exacerbated by the general lack of access to data from structures in a damaged state which remains a large barrier to effective diagnosis and prognosis [10].

Currently a programme of work by the authors is pursuing a physics-informed machine learning approach to attempt to address some of these issues in a structural dynamics setting. The aim is to bring together the flexibility and power of state-of-the-art machine learning techniques with more structured and insightful physics-based models derived from domain expertise. This reflects a natural wish that any inferences over our structures will be informed by both our engineering knowledge and relevant monitoring data available.

The potential means of combining physics-based models and data-driven algorithms are many, ranging from employing ML methods for parameter estimation [11, 12], to using them as surrogates or emulators [13–15]. Of interest here are methods where the explanatory power of a model is shared between physics-based and data-driven components. We will often refer to these approaches as ‘grey-box’ models (a combination of white and black-box components), but the term ‘hybrid modelling’ is equally applicable. The philosophy followed in our work is to embed

fundamental physical insight into a machine learning algorithm. In doing so, our aim is that the role of the machine learner is one of augmenting the explanatory power of the model rather than being employed to correct any potential error or bias in the physical foundation. This chapter will explore this idea in a Bayesian setting, introducing a number of different approaches and demonstrating their usefulness in an SHM setting.

2 Grey-box models, overview and literature

The term ‘grey-box model’ is perhaps most familiar to those from a control engineering background. Sohlberg [16, 17] provides a useful review and overview of grey-box models in this context². Figure 2 attempts to capture and summarise some of the currently available modelling approaches relevant for challenges in structural health monitoring on the white to black spectrum. Note that the “degree of greyness” of the models in the middle region will change according to implementation and application.

At the whiter end of the spectrum are modelling approaches where data are used for parameter estimation or model form selection, (with the buoyant field of equation discovery fitting in here, see [18, 19]). See also [20]. *Residual models* are those that use a data-driven approach to account for the observed difference between a physics-based model and measurements, with general form

$$y = \underbrace{f(x)}_{\text{White-box}} + \underbrace{\delta(x) + \varepsilon}_{\text{Black-box}} \quad (2)$$

where $f(x)$ is the output of the physical model, $\delta(x)$ is the model discrepancy and ε is the process noise (see for example [21–24]). The discrepancy term is often used to correct a misspecified physical model, giving rise to the term ‘bias correction’. Residual based approaches have proven effective across a range of SHM tasks including damage detection [25] and modal identification [26]. Here we are interested in residual modelling in the context of compensation for un-captured/missing behaviours in the physics-based model (discussed further in Section 3).

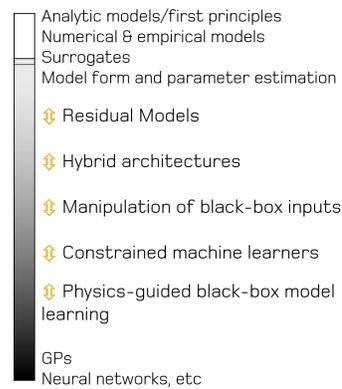


Fig. 2: Some modelling approaches on the white-black spectrum.

² The use of grey-box models within the control community is undergoing somewhat of a revival there and a good snap-shot of this can be gained by looking at the contributions to the most recent Nonlinear System Identification Benchmarks Workshop (<http://www.nonlinearbenchmark.org>)

The term *hybrid architectures* reflects the wider possibilities for combinations of white and black models (which could include the summation form of (2)). Section 5 will demonstrate one such example of combining data-driven and physics based models in a state-space setting.

The remainder of the spectrum contains models with structures that are data-driven/black-box in nature. Sohlberg [16]) describes *semi-physical modelling* as when features are subject to a nonlinear transformation before being used as inputs to a black-box model, we also refer to this as *input augmentation*, see [8, 27, 28] for more examples. We place these examples under the heading of *manipulation of black-box inputs*.

Section 4 of this chapter will discuss *constraints* for machine learning algorithms - these are methods that allow one to constrain the predictions of a machine learner so that they comply with physical assumptions. Excellent examples for Gaussian process regression are [29–31] and will be discussed in greater detail later.

The final grouping of grey-box approaches mentioned here are physics-guided black-box learners. These are methods that use physical insight to attempt to improve model optimisation and include the construction of physics-guided loss functions and the use of physics-guided initialisation. These will not be discussed further in this chapter but see e.g. [32–35] for more details.

2.0.1 Grey-box models for SHM

The remainder of the chapter will showcase some of the work of the authors on developing physics-informed machine learning approaches for SHM tasks. The developments here fall in the domain of residual and hybrid models (Sections 3 and 5), and constrained machine learners (Section 4). In order to provide an overview, a variety of methods and results are presented, however, the implementation details given here are necessarily very brief and we refer readers to the referenced papers and our webpage³ for specific details and more in-depth analysis. Reflecting the philosophy discussed in the introduction section, the approaches presented generally incorporate simple physics-based models or assumptions and rely on the machine learner for enhanced explanatory power and flexibility (i.e. we are operating towards the blacker end of the scale).

The machine learning approach used in the work shown here will be Gaussian process (GP) regression throughout. GPs have been shown to be a powerful tool for regression tasks [36] and are becoming common in SHM applications (see for example [2, 37–40]). Their use here and throughout the work of the authors is due to their (semi)non-parametric nature, their ability to function with a small number of training points, and most importantly, the Bayesian framework within which they naturally work. The Gaussian process formulation provides a predictive *distribution* rather than a single prediction point, allowing confidence intervals to be calculated and uncertainty to be propagated forward into any following analysis (see [9] for

³ <https://drg-greybox.github.io/>

example). As the use of GPs is now quite common, their fundamental formulation will not be introduced here, but the mathematical machinery required is briefly summarised in the Appendix - we refer unfamiliar readers to [36].

In the first examples shown here, the use of priors in the Bayesian framework is exploited as an appropriate and intuitive means of incorporating physical insight into a machine learning algorithm. In later sections we consider the construction of constraints for GPs and, separately, their incorporation into a state-space formulation (this latter example relies more heavily on physics-based machinery than the other examples).

3 Be more Bayes

A Bayesian philosophy is one that employs evidence from data to update prior beliefs or assumptions, and has been widely adopted across disciplines, including SHM. However, most commonly, uninformative priors are utilised that do not reflect the knowledge that we have as engineers of the systems we are interested in modelling.

The formulation of a Gaussian process regression requires the selection of a mean and covariance function which form the prior process. The process is then conditioned with training data to provide a posterior mean and covariance as the model prediction. In the standard approach, no prior knowledge is assumed; a zero mean function is selected alongside a generic covariance function such as a squared-exponential or one from the Matérn class which provide a flexible process to fit to most data.

In this section we will first employ simple physical models as prior mean functions to a GP and show how they may improve the extrapolative capability of the model. This simple means of incorporating prior knowledge is equivalent to using a GP with a zero mean prior to model the difference between the measured data and the physical model prediction, and hence can be classed as a residual approach (see Section 2). At the end of this section, we will show how some knowledge of a system may be used to derive useful covariance functions in a regression setting.

3.1 Prior mean functions - residual modelling

3.1.1 Performance monitoring of a cable-stayed bridge

The Tamar bridge is a cable-supported suspension bridge connecting Saltash and Plymouth in the South West of England which has been monitored by the Vibration Engineering Section at the University of Exeter [41]. The interest here is in the development of a model to predict bridge deck deflections that can be used as a performance indicator (see [42, 43]). The variation in deck deflections are driven by a number of factors, including fluctuating temperature and loading from traffic

(which are included as inputs to the model). Figure 3a shows the regression target considered in this example, which is a longitudinal deflection. The monitoring period shown is from September (Autumn) to January (Winter). In this figure one can see short term fluctuations (daily) and a longer term trend which is seasonal and driven by the increased hogging of the bridge deck as the ambient temperature decreases into the winter months. To mimic the situation where only a limited period of monitoring data is available for the establishment of an SHM algorithm, data from the initial month of the monitoring period is used to establish a GP regression model for deflection prediction (see [44] for more details).

A GP prediction, using the standard approach of a zero mean prior, is shown in 3b. Here one can see exactly the behaviour that is expected; the model is able to predict the deck deflections well in and around the training period, but is unable to predict the deflections in colder periods towards the end of the time series. The confidence intervals widen to reflect that the inputs to the model towards the end of the period are different from those in the training set - this demonstrates the usefulness of the GP approach, as one knows to place less trust in the predictions from this period.

To formulate a grey-box model for this scenario, a physics-informed prior mean function is adopted that encodes the expected linear expansion behaviour of stay-cables with temperature [45]. Figure 3c shows the GP prediction with a linear prior mean function, where one can see a significant enhancement of predictive capability across the monitoring period. Where temperatures are at their lowest, the model predictions fall back on the prior mean function allowing some extrapolative capability. The prediction error is significantly smaller for the grey-box model in this case; the ‘black-box’ nMSE is 68.65, whereas the GP with the physics-informed mean function has an nMSE of 7.33.

3.1.2 Residual modelling for wave loading prediction

In this example, we follow a similar approach of adopting a physics-informed mean function, this time with a dynamic Gaussian process formulation, a GP-NARX [46], to enhance predictive capability for a wave loading assessment. The monitoring or prediction of the loads a structure experiences in service is an important ingredient for health assessment, particularly where one wishes to infer e.g. fatigue damage accrued.

The implementation of residual modelling is most effective where the assumptions and limitations of the white-box model are well understood. As a widely used method for wave loading prediction, Morison’s Equation [47] is employed here as a physics-informed mean function. This empirical law is known to simplify the behaviour of wave loading, not accounting for effects such as vortex shedding or other complex behaviours [48] and will typically have residual errors in the region of 20% [49]. Here we consider the addition of a data-based GP-NARX to a simplified version of Morison’s Equation in an attempt to account for these missing phenomena. The model used is:

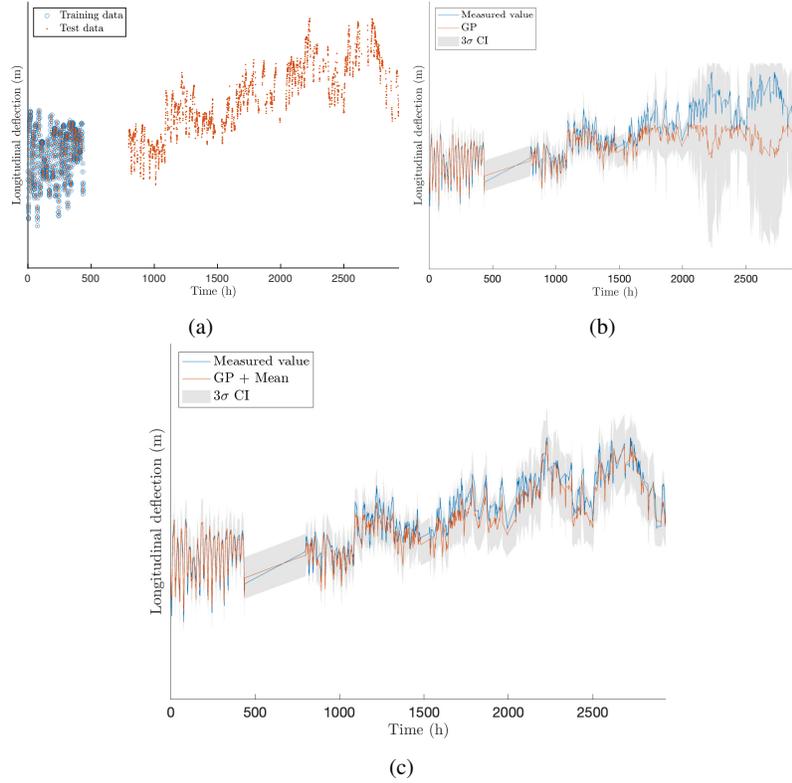


Fig. 3: Model for bridge deck deflections; (a) shows the training and test datasets for the GPs, (b) is a GP prediction with a zero mean function prior and (c) shows the prediction when a simple physics-informed mean function is incorporated. See [44] for more details.

$$y_t = \underbrace{C'_d U_t |U_t| + C'_m \dot{U}_t}_{\text{Morison's Equation}} + \underbrace{f([u_t, u_{t-1}, \dots, u_{t-l_u}, y_{t-1}, y_{t-2}, \dots, y_{t-l_y}])}_{\text{GP-NARX}} + \varepsilon \quad (3)$$

where y_t is the wave force, C'_d is the drag coefficient, C'_m is the inertia coefficient, U is the wave velocity, \dot{U} is the wave acceleration, $u_{t:t-l_u}$ are lagged exogenous inputs and $y_{t-1:t-l_y}$ are the lagged wave force, see [50] for more details (this paper also shows an example of an input augmentation model, where Morison's equation is used as an additional input to the GP-NARX).

The Christchurch bay dataset is used here as an example to demonstrate the approach [51]. To explore the generalisation capability with and without the physics-informed mean, different training sets for the GP-NARX are considered with increasing levels of coverage of the input space. A comparison of model errors (nMSE) with different training datasets is shown in Figure 4. The coverage level is indicated

as a percentage of the behaviour observed in the testing set that is also encountered in the training set [50].

As in the Tamar Bridge example, the model structure offers a significant improvement in extrapolation, where testing conditions are different from those in training dataset. Following this approach allows predictions to be informed by the prior mean in the absence of evidence from data. Clearly the prior specification is very important in this case and a misspecified prior could do more harm than good. Once again we advocate the use of simple and well founded physics-based models in an attempt to avoid this issue.

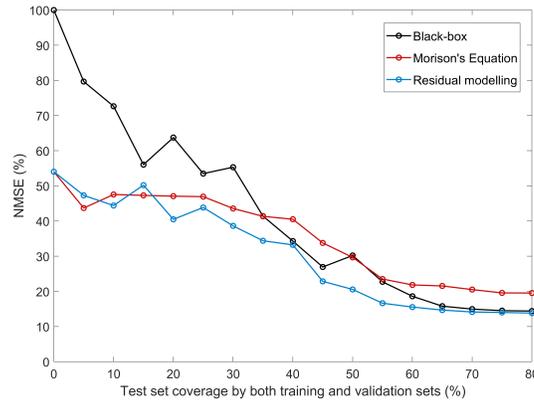


Fig. 4: A comparison of wave loading prediction model NMSEs vs test set coverage. Increasing coverage of the test set by the training and validation sets results in an increased level of model interpolation. See [50] for more details.

3.2 Physics-derived covariance functions

As discussed above, in a standard approach to GP regression, a generic covariance function such as a squared-exponential or one from the Matérn class is selected as a prior. In the posterior GP, the mean is a weighted sum of observations in the training set (see appendix), with the weightings provided by the covariance function and associated matrix. These commonly used functions encode that the covariance between points with similar inputs will be high and this allows the model to be data-driven in nature.

In the case where one has some knowledge of a process of interest, it is possible to derive a covariance function that reflects this. As an example, in [52], a composite covariance function is designed to reflect the characteristics of the guided waves being modelled.

For some stochastic processes, the (auto)covariance can be directly derived from the equation of motion of a system. An example relevant for vibration-based SHM is the single degree of freedom (SDOF) oscillator

$$m\ddot{y}(t) + c\dot{y}(t) + ky(t) = F(t) \quad (4)$$

with mass, damping and stiffness parameters, m , c , k respectively driven by a forcing process $F(t)$. In the case where the forcing is Gaussian white noise, the response Y is a Gaussian process with (auto)covariance

$$\phi_Y(\tau) = \mathbb{E}[Y(t_1)Y(t_2)] = \frac{\sigma^2}{4m^2\zeta\omega_n^3} e^{-\zeta\omega_n|\tau|} (\cos(\omega_d\tau) + \frac{\zeta\omega_n}{\omega_d} \sin(\omega_d|\tau|)) \quad (5)$$

where standard notation has been used; $\omega_n = \sqrt{k/m}$, the natural frequency, $\zeta = c/2\sqrt{km}$, the damping ratio, $\omega_d = \omega_n\sqrt{1-\zeta^2}$, the damped natural frequency. See [53] and also [54, 55].

This covariance function can be readily used in the regression context and provides a useful prior process for oscillatory systems with a response dominated by a single frequency. This form of covariance function can be described as *expressive* [56] and proves useful even when the equation of motion of the system of interest differs from an SDOF linear assumption.

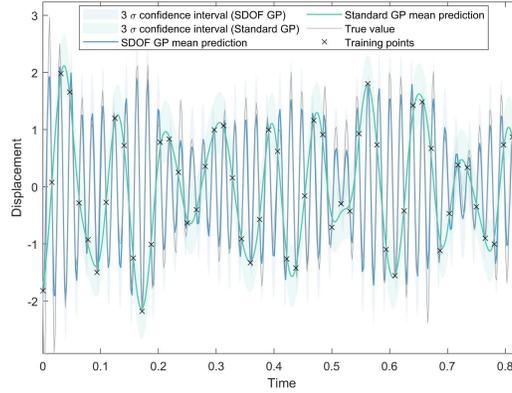


Fig. 5: A comparison of a GP regression using derived and squared-exponential covariance function. See [53] for more details.

Figure 5 shows an example of a GP regression for a system with a cubic nonlinearity. Here the linear prior provides an appropriate structure for the regression and is flexible enough to incorporate the nonlinearity in the response. This is a simulated example with the GP training data shown with crosses in the figure (every 8th point) - here the nMSE is 8.09. For comparison, a GP with a squared-exponential (SE) co-

variance function is established with the same training data. The SE process smooths through the data as expected ($nMSE=66.7$), whereas the derived covariance provides structure through the prior, resulting in good prediction during interpolation. The hyperparameters in the SDOF covariance function are physically interpretable, we are, therefore, able to guide their optimisation by providing the likely ranges for the system of interest. Here the benefit of being able to prescribe the likely frequency content of the system within the prior is clear and provides much advantage over the black-box approach - see [53] for more details.

4 Constrained Gaussian processes

In scenarios where one lacks significant knowledge of the governing equations and solutions, grey-box methods that tend towards the black end of the spectrum can be particularly useful. An example of such approaches are constrained machine learners, which, very generally, aim to embed physical constraints into the learning procedure such that predictions made by the black-box model then adhere to these constraints.

In the context of Gaussian process regression, there are a number of ways of constraining predictions, the simplest of which is to include data from boundaries within the model training. Other methods rely on applying constraints to the covariance function in a multiple output setting (see e.g. [29–31] and [57] where we employ derivative boundaries for beam deflection predictions).

Here we show an example of building known boundaries (geometry) into a GP regression via a sparse approximation of the covariance function. The approximation method relies on an eigendecomposition of the Laplace operator of a fixed domain [58]:

$$k(\mathbf{x}, \mathbf{x}') \approx \sum_i^m S(\sqrt{\lambda_i}) \phi_i(\mathbf{x}) \phi_i(\mathbf{x}'), \quad (6)$$

with ϕ_i and λ_i the eigenfunctions and values, and S the spectral density of the covariance function. If one chooses the fixed domain to reflect the geometry of the problem of interest, then inference with this model is appropriately bounded (see [59, 60] for more details).

As an example here we employ constraints for a crack localisation problem via measurement of Acoustic Emission (AE). The localisation approach taken is to use artificial source excitations and an interpolating

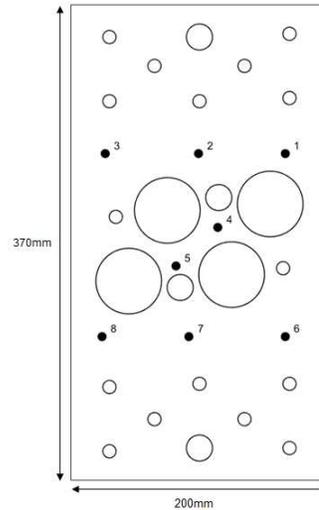


Fig. 6: Test structure schematic with sensor locations, recreated from [61].

GP to provide a map of the differences in times of arrival (ΔT) of AE sources to fixed sensor pairings across the surface of the structure [61, 62]. Once constructed, the map can be used to assess the most likely location of any new AE sources. The bounded GP approximation allows one to build in the geometry of the structure under consideration.

To investigate the predictive capability of the constrained GP, a case study using a plate with a number of holes in it is adopted. The holes, as shown in Figure 6, provide complexity to the modelling challenge, introducing several complex phenomena such as wave mode conversion and signal reflection. Depending on the location of the source and sensor, the holes may also shield a direct propagation path to the receiver [61], adding further complication.

Neumann boundaries are imposed here around each hole and at the edge of the plate. To compare the performance of the standard and bounded GPs, differing amounts/coverage of artificial source excitations were used for model training. The initial characterisation of a structure via artificial source excitation can be expensive and time consuming, for structures in operation it may also be infeasible to access all areas/components. To mimic the scenario where it is not possible to collect artificial source excitations across a whole structure, here we restrict the training grid to excitation points in the middle of the plate. Figure 7 compares the performance of the standard and bounded GPs with training sets of varying grid densities. For each training set, the prediction error (nMSE) on the test set is averaged across every sensor pair (there are 8 sensors).

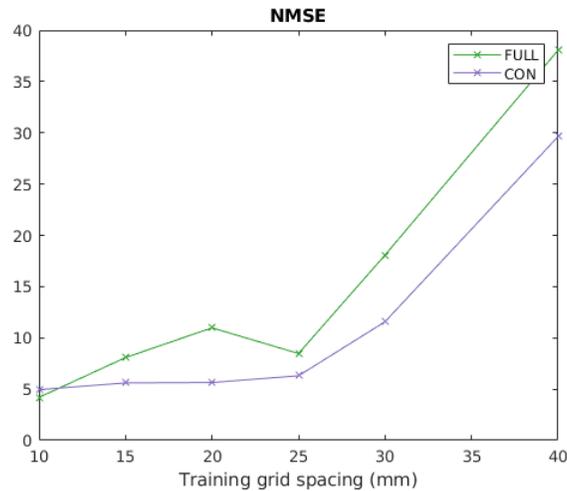


Fig. 7: Comparison between models errors for standard and bounded GPs for AE source localisation study. The nMSE is averaged across all sensor pair models for each training set considered.

From Figure 7 one can see that as the training set size reduces, the constrained GP consistently outperforms the standard full GP. This is particularly encouraging as the bounded GP remains a sparse approximation. As is consistent with our earlier observations, the inbuilt physical insight aids inference where training data are fewer. Figure 8 shows the difference in prediction error across the plate for the standard and bounded GPs for the 20mm spacing training case and a single sensor pairing. In this case, the squared error of the full GP is subtracted from the squared error of the constrained GP, i.e. positive values indicate a larger error in the full GP, whilst a negative value expresses a larger error in the constrained GP.

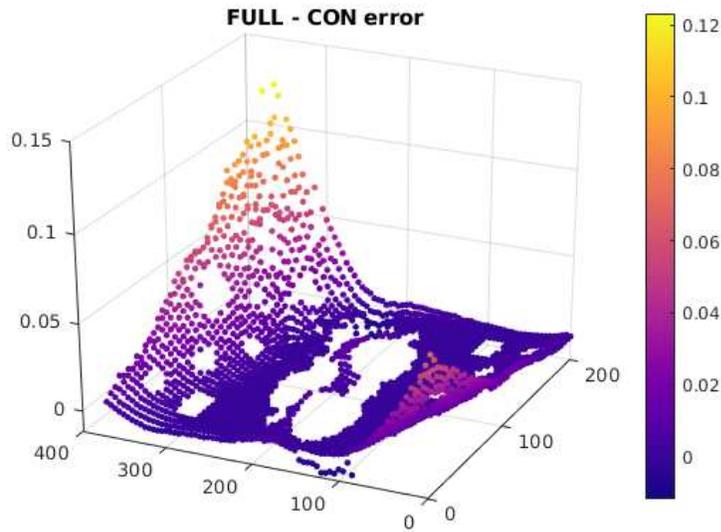


Fig. 8: Mapping of the difference between the full GP squared error and the constrained GP squared error across the test set for sensor pair 4-8.

The figure highlights the locations on the plate where the constrained GP more accurately predicts the true ΔT values. As expected, the locations at which this effect is most prominent are those that move further away from the training points, and particularly towards the extremities of the domain. At these locations, it is clear that the additional physical insight provided by the constrained GP is able to enhance the predictive performance in comparison to the pure black-box model.

5 Gaussian processes in a state-space approach

One of the canonical forms for dynamic models, in structural mechanical systems and beyond, is the state-space representation of the behaviour of interest.

In the context of this work, the state space model (SSM) is considered to be a probabilistic object defined by two key probability densities; a transition density $p(\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{u}_t)$ and observation density $p(\mathbf{y}_t | \mathbf{x}_t, \mathbf{u}_t)$. The transition density relates the hidden states at a given time \mathbf{x}_{t+1} to their previous possible values⁴, $\mathbf{x}_{0:t}$ and previous external inputs to the system, e.g. forcing, $\mathbf{u}_{0:t}$. The observation model relates available measurements \mathbf{y}_t to the hidden states \mathbf{x}_t , which may also be dependent on the external inputs at that time \mathbf{u}_t .

The state space formulation can be used to properly account for measurement noise (filtering and smoothing), and is commonly used for parameter estimation (the well-known Kalman filter is a closed form solution for linear and Gaussian systems). Use of the state-space models as a grey-box formulation in this setting is common within the control community [63, 64].

Here, we are interested in the case where we only have partial knowledge of a system - this could take the form of missing or incorrect physics in the equations of motion, or could be a lack of access to key measurements such as the force a system undergoes. In Section 3 we considered a GP-NARX formulation for wave loading prediction, with the ultimate aim of informing a fatigue assessment. The state-space formulation shown here offers an alternative means for load estimation which simultaneously provides parameter and state estimation in a Bayesian setting.

Joint input-state and input-state-parameter problems have seen growing interest in recent years, see, for example [65–69]. The approach shown here is one that considers a representation of a Gaussian process within a state space formulation to model the unknown forcing (following [70, 71]). This is achieved by deriving the transfer function of a Matérn kernel (via its spectral density), which provides a flexible model component to account for the unmeasured behaviour. Inference over the state space model is via Markov Chain Monte Carlo to provide distributions for parameter and hyperparameter estimations.

Figure 9 shows an example of force recovery for a simulated multi-degree of freedom system excited by a forcing time history from the Christchurch bay example discussed in Section 3. Here one can see that the force has been accurately inferred, the nMSE in this case is 1.15. For more details and analysis see [40].

The inference problem becomes significantly more challenging for nonlinear systems, and even more so if our knowledge of that nonlinearity is incomplete. Recently [72] has attempted this extension for the input-state estimation case for a known nonlinearity. The difficulty in inference is met there by employing a methodology based on Sequential Monte Carlo, specifically Particle Gibbs with Ancestor Sampling, to allow recovery of the states and the hyperparameters of the GP.

⁴ The notation *subscript* $a : b$ is used to denote values in that range inclusively, e.g. $\mathbf{x}_{0:t}$ is the value of the states \mathbf{x} at all times from $t = 0$ to $t = t$.

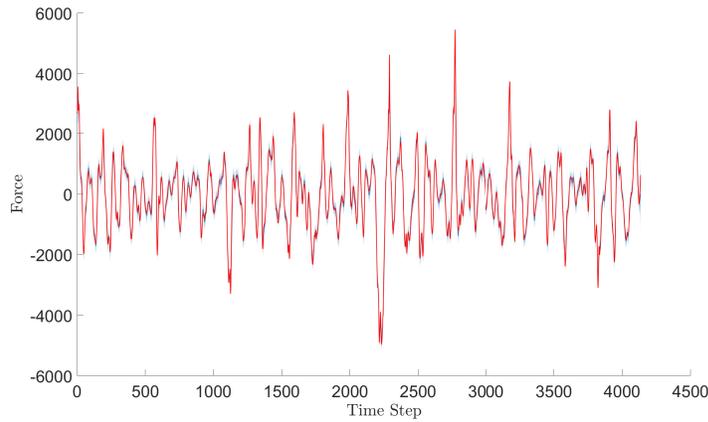


Fig. 9: Force estimation for MDOF simulation under Christchurch bay loading time history[40]

In the face of an unknown nonlinearity this framework may also be employed, where the GP may be used to account for missing behaviours from the assumed equations of motion. In [73], it is shown how this approach can be applied to a Duffing oscillator to learn the unknown cubic component of the model in a Bayesian manner without requiring prior knowledge of the nonlinear function. There are two particular advantages to this approach, the first is that it allows nonlinear system identification with a linear model, the Kalman filter and RTS smoother, which has significant computational advantages and arguably makes “fully Bayesian” inference industrially feasible. The second benefit is that it allows a user to apply very weak prior knowledge about the nonlinearity; in the language of grey-box models, there is a strong white box component (the second order linear system) but the form of the nonlinearity is the very flexible nonparametric GP. Contrast this approach with the purely black-box alternative of the GP-SSM, see [74], which suffers from significant nonidentifiability and computational challenges.

6 Conclusions

This chapter has introduced and demonstrated physics-informed machine learning methods suitable for SHM problems and inference in structural dynamics more generally. The methods allow the embedding of one’s physical insight of a structure or system into a data-driven assessment. The resulting models have proven to be particularly useful in situations where training data are not available across the operational envelope - a common occurrence in structural monitoring campaigns.

The Bayesian approach adopted in Section 3 allows predictions to fall back on a prior physical model in the absence of evidence from data. This pragmatic approach proved useful in the examples shown here but does rely on trusting the physical model in extrapolation. The ability to constrain the Gaussian process prior to known boundary conditions shown in Section 4 requires less physical insight and gives both an improved modelling performance, as well as providing the guarantee that predictions made adhere to known underlying physical laws of the system under consideration. At the whiter end of the spectrum, the state-space examples discussed in Section 5 provide a principled means of inference over structures with unknown forcing or nonlinearities. Some examples of where the presented methodology may be of benefit could include better understanding of fatigue damage accrual and parameter identification for e.g. novelty/damage detection.

As well as providing an enhanced predictive capability, the models introduced here have the benefit of being more readily interpretable than their purely black-box counterparts. In the past, a barrier to the uptake of SHM technology has been the lack of trust owners and operators have in so-called black-box models. Perhaps naturally, there is a hesitancy to adopt algorithms not derived from physics-based models, but this may also be due, in part, to their misuse in the past. We hope that this will be ameliorated by more interpretable models which also have the benefit of being more easily optimised (Section 3 shows an example where the hyperparameters in the GP regression take on physical meaning).

Physics-informed machine learning is rapidly becoming a popular research field in its own right, with many promising results and avenues for investigation. This review paper [35] currently on arXiv has 300 references largely populated by papers from the last two years. It is likely that many of the emerging methods will prove useful in SHM. The work here has focussed on a Gaussian process framework, clearly the use of neural networks provide an alternative grey-box route, as these are also commonly used in our field. We look forward to seeing how these may be adopted for SHM tasks.

6.0.1 Acknowledgements

We would like to thank Keith Worden for his general support and also particularly for provision of the Christchurch Bay data and support of the state-space work. Additionally, thanks is offered to James Hensman, Mark Eaton, Robin Mills, Gareth Pierce and Keith Worden for their work in acquiring the AE data set used here. We would like to thank Ki-Young Koo and James Brownjohn in the Vibration Engineering Section at the University of Exeter for provision of the data from the Tamar Bridge. Thanks also to Steve Reed, formerly of DSTL for the provision of the Tucano data. Finally, the authors would like to acknowledge the support of the EPSRC, particularly through grant reference number EP/S001565/1, and Ramboll Energy for their support of SG and DP.

7 Gaussian Process Regression

Here we follow the notation used in [36]; $k(\mathbf{x}_p, \mathbf{x}_q)$ defines a covariance matrix K_{pq} , with elements evaluated at the points \mathbf{x}_p and \mathbf{x}_q , where \mathbf{x}_i may be multivariate.

Assuming a zero-mean function, the joint Gaussian distribution between measurements/observations \mathbf{y} with inputs X and unknown/testing targets \mathbf{y}^* with inputs X^* is

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{y}^* \end{bmatrix} \sim \mathcal{N}\left(0, \begin{bmatrix} K(X, X) + \sigma_n^2 I & K(X, X^*) \\ K(X^*, X) & K(X^*, X^*) \end{bmatrix}\right) \quad (7)$$

The distribution of the testing targets \mathbf{y}^* conditioned on the training data (which is what we use for prediction) is also Gaussian:

$$\begin{aligned} \mathbf{y}^* | X_*, X, \mathbf{y} &\sim \mathcal{N}(K(X^*, X)(K(X, X) + \sigma_n^2 I)^{-1} \mathbf{y}, \\ &K(X^*, X^*) - K(X^*, X)(K(X, X) + \sigma_n^2 I)^{-1} K(X, X^*)) \end{aligned} \quad (8)$$

See [36] for the derivation. The mean and covariance here are that of the posterior Gaussian process. In this work covariance function hyperparameters are sought by maximising the marginal likelihood of the predictions

$$\log p(\mathbf{y}|X, \theta) = -\frac{1}{2} \mathbf{y}^T K^{-1} \mathbf{y} - \frac{1}{2} \log |K| - \frac{n}{2} \log 2\pi \quad (9)$$

via a particle swarm optimisation⁵ [75].

References

- [1] Charles R Farrar and Keith Worden. An introduction to structural health monitoring. *New Trends in Vibration Based Structural Health Monitoring*, pages 1–17, 2010.
- [2] Geoffrey Holmes, Pia Sartor, Stephen Reed, Paul Southern, Keith Worden, and Elizabeth Cross. Prediction of landing gear loads using machine learning techniques. *Structural Health Monitoring*, 15(5):568–582, 2016.
- [3] Hoon Sohn, Charles R Farrar, Francois M Hemez, Devin D Shunk, Daniel W Stinemat, Brett R Nadler, and Jerry J Czarnecki. A review of structural health monitoring literature: 1996–2001. *Los Alamos National Laboratory, USA*, 1, 2003.
- [4] Ellen Simoen, Guido De Roeck, and Geert Lombaert. Dealing with uncertainty in model updating for damage assessment: A review. *Mechanical Systems and Signal Processing*, 56:123–149, 2015.
- [5] Paul Gardner. *On novel approaches to model-based structural health monitoring*. PhD thesis, University of Sheffield, 2018.

⁵ Yes, we could be more Bayes here

- [6] Paul Gardner, Charles Lord, and Robert J Barthorpe. A unifying framework for probabilistic validation metrics. *Journal of Verification, Validation and Uncertainty Quantification*, 4(3), 2019.
- [7] Paul Gardner, Charles Lord, and RJ Barthorpe. Bayesian history matching for forward model-driven structural health monitoring. In *Model Validation and Uncertainty Quantification, Volume 3*, pages 175–183. Springer, 2019.
- [8] Ramon Fuentes, Elizabeth Cross, Andrew Halfpenny, Keith Worden, and Robert J Barthorpe. Aircraft parametric structural load monitoring using gaussian process regression. In *Proceedings of the European Workshop on Structural Health Monitoring 2014, Nantes*, 2014.
- [9] S. J. Gibson, T. J. Rogers, and E. J. Cross. Data-driven strain prediction models and fatigue damage accumulation. In *Proceedings of the 29th International Conference on Noise and Vibration Engineering (ISMA 2020)*, 2020.
- [10] Robert James Barthorpe. *On model-and data-based approaches to structural health monitoring*. PhD thesis, University of Sheffield, 2010.
- [11] Charles R Farrar, Keith Worden, Michael D Todd, Gyuhae Park, Jonathon Nichols, Douglas E Adams, Matthew T Bement, and Kevin Farinholt. Nonlinear system identification for damage detection. Technical report, Los Alamos National Laboratory (LANL), Los Alamos, NM, 2007.
- [12] Gaetan Kerschen, Keith Worden, Alexander F Vakakis, and Jean-Claude Golinval. Past, present and future of nonlinear system identification in structural dynamics. *Mechanical systems and signal processing*, 20(3):505–592, 2006.
- [13] Jeremy Oakley and Anthony O’Hagan. Bayesian inference for the uncertainty distribution of computer model outputs. *Biometrika*, 89(4):769–784, 2002.
- [14] Patrick E Leser, Jacob D Hochhalter, James E Warner, John A Newman, William P Leser, Paul A Wawrzynek, and Fuh-Gwo Yuan. Probabilistic fatigue damage prognosis using surrogate models trained via three-dimensional finite element analysis. *Structural Health Monitoring*, 16(3):291–308, 2017.
- [15] Tom Beucler, Michael Pritchard, Stephan Rasp, Jordan Ott, Pierre Baldi, and Pierre Gentine. Enforcing analytic constraints in neural-networks emulating physical systems. *arXiv preprint arXiv:1909.00912*, 2019.
- [16] Björn Sohlberg and Elling W Jacobsen. Grey box modelling—branches and experiences. *IFAC Proceedings Volumes*, 41(2):11415–11420, 2008.
- [17] Bjorn Sohlberg. *Supervision and control for industrial processes: using grey box models, predictive control and fault detection methods*. Springer Science & Business Media, 2012.
- [18] R Fuentes, R Nayek, P Gardner, N Dervilis, T Rogers, K Worden, and EJ Cross. Equation discovery for nonlinear dynamical systems: A bayesian viewpoint. *Mechanical Systems and Signal Processing*, 154:107528, 2021.
- [19] Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the national academy of sciences*, 113(15):3932–3937, 2016.
- [20] Jean-Philippe Noël and Johan Schoukens. Grey-box state-space identification of nonlinear mechanical vibrations. *International Journal of Control*, 91(5): 1118–1139, 2018.

- [21] Marc C Kennedy and Anthony O’Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425–464, 2001.
- [22] Chon Lok Lei, Sanmitra Ghosh, Dominic G Whittaker, Yasser Aboelkassem, Kylie A Beattie, Chris D Cantwell, Tammo Delhaas, Charles Houston, Gustavo Montes Novaes, Alexander V Panfilov, et al. Considering discrepancy when calibrating a mechanistic electrophysiology model. *Philosophical Transactions of the Royal Society A*, 378(2173):20190349, 2020.
- [23] J Brynjarsdottir and A O Hagan. Learning about physical parameters: The importance of model discrepancy. *Inverse Problems*, 30(11):114007, 2014.
- [24] K Worden, CX Wong, U Parlitz, A Hornstein, D Engster, Tegoeh Tjahjowidodo, Farid Al-Bender, DD Rizos, and SD Fassois. Identification of pre-sliding and sliding friction dynamics: Grey box and black-box models. *Mechanical systems and signal Processing*, 21(1):514–534, 2007.
- [25] O. Arda Vanli and Sungmoon Jung. Statistical updating of finite element model with lamb wave sensing data for damage detection problems. *Mechanical Systems and Signal Processing*, 42:137 – 151, 2014.
- [26] H.-P. Wan and W.-X. Ren. A residual-based Gaussian process model framework for finite element model updating. *Computers and Structures*, 156:149–159, 2015.
- [27] TJ Rogers, GR Holmes, EJ Cross, and K Worden. On a grey box modelling framework for nonlinear system identification. In *Special Topics in Structural Dynamics, Volume 6*, pages 167–178. Springer, 2017.
- [28] K Worden, RJ Barthorpe, EJ Cross, N Dervilis, GR Holmes, G Manson, and TJ Rogers. On evolutionary system identification with applications to nonlinear benchmarks. *Mechanical Systems and Signal Processing*, 112:194–232, 2018.
- [29] Arno Solin, Manon Kok, Niklas Wahlström, Thomas B Schön, and Simo Särkkä. Modeling and interpolation of the ambient magnetic field by gaussian processes. *IEEE Transactions on robotics*, 34(4):1112–1127, 2018.
- [30] Niklas Wahlström, Manon Kok, Thomas B Schön, and Fredrik Gustafsson. Modeling magnetic fields using gaussian processes. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 3522–3526. IEEE, 2013.
- [31] Carl Jidling, Johannes Hendriks, Niklas Wahlström, Alexander Gregg, Thomas B Schön, Christopher Wensrich, and Adrian Wills. Probabilistic modelling and reconstruction of strain. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 436:141–155, 2018.
- [32] Anuj Karpatne, William Watkins, Jordan Read, and Vipin Kumar. Physics-guided neural networks (pgnn): An application in lake temperature modeling. *arXiv preprint arXiv:1710.11431*, 2017.
- [33] Maziar Raissi, Zhicheng Wang, Michael S Triantafyllou, and George Em Karniadakis. Deep learning of vortex-induced vibrations. *Journal of Fluid Mechanics*, 861:119–137, 2019.

- [34] Anuj Karpatne, Gowtham Atluri, James H Faghmous, Michael Steinbach, Arindam Banerjee, Auroop Ganguly, Shashi Shekhar, Nagiza Samatova, and Vipin Kumar. Theory-guided data science: A new paradigm for scientific discovery from data. *IEEE Transactions on knowledge and data engineering*, 29(10):2318–2331, 2017.
- [35] Jared Willard, Xiaowei Jia, Shaoming Xu, Michael Steinbach, and Vipin Kumar. Integrating physics-based modeling with machine learning: A survey. *arXiv preprint arXiv:2003.04919*, 2020.
- [36] C.E. Rasmussen and C.K.I. Williams. *Gaussian Processes for Machine Learning*, volume 38. The MIT Press, Cambridge, MA, USA, 2006.
- [37] Lawrence A Bull, Paul Gardner, Timothy J Rogers, Elizabeth J Cross, Nikolaos Dervilis, and Keith Worden. Probabilistic inference for structural health monitoring: New modes of learning from data. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*, 7(1):03120003, 2020.
- [38] Hua-Ping Wan and Yi-Qing Ni. Bayesian modeling approach for forecast of structural stress response using structural health monitoring data. *Journal of Structural Engineering*, 144(9):04018130, 2018.
- [39] Jyrki Kullaa. Distinguishing between sensor fault, structural damage, and environmental or operational effects in structural health monitoring. *Mechanical Systems and Signal Processing*, 25(8):2976–2989, 2011.
- [40] TJ Rogers, K Worden, and EJ Cross. On the application of Gaussian process latent force models for joint input-state-parameter estimation: With a view to Bayesian operational identification. *Mechanical Systems and Signal Processing*, 140:106580, 2020.
- [41] Ki-Young Koo, JMW Brownjohn, DI List, and R Cole. Structural health monitoring of the tamar suspension bridge. *Structural Control and Health Monitoring*, 20(4):609–625, 2013.
- [42] E Cross, K Worden, Ki Young Koo, and James Brownjohn. *Filtering environmental load effects to enhance novelty detection on cable-supported bridge performance*, pages 745–752. Bridge Maintenance, Safety and Management. CRC Press, 06 2012. doi: 10.1201/b12352-101.
- [43] Elizabeth Cross. *On structural health monitoring in changing environmental and operational conditions*. PhD thesis, University of Sheffield, 2012.
- [44] S. Zhang, T. J. Rogers, and E. J. Cross. Gaussian process based grey-box modelling for SHM of structures under fluctuating environmental conditions. In *Proceedings of 10th European Workshop on Structural Health Monitoring (EWSHM 2020)*, 2020.
- [45] Robert Westgate. *Environmental effects on a suspension bridge’s performance*. PhD thesis, University of Sheffield, 2012.
- [46] K Worden, WE Becker, TJ Rogers, and EJ Cross. On the confidence bounds of gaussian process narx models and their higher-order frequency response functions. *Mechanical Systems and Signal Processing*, 104:188–223, 2018.
- [47] J. R. Morison, Joseph W Johnson, and Samuel A. Schaaf. The force exerted by surface waves on piles. *Petroleum Transactions*, 189:149–157, 1950.

- [48] E. Guilmineau and P. Queutey. A numerical simulation of vortex shedding from an oscillating circular cylinder. *Journal of Fluids and Structures*, 16: 773–794, 2002.
- [49] A. M. M. Wood and C. A. Fleming. *Coastal Hydraulics*. Macmillan Education, Limited, 1981.
- [50] Grey-box models for wave loading prediction. *Mechanical Systems and Signal Processing*, 159:107741, 2021. ISSN 0888-3270. doi: <https://doi.org/10.1016/j.ymssp.2021.107741>.
- [51] G Najafian, RG Tickell, R Burrows, and JR Bishop. The uk christchurch bay compliant cylinder project: analysis and interpretation of morison wave force and response data. *Applied Ocean Research*, 22(3):129–153, 2000.
- [52] Marcus Haywood-Alexander, Nikolaos Dervilis, Keith Worden, Elizabeth J Cross, Robin S Mills, and Timothy J Rogers. Structured machine learning tools for modelling characteristics of guided waves. *arXiv preprint arXiv:2101.01506*, 2021.
- [53] E. J. Cross and T. J. Rogers. Physics-derived covariance functions for machine learning in structural dynamics. In *Submitted to SYSID 2021*, 2021.
- [54] Athanasios Papoulis. *Probability, random variables, and stochastic processes*. McGraw-Hill Education, 1965.
- [55] TK Caughey. Nonlinear theory of random vibrations. In *Advances in applied mechanics*, volume 11, pages 209–253. Elsevier, 1971.
- [56] Andrew Wilson and Ryan Adams. Gaussian process kernels for pattern discovery and extrapolation. In *International Conference on Machine Learning*, pages 1067–1075, 2013.
- [57] Cross E.J., Rogers T.J., and Gibbons T.J. Grey-box modelling for structural health monitoring: Physical constraints on machine learning algorithms. In *Structural Health Monitoring 2019: Enabling Intelligent Life-Cycle Health Management for Industry Internet of Things (IIOT) - Proceedings of the 12th International Workshop on Structural Health Monitoring*, volume 2, pages 2136–2145, 1 2019. ISBN 9781605956015.
- [58] Arno Solin and Simo Särkkä. Hilbert space methods for reduced-rank Gaussian process regression. *Statistics and Computing*, 30(2):419–446, 2020.
- [59] M. R. Jones, T. J. Rogers, P. A. Gardner, and E. J. Cross. Constraining Gaussian processes for grey-box acoustic emission source localisation. In *Proceedings of the 29th International Conference on Noise and Vibration Engineering (ISMA 2020)*, 2020.
- [60] Matthew R Jones, Timothy J Rogers, Ian E Martinez, and Elizabeth J Cross. Bayesian localisation of acoustic emission sources for wind turbine bearings. In *Health Monitoring of Structural and Biological Systems XV*, volume 11593, page 115932D. International Society for Optics and Photonics, 2021.
- [61] James Hensman, Robin Mills, SG Pierce, Keith Worden, and M Eaton. Locating acoustic emission sources in complex structures using Gaussian processes. *Mechanical Systems and Signal Processing*, 24(1):211–223, 2010.

- [62] Matthew R Jones, Tim J Rogers, Keith Worden, and Elizabeth J Cross. A Bayesian methodology for localising acoustic emission sources in complex structures. *arXiv preprint arXiv:2012.11058*, 2020.
- [63] Niels Rode Kristensen, Henrik Madsen, and Sten Bay Jørgensen. Parameter estimation in stochastic grey-box models. *Automatica*, 40(2):225–237, 2004.
- [64] Herbert JAF Tulleken. Grey-box modelling and identification using physical knowledge and bayesian techniques. *Automatica*, 29(2):285–308, 1993.
- [65] E Lourens, Costas Papadimitriou, Steven Gillijns, Edwin Reynders, Guido De Roeck, and Geert Lombaert. Joint input-response estimation for structural systems based on reduced-order models and vibration data from a limited number of sensors. *Mechanical Systems and Signal Processing*, 29:310–327, 2012.
- [66] Saeed Eftekhari Azam, Eleni Chatzi, and Costas Papadimitriou. A dual kalman filter approach for state estimation via output-only acceleration measurements. *Mechanical Systems and Signal Processing*, 60:866–886, 2015.
- [67] Frank Naets, Jan Croes, and Wim Desmet. An online coupled state/input/parameter estimation approach for structural dynamics. *Computer Methods in Applied Mechanics and Engineering*, 283:1167–1188, 2015.
- [68] K Maes, Freddie Karlsson, and G Lombaert. Tracking of inputs, states and parameters of linear structural dynamic systems. *Mechanical Systems and Signal Processing*, 130:755–775, 2019.
- [69] Vasilis K Dertimanis, EN Chatzi, S Eftekhari Azam, and Costas Papadimitriou. Input-state-parameter estimation of structural systems from limited output information. *Mechanical Systems and Signal Processing*, 126:711–746, 2019.
- [70] Mauricio Alvarez, David Luengo, and Neil Lawrence. Latent force models. In *Artificial Intelligence and Statistics*, pages 9–16, 2009.
- [71] Jouni Hartikainen and Simo Sarkka. Sequential inference for latent force models. *arXiv preprint arXiv:1202.3730*, 2012.
- [72] Timothy J Rogers, Keith Worden, and Elizabeth J Cross. Bayesian joint input-state estimation for nonlinear systems. *Vibration*, 3(3):281–303, 2020.
- [73] T. Friis, R. Brincker, and T.J. Rogers. On the application of Gaussian process latent force models for Bayesian identification of the Duffing system. In *Proceedings of ISMA 2020 - International Conference on Noise and Vibration Engineering and USD 2020 - International Conference on Uncertainty in Structural Dynamics*, 2020.
- [74] Roger Frigola. *Bayesian time series learning with Gaussian processes*. PhD thesis, University of Cambridge, 2015.
- [75] Timothy J Rogers. *Towards Bayesian System Identification: With Application to SHM of Offshore Structures*. PhD thesis, University of Sheffield, 2019.