



This is a repository copy of *Structured machine learning tools for modelling characteristics of guided waves*.

White Rose Research Online URL for this paper:
<https://eprints.whiterose.ac.uk/169792/>

Version: Accepted Version

Article:

Haywood-Alexander, M., Dervilis, N. orcid.org/0000-0002-5712-7323, Worden, K. orcid.org/0000-0002-1035-238X et al. (3 more authors) (2021) Structured machine learning tools for modelling characteristics of guided waves. *Mechanical Systems and Signal Processing*, 156. 107628. ISSN 0888-3270

<https://doi.org/10.1016/j.ymssp.2021.107628>

Article available under the terms of the CC-BY-NC-ND licence (<https://creativecommons.org/licenses/by-nc-nd/4.0/>).

Reuse

This article is distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs (CC BY-NC-ND) licence. This licence only allows you to download this work and share it with others as long as you credit the authors, but you can't change the article in any way or use it commercially. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

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



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

Structured Machine Learning Tools for Modelling Characteristics of Guided Waves

Marcus Haywood-Alexander^{a,*}, Nikolaos Dervilis^a, Keith Worden^a, Elizabeth
J. Cross^a, Robin S. Mills^b, Timothy J. Rogers^a

^a*Dynamics Research Group, Department of Mechanical Engineering, The University of
Sheffield, Mappin Building, Mappin Street, Sheffield, S1 3JD, United Kingdom*

^b*Laboratory for Verification and Validation (LVV), Europa Avenue, Sheffield S9 1ZA,
United Kingdom*

Abstract

The use of ultrasonic guided waves to probe the materials/structures for damage continues to increase in popularity for non-destructive evaluation (NDE) and structural health monitoring (SHM). The use of high-frequency waves such as these offers an advantage over low-frequency methods from their ability to detect damage on a smaller scale. However, in order to assess damage in a structure, and implement any NDE or SHM tool, knowledge of the behaviour of a guided wave throughout the material/structure is important (especially when designing sensor placement for SHM systems). Determining this behaviour is extremely difficult in complex materials, such as fibre-matrix composites, where unique phenomena such as continuous mode conversion takes place. This paper introduces a novel method for modelling the feature-space of guided waves in a composite material. This technique is based on a data-driven model, where prior physical knowledge can be used to create structured machine learning tools; where constraints are applied to provide said structure. The method shown makes use of Gaussian processes, a full Bayesian analysis tool, and in this paper it is shown how physical knowledge of the guided waves can be utilised in modelling using an ML tool. This paper shows that through careful consideration when applying machine learning techniques, more robust models can be generated which offer advantages such as extrapolation ability and physical interpretation.

Keywords: guided waves, feature space modelling, machine learning,
structural health monitoring, composite plate waves

1. Introduction

- 2 In engineering applications, the use of complex materials, such as composite
or porous materials can offer benefits thanks to their high strength-to-weight
4 ratio [1], amongst other advantages. However, with the increased usage of such
materials, also comes the ability of building larger structures; an example being
6 the ever-increasing size of wind turbines due to the use of glass-fibre reinforced

*Corresponding Author

Email address: `mhaywood-alexander1@sheffield.ac.uk` (Marcus Haywood-Alexander)

polymer for the blades [2]. With larger structures also comes an increasing need
8 for robust structural health monitoring systems in order to prolong life and
reduce costs associated with maintenance, down time and repair of large-scale
10 structures [3]. There are numerous strategies to perform structural health
monitoring, many of which involve the use of frequency-based behaviour of
12 the system [4]. Increasingly, high-frequency methods are being used, as they
have an ability to detect defects on a smaller scale thanks to the analogy with
14 diffraction theory, where a change in waves only occurs when they pass through
diffractors of small enough size relative to their wavelength. One such example of
16 these is to use *ultrasonic guided waves* (UGWs), where a guided wave is induced
within structure which acts as a wave-guide, and analysis of the ‘wave-packet’
18 as it arrives in certain locations can give indications of inhomogeneities in the
material.

20 There have been many advances in the modelling of composite materials,
such as in terms of the fatigue damage [5]; a comparative review of state-of-the-art
22 modelling methodologies for damage in composites has been made by Orifici *et al.*
[6], in which they discuss many issues such as length scales and implicit
24 modelling. Time-space modelling of guided waves can be solved analytically
for isotropic materials [7] or using numerical methods for layered composites
26 [8]; however, the complexity of these calculations becomes substantial when
modelling the interaction with damage or for a fibre-composite. One primary
28 issue with composites is the phenomena of *continuous mode conversion* (CMC)
[9, 10]. This can be conceptualised as secondary guide behaviour of the fibres
30 within the material, which for quasi-isotropic materials, with increased randomness
of lamina orientation, it may be possible to model with a set lay-up configuration
32 [11].

Along with the increase in available computing power, recent years have
34 shown a continuous increase in adoption of machine learning (ML) techniques
applied to engineering modelling problems in order to overcome barriers in
36 physic-based modelling, though not always successfully [12]. Many of these
methods are applied to a dataset to generate a model of certain behaviour;
38 however, this can limit applications. It is only reasonable to assume that a model
fits the *specific* material/structure, and generalisation cannot be implemented,
40 so changes in design or environment would require retraining of the model.

In this paper, a series of structured machine learning tools are presented in
42 which physical knowledge is embedded by differing means. It is important to
note this does not generate a directly interpretable model as material properties
44 are not introduced or extracted from the methodology. Instead, belief is embedded
from prior knowledge of the physics which control guided wave features, creating
46 structured tools that are constrained by this physical knowledge.

Using this novel method, study is done to investigate a novel path of understanding
48 and predicting some behaviour characteristics of guided waves by utilising enhanced
machine learning tools that capture the uncertainty of modelling the attenuation
of such waves (important when designing ultrasonic monitoring systems on large
50 scale structures like wind turbine blades). This paper focusses on one guided
wave feature, the maximum amplitude of the Hilbert envelope of a wave-packet
52 and although this paper focusses on features like energy attenuation, there is
nothing to stop the implementation of this new view of modelling on other
54 guided waves features.

56 This paper will discuss, briefly, the pertinent physics of guided waves in

Section 2 and then introduce the experimental setup used to demonstrate the methodology in Section 3. The main contribution of this paper is presented in Section 4; in this section, increasingly sophisticated models are generated which show how physical knowledge of the guided waves can help a machine learning strategy. This discussion begins with a purely physical one-dimensional attention model which is then extended to the two-dimensional case for modelling features across the surface of a composite plate. These models are compared, demonstrated and discussed in context of experimental data in Section 5. Finally, conclusions are made in Section 6 and possible future works are presented.

2. Physics of guided waves

Guided waves are used in several engineering applications, such as non-destructive evaluation (NDE) and structural health monitoring (SHM); with prior knowledge of how these waves behave, defects and inhomogeneities can be detected in structures. Such waves undergo an interesting phenomena when they occur in particular structures, such as rods, hollow cylinders and plates; they propagate primarily in the longitudinal direction perpendicular to oscillation and are known as guided waves. When such waves are guided due to propagation along the surface of a medium, they are called Rayleigh waves. However, if a wave travels in a bounded medium, perpendicular to two surfaces, where the wavelength is sufficiently long compared to the distance between these surfaces, often exhibited in plates, it is called a Lamb wave. Overviews of the derivations and characteristics of such waves are well described by Rose [13] and Worden [14].

2.1. Attenuation of guided waves

As Rayleigh waves propagate along a surface of a structure, their amplitude A , decays with propagation distance [7], x , by

$$A(x) \propto \frac{1}{\sqrt{k_R x}} \quad (1)$$

where k_R is the real wavenumber. The attenuation of Lamb waves depends on many factors, although Pollock [15] states the four most important ones to be:

- (i) geometric spreading,
- (ii) material damping,
- (iii) dissipation into adjacent media,
- (iv) wave dispersion,

Attenuation of Lamb waves has been accurately modelled by the inclusion of proportional damping through numerical and experimental studies [16, 17], and the effect of propagation distance on the amplitude of the wave has been described for geometric spreading [17] briefly,

$$A(x) \propto A_0 \sqrt{x_0/x} \quad (2)$$

and material damping [13] as,

$$A(x) \propto A_0 \exp(-\zeta_i x) \quad (3)$$

86 where A_0 & x_0 are the amplitude and distance at an initial location from a
point source, and ζ_i is the attenuation coefficient of the viscoelastic medium.
88 Due to dispersive characteristics of Lamb waves, ζ_i is dependent on the central
frequency of the wave. A key factor of guided waves in plates is the variety
90 of wave modes that propagate within a single *wave-packet*, which are split into
two main types: symmetric, S_n , and antisymmetric, A_n , each of which has an
92 increasing number of modes with increasing *frequency-thickness* of the plates in
question.

94 2.2. Anisotropic media and guided wave feature-spaces

When Lamb waves propagate in anisotropic media, the modelling and solutions
96 become very complex, even more so when attempting to model their interaction
with defects [18, 19, 20]. A crucial characteristic of guided waves in fibrous
98 materials such as carbon-fibre-reinforced-polymer (CFRP) is the phenomena
of Continuous Mode Conversion (CMC) [15], as shown by Mook *et al* [9] and
100 Willberg *et al* [10], where the boundaries of layers or weaves cause conversion of
 S_0 modes into A_0 with frequent enough occurrence that they can be considered
102 continuous along the propagation path. At propagation paths through the fibres,
the energy of all modes is reduced thanks to this phenomena.

104 In order to use guided waves in damage detection strategies, a baseline state
must first be determined, which is often done through the use of features which
106 change in the presence of damage [4]. Determination of a baseline model can
be either model driven, data driven, or a combination of both. As discussed
108 with the example of guided waves in a fibre composite, analytical/numerical
models can be difficult to develop and may not be robust enough to reasonably
110 assume an accurate baseline state. Thus, there has been an increase in the use
of purely data-driven models to determine a baseline state [21]; however, it is
112 then reasonable to assume this baseline state *only* applies to nominally-identical
scenarios/structures from which the data are collected. Since models generated
114 from such methods are general and not specific to the scenario of interest,
enhancing the data-driven models may be valuable in order to offer advantages,
116 such as extrapolability, whilst maintaining accuracy from real data.

As mentioned above, this paper focusses on one guided wave feature, which
118 is indicative of the amplitude of the first asymmetric mode. However, this
does not limit the ideas, or even methods, of the work to just this feature; any
120 energy-based features may be modelled with the same or a similar strategy, and
the process of incorporating prior knowledge for other physics-based features
122 can follow quite readily. For example, the time-of-flight modelling could be
augmented by using prior knowledge of the dispersion characteristics of guided
124 waves in the material.

3. Experiment

126 Guided waves were initiated in a CFRP plate by excitation of a piezo-electric
transducer (PZT), the location of which can be seen in Fig. 1a. The PZT was
128 actuated with a square pulse of temporal width t_{pw} , which would result in a
frequency excitation range up to $1/t_{pw}$ Hz being excited, allowing multiple wave
130 modes to be excited within the plate. A Polytec scanning laser vibrometer was
used to measure the out-of-plane surface velocity of the induced wave-packets

132 on the opposite side to the PZT, where the recording start was synchronised
 134 with the function generator attached to the PZT. Each wave-packet was then
 136 fed through a simple feature extraction process to generate a two-dimensional
 feature-space map of the maximum of the Hilbert envelope, h_m , over the surface
 of the plate. Specific details of the experimental setup are shown in Section 3.

Plate dimensions	300mm x 300mm x 1mm
Layup	[90/0/90] _s , Epoxy matrix
PZT Location	150mm x 150mm
Pulse width	1 μ s
Excitation bandwidth	1MHz
Signal record length	4ms
Pre-trigger	400 μ s
No. scan points	8314
No. of averages	50

Table 1: Details of experimental setup used to acquire feature-space data.

The results of the experiment showing the raw data feature-space map of
 138 h_m over the surface of the plate can be seen in Fig. 1b. One can clearly see the
 140 effect of the fibres on the amplitude of the first asymmetric mode; thanks to
 the phenomena of continuous mode conversion, the amplitude is greater when
 142 propagating along the fibres compared to when propagating across the fibres.
 This experimental procedure, though simple, is pivotal for the advanced tools,
 144 which will be presented later, in order to aid ‘black-box’ data-based ML tools
 for development of a informed, data-driven (IDD) model. It is important to
 146 note the noisy data produced as a result of the low-signal-to-noise ratio that
 occurs from poor reflectivity of the material.

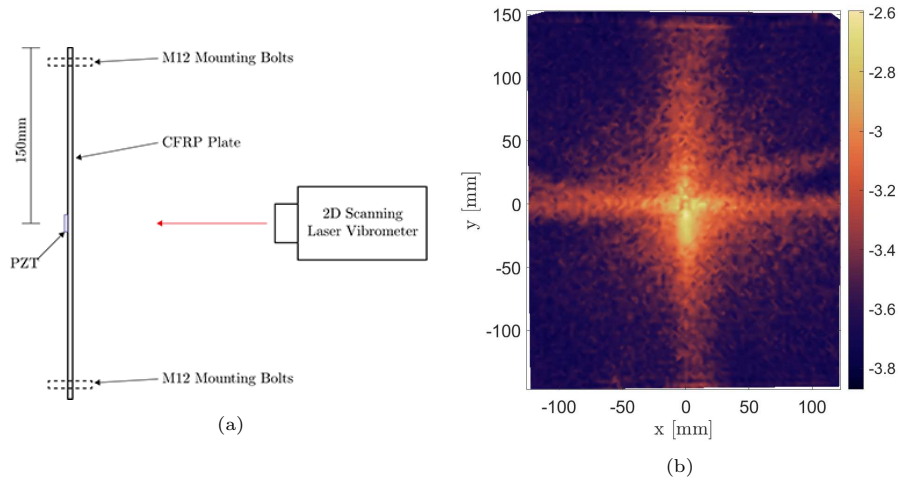


Figure 1: (a) Diagram showing a top down view of the experimental setup and location of PZT on the 300mm x 300mm CFRP plate and (b) results of feature-space map of h_m over the surface of the plate from raw data, represented in \log_{10} scale for viewing purposes.

4. Structured machine learning tools

148 This section will discuss how, on the basis of physical understanding of the
 150 guided waves, data-based models of guided wave features in composites can be
 developed. Initially, a one-dimensional attenuation model is considered. This
 152 choice allows insight into the effect and contribution of the first two attenuation
 mechanisms described by Pollock[15], geometric spreading and material damping.
 Following this investigation, two-dimensional Gaussian process models are considered.
 154 After considering the effectiveness of a ‘black-box’ approach where the model
 is purely based on the data it has seen. The knowledge of the guided waves
 156 is included through two techniques; one of which is through incorporation of
 a mean function in the radial direction from the source. The mean function
 158 models the one-dimensional attenuation, i.e. it models the geometric spreading
 and material damping in the composite. The second approach is to include the
 160 physical knowledge of the guided waves through modifications to the kernel of
 the Gaussian process.

162 4.1. One-dimensional attenuation modelling

The one-dimensional model of the wave attenuation here is based on a
 Bayesian linear regression (BLR). A full description and some derivations for
 this method are provided by Murphy [22], but a short introduction is given in
 Appendix A. Several basis expansions of propagation distance x were tested
 in combination in this work, the full model is shown in Eq. (4). These basis
 expansions correspond to different attenuation mechanisms associated with damping,
 geometric spreading and a combination of the two; these models are shown in
 Table 2. The parameters ($\Phi = \{\phi_1, \phi_2, \phi_3\}$) are switching parameters ($\phi_n =$
 0, 1) which control the combination of functions shown in Eq. (5b).

$$A(x) = \beta_1^{\phi_1} (\exp(-\beta_2 x))^{\phi_2} (x^{-1/2})^{\phi_3} \quad (4)$$

164 The model shown above is not a linear form that can be represented by Eq.
 (A.1), however, by taking the natural log of Eq. (4) a *linear-in-the-parameters*
 model can be developed. This is shown below in Eq. (5b).

$$f(x) = \ln(A(x)) = (\ln(\beta_1)) \cdot \phi_1 - (\beta_2 x) \cdot \phi_2 + (\ln(x^{-1/2})) \cdot \phi_3 \quad (5a)$$

$$f(x) = \ln(A(x)) = w_1 \cdot \phi_1 - (w_2 x) \cdot \phi_2 + (\ln(x^{-1/2})) \cdot \phi_3 \quad (5b)$$

166 The values of β_1 and β_2 can be recovered by exponentiating w_1 and w_2 respectively.

Model Basis	Linear Form	Basis Parameters
$A_1(x) = \beta_1 \exp(-\beta_2 x)$	$f(x) = \ln(\beta_1) - \beta_2 x$	$\Phi_1 = [1, 1, 0]$
$A_2(x) = \beta_1 x^{-1/2}$	$f(x) = \ln(\beta_1) + \ln(x^{-1/2})$	$\Phi_2 = [1, 0, 1]$
$A_3(x) = \beta_1 \exp(-\beta_2 x) x^{-1/2}$	$f(x) = \ln(\beta_1) - \beta_2 x + \ln(x^{-1/2})$	$\Phi_3 = [1, 1, 1]$

Table 2: Bayesian linear regression model basis expansions

4.2. Two-dimensional attenuation modelling

168 So far, it has been shown how physical attenuation phenomena can be
169 modelled as a Bayesian linear regression along one dimension. Now, attention
170 turns to modelling the two-dimensional input feature space, which will be shown
171 in Section 5 on a composite plate. If such a plate were homogeneous, modelling
172 of the features along any one direction would provide an adequate model of the
173 two-dimensional feature space. However, in non-homogeneous materials — such
174 as composites — this is no longer sufficient. Instead, the attenuation changes
175 with direction, and therefore the model of the space must be able to capture
176 changes in behaviour across the two dimensional field. For waves propagating
177 from a point source, it can be helpful to think that there is a radial and angular
178 component to the *function* over the space which describes the feature of interest.

To build such a model it is necessary to have a tool which can model data
180 across a two-dimensional space on the basis of observed data and which can be
181 guided by belief about the physical phenomena. For this, a machine learning
182 approach is adopted; the tool chosen for the job is a *Gaussian process* (GP).
183 The Gaussian process is a flexible Bayesian regression method which works by
184 placing a prior over functions, which is then updated, on the basis of data, to
185 return a posterior distribution over functions [23, 24]. A brief introduction to
186 Gaussian processes can be found in Appendix B.

The kernel used in the GP is a significant modelling choice, and modifications
188 of these provides structure through embedding prior belief of the model. These
189 kernels are computed as any other kernel; linear pair-wise distances between
190 points to form a covariance matrix. More detailed theory can be found in
191 Appendix B and practical implementation can be found in [23]. There are a
192 number of choices available for the kernel function, each of which embeds a
193 different prior belief as to which *family of functions* $f(\mathbf{x})$ is drawn from. For
194 example, if a linear kernel is used, the solution to a Bayesian linear regression
195 is recovered. More commonly, nonlinear kernels will be chosen, as many tasks
196 require regression of nonlinear functions; popular choices include the use of the
197 Squared-Exponential (SE) kernel or the Matérn class of kernels. In most cases
198 the mean function is set to zero in the prior; however, in this work it will be
199 important to consider if the mean functional behaviour can be specified by the
200 physics of the guided waves.

An important characteristic of GPs is that standard stationary kernels operate
202 based on the Euclidean distance between two points, and so map covariances
203 well when using a Cartesian space input. However, the physics and behaviour
204 of guided waves is described here using the polar coordinate system, as they are
205 emitted from a point source. Therefore, this attribute must be considered when
206 utilising GPs for modelling the feature-space of guided waves. Padonou and
207 Roustant [25] outline a method of applying GPs to a polar input space, where
208 separate kernels are applied to the angular and radial dimensions separately,
before combing to generate an overall covariance function.

210 4.3. General nonlinear kernels

It will be important to consider how the GP would model the data if no
212 restrictions were placed on it with respect to the physical behaviour of the guided
213 waves. This type of model will provide a benchmark against which the proposed
214 models can be compared. Two important properties, which certain kernels

216 possess, are stationarity and isotropy. A *stationary* kernel is only dependent
upon the difference between any two points, not the absolute values of those
218 points. An *isotropic* kernel is invariant to translation or rotation of the input
data; practically, this appears as the covariance being only dependent on the
220 absolute difference between two data points [23]. These properties will be
important when discussing what is desired from a kernel to model the features
of guided waves.

One such stationary and isotropic kernel, is the popular squared-exponential
(SE) kernel [23]. This kernel is given by,

$$k_{\text{SE}}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp \left\{ -\frac{\|\mathbf{x} - \mathbf{x}'\|_2^2}{2\ell^2} \right\} \quad (6)$$

An alternative general nonlinear kernel is the Matérn 5/2 kernel (as applied here
to the radial dimension); this well-established kernel is used as it offers smooth
shapes and is defined as,

$$k_{\text{mat}}(\mathbf{x}, \mathbf{x}') = \left(1 + \frac{\sqrt{5}|\mathbf{x} - \mathbf{x}'|}{\ell} + \frac{5|\mathbf{x} - \mathbf{x}'|^2}{3\ell^2} \right) \exp \left(-\frac{\sqrt{5}|\mathbf{x} - \mathbf{x}'|}{\ell} \right) \quad (7)$$

222 The simplest and perhaps most obvious choice for mapping the features
across a composite plate would be to set $\mathbf{x} = \{x, y\}$, the cartesian coordinates
224 of a location on the plate. This method imposes the prior belief that the feature
being modelled across the plate will vary smoothly in a nonlinear manner with
226 respect to the x and y coordinate.

However, since it is known here that the waves are generated from a point
source in the centre of the plate and that these will propagate from that point,
the behaviour would be better modelled in a set of polar coordinates. For the
general case, here the work of Padonou and Roustant [25] is followed, and a
detailed description is provided in Appendix C. An important distinction of
the polar kernel is the definition of distances between points, specifically in the
angular dimension where, for example, a point with angle 359° should have a
high covariance with 1° if the radii are also close. The general polar kernel is
defined as,

$$k_2(\mathbf{x}, \mathbf{x}') = \sigma_f^2 (1 + \sigma_{f,r}^2 k_{\text{mat}}(\rho, \rho')) (1 + \sigma_{f,a}^2 k_{\text{W}}(\theta, \theta')) \quad (8)$$

228 where $\sigma_{f,m}$ and $\sigma_{f,a}$ act as weights representing the influence of changes in each
dimension on a change in the output.

Since the GP is a generative model over functions, it is possible to sample
230 realisations of possible functions from the distribution, allowing the user to
understand the type of functions that the kernel will generate. Using this
232 method, four prior realisations were generated from the polar kernel shown
in Eq. (8) and can be seen in Fig. 2. It can be seen that the functions generated
234 appear to operate separately on each dimension ρ and θ . A key characteristic
to note, is that there is no discontinuity as θ moves through 2π to zero in
236 the angular direction; this is as a result of the Wendland- C^2 function kernel
(Eqs. (C.1) and (C.2)). Further discussions on the characteristics of such a
238 kernel can be found in [25]. For this work, the polar kernel will serve as an
alternate model where there is very little restriction placed on the functions
240 that can be modelled.




Figure 2: Selection of four random polar space functions randomly generated from covariance function represented in Eq. (8), with hyperparameter values of $\tau = 4$, $l = 1$ and $\sigma_{f,m}^2 = \sigma_{f,a}^2 = 1$.

4.4. Mean function modelling of attenuation

242 From this point onwards, the model learning will make use of prior knowledge
of guided wave propagation. This informed model learning will begin by considering
244 how a mean function $m(X)$ can be used to introduce a physical basis to the
model. There is no restriction on this mean function given that it is known.
246 Mathematically, it is trivial to include the mean function (if known) through
simply subtracting the expected mean function from the target data and training
248 on the residuals,

$$\tilde{\mathbf{y}} = \mathbf{y} - m(\mathbf{x}) \quad (9)$$

This can be interpreted as learning the difference or discrepancy between
this chosen mean function and the generating function of the data. In this
scenario, the mean function $m(\mathbf{x}) = \phi(\mathbf{x})\mathbf{w}$ is the model described in Section 4.2,
the one-dimensional Bayesian linear regression model. Since the weights of the
model vary depending on propagation direction with respect to fibre orientation,
it is necessary to simultaneously learn the distribution of \mathbf{w} the weights of the
BLR and the hyperparameters of the GP. Therefore, the linearised form shown
in Eq. (5b) is used and the target data becomes,

$$\tilde{\mathbf{y}} = \ln(\mathbf{y}) - \ln(\rho^{-1/2}) - m(\mathbf{x}) \quad (10)$$

250 and the steps for training and expected values of the mean and variance can
be followed in [23]. This solution can be interpreted as finding the mean one
dimensional behaviour across all propagation directions. This mean behaviour

252 is then compensated by the GP to fit the observed data and learn the latent
 253 function which describes it.

254 4.5. Functional priors through kernel design

255 An alternative to using the mean function to include what is known about
 256 guided waves, is to modify the kernel. It is possible to restrict the *family of*
 257 *functions a priori* to generate only functions which are plausible, given physical
 258 understanding of the guided waves. This restriction is the key advantage of the
 proposed Bayesian approaches in this paper.

260 All of the models which are generated from this point onwards will consider
 the propagation of the guided waves to occur in a polar coordinate system, where
 262 the source is located at radius $\rho = 0$, i.e. the source is at the origin. Since the
 kernel which defines the process will be composed of the ANOVA combination
 264 of the radial component and the angular component, it is possible to consider
 how to modify each of these components individually.

266 4.5.1. Imposing rotational symmetry in the feature space

It can be seen in Fig. 1b, that the energy of the wave exhibits a symmetry on
 the plate. Physically, this makes sense given what is known about the symmetry
 in the orientation of the fibres in the lay-up; it is therefore desirable to exploit
 this in the kernel. Eq. (11) was designed to model this symmetric behaviour.
 The strictly periodic kernel is applied to the angular dimension, where n can be
 altered to include the number of axes of symmetry,

$$k_{\text{sym}}(\theta, \theta') = (\alpha_1 + \alpha_2 \cos(2nd_2)), \quad n \geq 1 \quad (11)$$

268 $d_2 = \arccos(\cos(\theta - \theta'))$ is the geodesic distance, n the number of symmetry
 axes required, α_1 is the offset term, and α_2 the amplitude hyperparameter.

Sample functions from this prior are shown in Fig. 3a. These samples show
 270 the strict periodicity that this kernel imposes. Notably, this form of kernel
 does not enforce the phase of the function. Instead, as data are observed,
 272 the posterior distribution — or function that is learnt — adapts to the phase
 information in the data. This adaptability is a benefit since enforcing phase
 274 within the kernel may lead to issues if the fibre orientation is not known exactly.

The primary issue with the strictly periodic kernel is the assumption of
 276 consistent rate of variation in the function with θ , which may make Eq. (11) too
 restrictive to model accurately the guided wave feature space. It is clear to see
 278 in Fig. 1b, that the energy of the wave decays rapidly in the transition between
 propagation along fibres and across fibres as the angle θ varies. Therefore,
 280 flexibility was introduced by an additive combination of the strictly-periodic
 kernel and the squared-exponential kernel previously shown in Eq. (6).

As well as allowing varying rate of change with θ , a combined kernel also
 reduces some of the restrictions that are imposed with the pure strictly-periodic
 kernel Eq. (11). An additive combination was used as opposed to multiplicative,
 as this does not generate large variations in amplitude between period peaks and
 allows for the kernel to capture the symmetry, while still allowing some variation
 to take place. The additive combination can be considered an ‘OR’ operation
 [26], the resulting kernel applied to the angular dimension is given by,

$$k_{\text{ang}}(\theta, \theta') = \sigma_{f,\text{sqe}}^2 \exp\left(-\frac{d_2^2}{l_1^2}\right) + \sigma_{f,\text{sym}}^2 (\alpha_1 + \alpha_2 \cos(nd_2)) \quad (12)$$

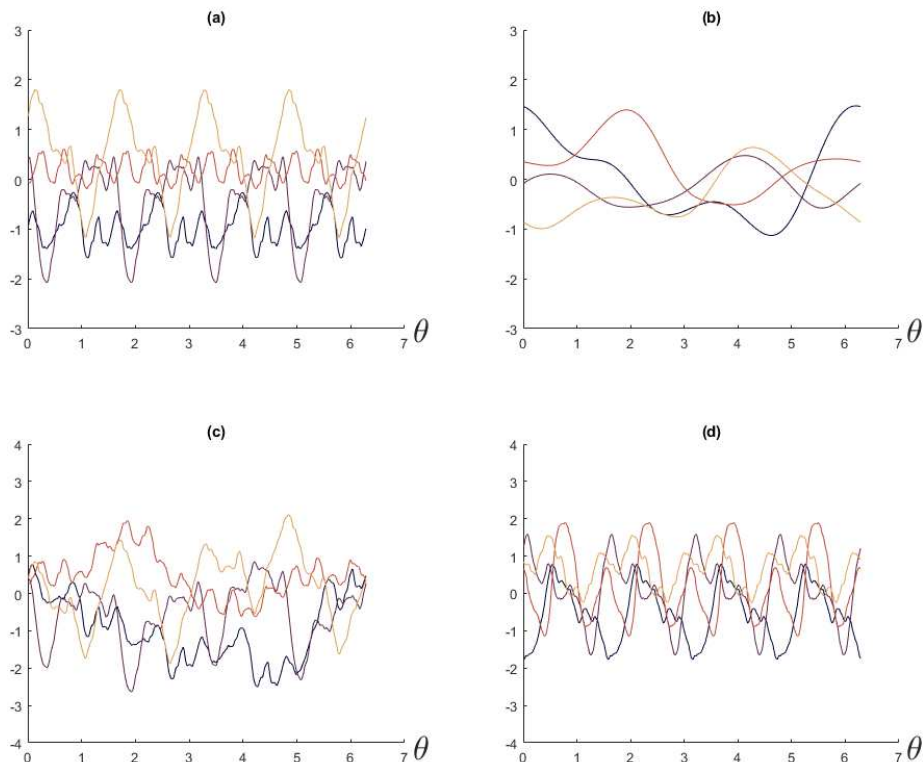


Figure 3: Selection of random priors for the angular kernel designs, over a full circle range, for (a) strictly-periodic kernel (Eq. (11)), (b) squared-exponential kernel (Eq. (6)), (c) multiplied combination and (d) additive combination of the kernels (Eq. (12)). Each line represents a different random function drawn from these priors.

282 An important point of the resulting kernel is that it is stationary, as it is only
 284 proportional to the distance between points, rather than their values. This
 means that the kernel is unaffected by translation or rotation of the coordinates,
 a key advantage when modelling in the angular dimension.

286 4.5.2. Kernel forms for radial attenuation

As discussed in Section 2, two known and documented attenuation mechanisms
 can be modelled; these are shown in Eqs. (2) and (3). Attenuation effects due
 to viscoelastic damping of a material can be embedded into the priors through
 the use of an *exponential decay* (ED) kernel, where propagation distance x
 is replaced with ρ ,

$$k_{\text{ed}}(\rho, \rho') = \exp(-\rho l) \cdot \exp(-\rho'^{\top} l) \quad (13)$$

Attenuation effects due to geometric spreading can be modelled through the use
 of a polynomial kernel (Eq. 14), where $p = -1/2$, in order to represent Eq. 2.

$$k_{\text{pol}}(\rho, \rho') = (\rho \cdot \rho'^{\top})^p \quad (14)$$

288 Prior draws from both the ED and polynomial kernels can be seen in Figs.
 4(a) and 4(b) respectively. Both kernel functions embed decay with respect to
 propagation distance ρ , but each model shows a different mechanism for this




Figure 4: Selection of random priors for the radial kernel designs, for (a) exponential decay kernel (Eq. (13)), (b) square-root decay (Eq. (14)), (c) multiplicative combination (Eq. (15)) and (d) additive combination of the kernels. Each line represents a different random function drawn from these priors.

290 decay. The kernel chosen to represent decay due to the geometric spreading will always tend to infinity as $\rho \rightarrow 0$ and this limitation should be considered.

To illustrate the use of these kernels, multiplicative and additive combination of these two attenuation mechanism kernels can be seen in Figs. 4(c) and 4(d) respectively. A multiplicative combination of the attenuation mechanisms aligns more closely with physical understanding and the discussion presented in [17], since the kernels operate to reduce the energy in the wave simultaneously and do not subtract energy, but rather reduce it. As such, a kernel to model attenuation along the radial direction is proposed as the multiplicative combination of the exponential decay (Eq. (13)) and polynomial (Eq. (14)) kernels; this is

$$k_{\text{rad}}(\rho, \rho') = \sigma_{f,r}^2 (\rho \cdot \rho'^{\top})^p \cdot (\exp(-\rho l_2) \cdot \exp(-\rho'^{\top} l_2)) \quad (15)$$

292 4.5.3. Combined two-dimensional kernel

294 It has been shown how understanding of the physical processes involved in attenuation of guided waves can be used to impose prior belief in the GP machine learning model along each of the radial and angular dimensions. It remains to explain how these may be combined to form a meaningful prior over the two-dimensional feature space.

296 Following closely the approach of [25] for the general nonlinear kernel in polar coordinates, the two kernels described in Eqs. (12) and (15) will be combined

using an *ANOVA* approach,

$$k_3(\mathbf{x}, \mathbf{x}') = (1 + \sigma_{f,a}^2 k_{\text{ang}}(\theta, \theta')) (1 + \sigma_{f,r}^2 k_{\text{rad}}(\rho, \rho')) \quad (16)$$

where $\mathbf{x} = \{\{\theta, \rho\}_1, \dots, \{\theta, \rho\}_m\}$.

Again, it is possible to draw samples of the functions, now in the two-dimensional space, to visualise the restrictions which have been placed on the functions that can be modelled. Four prior draws from a GP with zero mean and the covariance defined in Eq. (16) can be seen in Fig. 5. It should be noted at this point that the input units and output values are non-dimensional, and the figures showcase key characteristics imposed by the kernels by displaying functions that are samples from an arbitrary prior. It is reassuring that these prior draws match, at least visually, the behaviour that would be expected in the data being modelled. This type of *prior predictive checking* can be invaluable for confirming that the assumptions built into the model are reasonable. A key feature that can be seen is the symmetry that this introduced in the angular dimension without requiring a fixed phase to be specified. It can also be seen that slight variations from this symmetry are possible due to the inclusion of the squared-exponential kernel in Eq. (12). The decay in the radial direction as a result of the kernel shown in Eq. (15) is also clearly seen. As a result of these characteristics it has been shown how a GP kernel can be designed in such a way that it is applicable to modelling the attenuation of guided waves in a two-dimensional space.

A second kernel was also tested with the same radial component as described in Eq. (15), but with an alternative angular kernel. For this method, $k_{\text{ang}}(\theta, \theta')$ becomes a modified version of Eq. (C.2), where the geodesic distance is instead defined as $d_2(\theta, \theta') = \arccos(\cos(2n(\theta - \theta')))$, where n is again the number of symmetry lines required. This alteration was done to still enforce symmetry but allow a more flexible modelling of the functions being considered in the angular dimension. This kernel has the form,

$$k_4(\mathbf{x}, \mathbf{x}') = (1 + \sigma_{f,a}^2 k_{\text{W}}(\theta, \theta')) (1 + \sigma_{f,r}^2 k_{\text{rad}}(\rho, \rho')) \quad (17)$$

So far, only samples from each GP have been shown. However, it is possible to recover the distribution over the model in closed form. Since the mean function chosen in most models is simply zero across the complete space, the prior mean is not very informative. The prior variance, however, is of interest to consider. In Fig. 6 the prior variances of Eqs. (16) and (17) are shown. It is important to remember that this is the variance in predictions made by the model before the information from any data has been included. The variance for both of these kernels is seen to decay as the distance from the source increases. The model will tend towards infinity at $\rho = 0$ for two reasons; the kernels are non stationary [23], and due to the exponential decay included through Eq. (13), the function values themselves will tend towards infinity as seen in Fig. 5. This limitation of the models should be considered and care should be taken if predicting close to $\rho = 0$. However, in the experimental data used in this study the waves are generated by means of a piezoelectric actuator. This means that the source of the guided waves is not a point source and the models should not be used within the region covered by the piezoelectric actuator.

4.6. Overview of modelling approaches

Up to this point, it has been discussed how one might construct a GP kernel which can represent the behaviour of guided wave attenuation. It is




Figure 5: Selection of four random polar space functions, selected from 1000 random functions generated from covariance function represented in Eq. (16), with hyperparameter values of $\Theta_k = \{1, 10, 10, 1, 1, 1, 1, 0.001\}$. Here the functions are non-dimensional as the functions are samples from an arbitrary prior.

342 worth reviewing the models which will be compared when results are shown on
 experimental data. Table 3 shows a summary of all six models which will be
 344 compared in this study.

As a baseline, model A is the archetypal Gaussian process model with a
 346 zero-mean function and the use of the squared-exponential kernel operating on
 two inputs, the x and y coordinates on the plate. This provides a benchmark
 348 where no knowledge of the guided-waves is included.

The second model (B) is a demonstration of the use of the polar coordinate
 350 GP of [25]. This model also contains no specific reference to the physical
 mechanisms in guided-waves, but does make use of the knowledge that the
 352 guided-waves propagate radially from a source. The use of a polar coordinate
 system in this case is a sensible choice given the structure of the data being
 354 used. This method serves as another benchmark demonstrating an approach
 which requires very little understanding of the physical mechanisms involved in
 356 guided wave propagation.

Model C is the first model where a specific physical process is included.
 358 In this case, the model of wave attenuation $A_3(x)$ (Table 2) is used along the
 radial direction as the mean function. The kernel used is the same flexible polar
 360 kernel as in model B — that proposed in [25]. Importantly, this kernel remains
 flexible to influence the model in both the radial and the angular dimensions,




Figure 6: Prior variances over polar coordinate space from kernel E (left) and kernel F (right), represented by Eqs. (16) and (17) respectively. Variances are plotted in \log_{10} -scale for viewing purposes.

362 potentially correcting for any unmodeled phenomena along the radial dimension
 363 in the mean function.

364 Model D restricts the flexibility of Model C by removing the dependence
 365 of the kernel on the radial dimension. The GP used here relies on the mean
 366 function to capture all of the radial behaviour through $A_3(x)$ and the covariance
 367 to capture all the variation in the angular dimension. This implies that the
 368 data can be generated by some function, offset from the mean, which is only
 369 dependent upon the angle being considered. This model should be considered
 370 with care since it is highly restrictive.

371 Model E removes the use of the mean function; instead the knowledge of
 372 the guided waves is embedded directly in the kernel as a *functional prior*. This
 373 model enforces periodicity in the angular dimension and embeds the physical
 374 attenuation models in the radial direction by means of the kernel described in
 375 Eq. (16). Model F is very similar to the model E but with the kernel defined as in
 376 Eq. (17), with the modification to the angular component described previously.

Model	Mean $m(\mathbf{x})$	Covariance $k(\mathbf{x}, \mathbf{x}')$	Input space
A	0	$k_{\text{SQE}}(\mathbf{x}, \mathbf{x}')$	Cartesian
B	0	$k_2(\mathbf{x}, \mathbf{x}')$	Polar
C	$A_3(\rho)$	$k_2(\mathbf{x}, \mathbf{x}')$	Polar
D	$A_3(\rho)$	$k_{\text{W}}(\theta, \theta')$	Polar
E	0	$k_3(\mathbf{x}, \mathbf{x}')$	Polar
F	0	$k_4(\mathbf{x}, \mathbf{x}')$	Polar

Table 3: Table of GP strategies tested for feature-space mapping, showing the properties and characteristics of each model.

378 Appendix D provides a reference in which each of these kernel forms can be
 379 compared and in which the hyperparameters are listed. The reader may find
 380 this a useful companion if planning to reproduce the methodology from this
 work.

4.7. Hyperparameter learning

382 Thus far, the kernels of the GP have been presented as priors over the
 functions which that GP will generate. These modelling strategies have allowed

384 embedding of physical processes governing the behaviour of guided waves in
 385 a flexible and rigorous manner. However, each of these kernels has a small
 386 number of associated *hyperparameters* which govern the characteristics of the
 387 family of functions they represent. It is necessary, therefore, to review how a
 388 user may practically ascertain the values of these hyperparameters. As with
 389 many problems in machine learning, and indeed engineering, this boils down to
 390 an optimisation problem. The specific form of this problem will now be shown.

The hyperparameters vary depending upon the form of the kernel, but for
 392 generality Θ_k is considered to be the vector of hyperparameters for whichever
 393 kernel is being used. For example, in the case of the kernel proposed by [25],
 394 this vector is defined as $\Theta_k = \{l, \sigma_{f,r}^2, \sigma_{f,a}^2, \tau\}$; whereas, for the kernel in Eq.
 (16) which is used in model E, this vector is $\Theta_k = \{l_1, \alpha_1, \alpha_2, \sigma_{f,sqe}^2, \sigma_{f,sym}^2,$
 396 $l_2, \sigma_{f,r}^2, \sigma_{f,a}^2, \sigma_n^2\}$. The hyperparameters each control distinct and important
 characteristics to the kernel. As an example for the kernel used in model E (Eq.
 398 (16)), these characteristics are interpreted as follows:

- (i) l_1 and l_2 are the characteristic length scales of the SE kernel and ED kernel
 400 respectively.
- (ii) α_1 and α_2 represent the offset and scaling term for the sine-wave-based
 402 strictly-periodic kernel.
- (iii) σ_f^2 terms are scaling factors for individual kernels which control their
 404 relative importance when combined.
- (iv) σ_n^2 is the noise variance parameter related to the expected measurement
 406 noise on the signal.

In this work, the parameters of the mean function are also included in the
 408 optimisation routine to estimate them synchronously with the hyperparameters
 of the kernel. This optimisation problem is most commonly cast as the maximisation
 410 of the *marginal likelihood* of the Gaussian process which is available in closed
 form [23]. This optimisation strategy has several advantages such as leveraging
 412 the Bayesian Occam’s Razor [27]. For computational reasons this method is
 realised practically as a minimisation of the *negative log marginal likelihood*
 414 (NLML), more details of which can be found in [23] and Appendix B.

As such, the optimisation task at hand is formally to estimate,

$$\hat{\Theta} = \arg \min(-\log p(\mathbf{y}|\Theta)) \quad (18)$$

For all results in this paper, Eq. (18) was optimised using the *quantum*
 416 *particle swarm* technique [28]. Discussion of this choice can be found in [29].

5. Results

418 Having developed a number of approaches for modelling of guided-wave
 features, it is necessary to demonstrate the differences between each of these
 420 on an experimental dataset. The dataset is chosen is the composite plate
 described in Section 3. Initially, it will be shown how the one-dimensional
 422 models described in Section 4.2 can be applied along a fixed radial direction.
 Following this, the two-dimensional models laid out in Table 3 will also be
 424 compared.

5.1. One-dimensional attenuation models

426 The models shown in Section 4.2 are compared for two different cases; the
 428 first is where data are collected along the fibres in the weave and the second
 when across the fibres. The $NMSE$ is calculated for both the training set and
 the test data sets. Figure 7 and Table 4 show the results of this process. From
 430 the values of the $NMSE_{tr}$ and $NMSE_t$ for both propagation orientations, it
 appears that the multiplicative combination of geometric spreading and material
 432 damping is the most promising model for the attenuation of the energy of the
 first antisymmetric mode for all propagation directions in a CFRP plate. It is
 434 also interesting to note the increased quality of fit of model Φ_2 in comparison
 to Φ_1 indicates that geometric spreading has a larger effect on the attenuation
 436 of the energy than material damping.

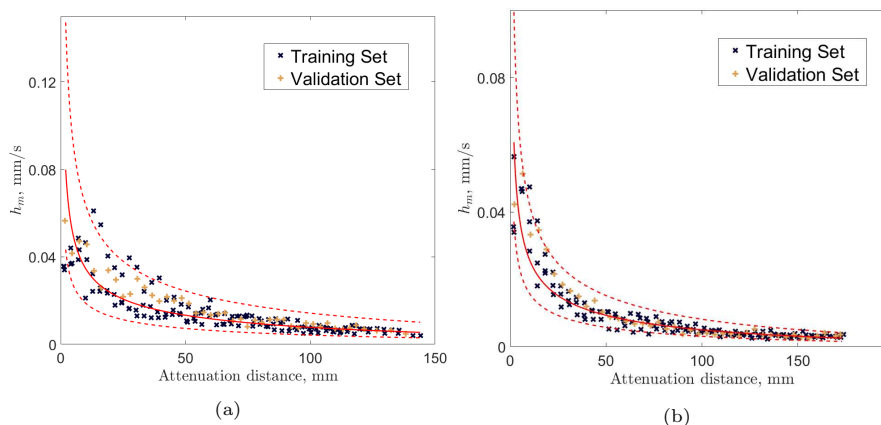


Figure 7: Models of attenuation of the energy of the first antisymmetric mode using model Φ_3 for propagation directions (a) along the fibres and (b) through the fibres. The estimated mean function is plotted with a solid red line and the 95% confidence interval is shown with dashed red lines.

438 When modelling the attenuation of the guided-waves propagating through
 the fibres, the 95% confidence intervals, seen in Fig. 7, are much smaller, which
 is due to the less noisy distribution of values for h_m around the estimated mean.
 440 Physically, this could be envisaged as the increased number of boundaries for the
 wave to propagate through causing more frequent mode conversion and so there
 442 is a more consistent energy dissipation as it tends towards *continuous mode*
conversion [9, 10]. In contrast, when propagating along the fibres, the wave
 444 mode is relatively uninterrupted in comparison, and so its initial energy has a
 strong effect on the resulting shape. Furthermore, attenuation due to geometric
 446 spreading is more likely to be described with Eq. (2) during CMC. However,
 this is currently an assumption, experimental validation of this is planned as
 448 future work.

As can be seen in Table 4, the confidence in the mean weights (which
 450 increases as σ_w decreases), is similarly large for model Φ_1 but is much more
 confident for model Φ_2 when being applied across the fibres. This observation
 452 may also form the basis for further investigation.

	β_1	β_2	σ^2	$\sigma_{\mathbf{w}}$	$NMSE_{tr}$	$NMSE_t$
Along fibres						
Φ_1	0.001523	0.01215	9.403	0.1532	0.13203	0.10548
Φ_2	0.004844		8.199	0.0640	0.21864	0.11664
Φ_3	0.005934	0.00288	6.558	0.1068	0.17489	0.14795
Through fibres						
Φ_1	0.000729	0.01277	7.464	0.1321	0.10191	0.09018
Φ_2	0.002332		4.224	0.0371	0.10084	0.09088
Φ_3	0.002685	0.00229	3.548	0.0628	0.08470	0.07404

Table 4: Table of results from 1D attenuation modelling using BLR; \mathbf{w}_1 & \mathbf{w}_2 are model weight parameters, σ^2 is the estimated variance of the function, $\sigma_{\mathbf{w}}$ is the estimated variance of the weights, $NMSE_{tr}$ is the *negative mean squared error* between the model and the dataset used for training, and $NMSE_t$ is between the model and the independent validation set.

5.2. Two-dimensional Gaussian process models

454 In this section, the various approaches to modelling the two dimensional
455 feature space summarised in Table 3 will be compared. Models are compared
456 visually and based upon a number of metrics as described in Appendix E. The
457 quantitative assessment of the models is discussed in Section 5.3. These metrics
458 are the log marginal likelihood LML , which is a measure of how well the model
459 fits the training data. Next, the predictive log likelihood of an independent
460 test set, considered in the case where every prediction is assumed independent
461 PLL_i , and when the predictions are assumed correlated PLL_c . Finally, the
462 normalised mean squared error $NMSE$ of the mean fit to the independent test
463 set is also computed. This final metric should be treated with care, since it
464 does not represent the quality of the uncertainty quantification in the model fit.
465 The most rigorous test of these models can be considered to be the correlated
466 predictive log likelihood which captures the full correlation of the predictive
467 model including the mean, variance and covariance predictions. For all graphical
468 representations shown, the data is presented in the \log_{10} scale, but the models
469 were all trained directly on the values of h_m . Therefore, the units for the figures
470 are in $\log_{10}(mm)$.

5.2.1. Uninformed Gaussian process models

472 Fig. 8 shows the mean predictions of the two uninformed GP models (A and
473 B in Table 3). It can be clearly seen, that even without specific prior knowledge,
474 the use of polar coordinate system (model B) offers a significant improvement
475 over the Cartesian approach (model A). Model A appears to be an ‘out-of-focus’
476 copy of the original data, whereas model B has generated a smoother function
477 which is more likely to represent the physical mechanisms by which the wave
478 operates. Even this simple consideration of the structure of the data being
479 modelled leads to far more consistent results from the model. The quality of
480 the fits for each of these models is compared qualitatively in Table 5 along with
481 the other GP models.

482 Since the wave attenuation data naturally follows (approximately) a polar
483 behaviour, one could envisage this problem as the GP trying to learn the
484 mapping of Cartesian to polar spaces as well as the mapping from the polar
485 to the feature-space. This two-stage mapping is being attempted through a

486 single kernel and significantly complicates the modelling problem, thus, it is
 488 likely to underperform a model specified in the correct space. Further to this,
 it is concerning that the Cartesian model may have attempted to model some
 of the structure in the measurement noise. This is again a topic of further
 490 investigation for the future.

However, as can be seen in Fig. 8b, by learning the model through a uninformed
 492 kernel operating on the polar coordinates, a more believable model of the feature
 space is learnt. These results show that even when implementing machine
 494 learning methods with no direct embedding of the physical process, the space
 in which the function operates must still be taken into account.




Figure 8: Results of uninformed ‘black-box’ GPs; (a) GP strategy A and (b) GP strategy B. The units are in $\log_{10}(mm)$.

496 5.2.2. Guided wave mean functions

Models C and D in Table 3 show the two approaches where mean function
 498 behaviour is included in the model to capture the expected behaviour of the
 guided waves. Both of these models make use of the third one-dimensional
 500 attenuation model from Section 4.2, which includes both geometric spreading
 and viscoelastic damping. Model C couples this mean behaviour in the radial
 502 direction with the flexible polar kernel used in model B, this allows the GP to
 learn functional behaviour in both the radial and angular dimensions. Model D
 504 is more restrictive and it is assumed that the mean function captures all of the
 radial behaviour and the GP only models functional behaviour in the angular
 506 dimension.

The resulting mean predictions on an independent test set for each of these
 508 models are shown in Fig. 9. Considering the prediction shown in Fig. 9a, a
 ‘banding’ effect is seen as a circular structure centred on the origin. This
 510 ‘banding’ is most prominent in areas of propagation through the fibres and
 less prominent in the central region, around the wave source. This banding can
 512 be explained by considering how the mean function is included in the model.
 The mean function $m(\mathbf{x})$ is likely to fit the mean of the ρ dimension well as
 514 the two dominant attenuation mechanisms captured. The inclusion of the mean

function in a Gaussian process can be imagined as subtracting this function from
 516 the relevant dataset. When the mean function captures most of the behaviour,
 only unstructured data should be left to be modelled by the GP covariance, i.e.
 518 the noise in the system. In the results shown here, the mean function fits the
 data well and the unstructured data along the radial ρ dimension is still modelled
 520 in Eq. (7). The GP can still attempt to find a structure in unstructured data.
 An interesting note from this result is that if there were functional information
 522 in the data still to be inferred, this would be picked up by the covariance kernel.
 In this case it is believed that the banding artefact may be due to the GP
 524 modelling structure in the noise on this realisation of the measurement, it is
 expected that if further training data were included, this effect would diminish.




Figure 9: Results of informed model generated using generic kernels and inclusion of a mean function in the GP, with (a) model C, the full polar kernel, and (b) model D, the GP only modelling angular behaviour. The units are in $\log_{10}(mm)$.

526 To avoid this issue, and assuming that the mean function models well the
 radial attenuation behaviour, model D does not include the radial component
 528 in its covariance kernel, so it has covariance given by Eq. (C.1). The results of
 training the model with exclusion of the radial kernel can be seen in Fig. 9b, in
 530 which it can be seen that the banding artefacts are no longer evident. However,
 the model also appears to lose accuracy as the value of h_m attenuates much more
 532 quickly away from the source than is seen in Fig. 1b. The loss of the banding
 artefacts demonstrates that it is the inclusion of this radial dimension in the
 534 kernel which leads to this phenomenon. These models will also be compared
 quantitatively once all models have been shown qualitatively.

536 5.2.3. *Kernels capturing guided wave behaviour*

The final two models (E and F) attempt to embed understanding of the
 538 guided-waves by directly modifying the prior belief in the model through kernel
 design.

540 The estimated mean predictions on an independent test set for each of these
 models are shown in Figs. 10a and 10b. It is clear in these results, how even the
 542 small changes between the two kernels can significantly impact the function
 space that is learnt. Comparatively, model E (Fig. 10a) leads to a much

544 ‘smoother’ result in comparison to model F (Fig. 10b). Fewer high-frequency
 components are seen in the angular dimension leading to this appearance. This
 546 difference is due to the differences in the prior belief imposed in the angular
 kernels for each of these models.




Figure 10: Results of informed model generation using non-generic kernels, here indicating the expected mean function over the entire space when modelling with (a) GP strategy E and (b) GP strategy F. The units are in $\log_{10}(mm)$.

548 An interesting result seen here is the spreading of energy away from the fibre
 orientation. This spreading could be physically explained by acknowledging the
 550 secondary-guiding characteristics of the fibres themselves. As the waves travel
 along the fibres independently, they will lose some energy into adjacent media
 552 (i.e. the epoxy matrix). This will manifest itself as energy spreading in a
 different direction to the fibre orientation.

554 Visually, kernel E appears to allow the capture of the spreading of the wave
 energy better. In Fig. 10b it appears that the decay of energy is not captured
 556 well and that energy is only propagating along the fibres, not across. The
 periodicity enforced in model F is obvious to see in the significantly lower
 558 value ‘band’ at approximately $\theta = \{20^\circ, 110^\circ, 200^\circ, 290^\circ\}$. The kernel used
 in model E offers greater flexibility in symmetry as a result of the additive
 560 combination of the SE kernel along the angular dimension. For both kernels, the
 predictive mean has less variation in the function modelling wave propagation
 562 along the fibres, and is a less smooth function when propagating through the
 fibres. This difference could be improved by increasing the signal-to-noise ratio
 564 of the experiment. This alteration will help model areas of high attenuation
 (i.e. low energy); since the energy of the wave decreases significantly away from
 566 the fibre orientation, the value of h_m may not exceed the noise floor. Currently,
 in regions with large ρ the data becomes unstructured and it is difficult to
 568 infer the function with as much confidence. It may also be worth exploring a
 heteroscedastic noise model in future work [30].

570 For both models E and F, the estimated variance over the surface generates
 similar results; this can be seen in Figs. 11a and 11b. Both models see a sharp
 572 increase in variance towards the centre, this can be explained by examining
 the one-dimensional attenuation kernel proposed in Eq. (15). The polynomial




Figure 11: Results of informed model generation using non-generic kernels, here indicating the expected variance over the entire space when modelling with (a) GP strategy E and (b) GP strategy F.

574 kernel Eq. (15) included in these models will result in functions that tend
to infinity with $\rho \rightarrow 0$. This kernel is used in multiplicative combination
576 with the exponential decay kernel Eq. (13), resulting functions with the same
characteristic. The second kernel design results in a slight increase in variance
578 at approximately $\theta = \{50^\circ, 140^\circ, 230^\circ, 320^\circ\}$. From physical interpretation of
how fibres affect the energy of the waves, the variance should not increase at this
580 orientation if not also at $\theta = \{40^\circ, 130^\circ, \text{etc.}\}$. For both kernel designs, there
also appears a greater variance in the energy of the wave when propagating
582 along the fibres; this is likely a result of the short range of θ in which the wave
has directly propagated away from the source along a single fibre.

584 5.3. Quantification of model performance

Until now, the modelling approaches for guided wave features have been
586 compared in a qualitative manner. It has been discussed how certain models
give rise to desirable characteristics in the latent functions being learnt, which
588 may obscure the physical behaviour.

Model	$m(\mathbf{x})$	$k(\mathbf{x}, \mathbf{x}')$	LML	PLL_i	PLL_c	NMSE
A	0	$k_{\text{SQE}}(\{x, y\}, \{x, y\}')$	24,301	3,745.1	3,770.9	4.424
B	0	$k_2(\{\rho, \theta\}, \{\rho, \theta\}')$	50,012	9,964.4	9,975.40	4.0567
C	$A_3(\rho)$	$k_2(\{\rho, \theta\}, \{\rho, \theta\}')$	19,872	13,047.7	4,221.4	1,037.7
D	$A_3(\rho)$	$k_{\text{W}}(\theta, \theta')$	14,950	11,176.6	3,326.9	29.9
E	0	$k_3(\{\rho, \theta\}, \{\rho, \theta\}')$	96,952	21,013.6	22,727.4	9.9422
F	0	$k_4(\{\rho, \theta\}, \{\rho, \theta\}')$	75,682	16,274.9	16,291.3	9.9133

Table 5: Table of qualitative assessment values for 2D GP modelling strategies tested, indicating the *log-marginal likelihood* (LML), *independent predictive log likelihood* (PLL_i), *co-dependent predictive log likelihood* (PLL_c), and *normalised mean square error* (NMSE). Best values for each metric are highlighted in bold.

Table 5 shows the results of the GP models tested against the metrics

590 previously described. These metrics allow insight into the accuracy and validity
of resulting models, and also provides an opportunity to discuss what might be
592 meant by a “good model”. It is important to consider that any given model is
only as good as what it will be ultimately used for. In the context of engineering,
594 specifically in NDE and SHM, these models will be used to make operational
decisions about the system. As such, it should be considered whether the end
596 user is most interested in the point-wise prediction accuracy in which case the
NMSE is the most appropriate metric. If instead, the models will be carried
598 forward into a risk-driven assessment, then capturing the full uncertainty in the
model is important and the correlated predictive log likelihood will be the most
600 appropriate metric. In this work, a number of possible assessment criteria are
presented; the onus is on the end user to choose the model which best captures
602 the characteristics of the data/function which are most important to them. It
is the opinion of the authors that the most robust measure of how well the
604 functional behaviour of the feature space is captured is the correlated predictive
log likelihood PLL_c .

606 Turning attention to specific results from these experiments, the results for
each of the six models under each metric are shown in Table 5. It can be seen
608 that models A and B result in the lowest NMSE scores, in other words that
their point-wise predictions are closest to the observed test data. This result
610 may be expected since they are the most flexible models. Interestingly, model
B, the informed polar coordinate model, also recovers the highest independent
612 predictive log likelihood. Interpreting this result; if one wants to predict only the
behaviour at single points on the plate and is not concerned with the correlation
614 between these predictions, then this is the optimal model (from those tested).
It is also seen in the log marginal likelihood LML that model B captures the
616 behaviour in the training data much better than model A.

The inclusion of the guided-wave attenuation models as mean functions in
618 the GP (models C and D) appears not to produce desirable effects. These
models perform worst in their representation of the training data LML and their
620 point-wise prediction capability $NMSE$. For this dataset, this is compelling
evidence that the inclusion of the guided wave attenuation mechanisms through
622 a mean function does not lead to a useful model. As discussed, the models may
be finding too much structure in the noise of the data used for training, especially
624 model C. This postulation is evident in the very high $NMSE$ score, which would
indicate that the model performs considerably worse than taking the mean of
626 the prediction data. However, it is also important to note the exceptionally
large NMSE for model C (representing an error of $\sim 1038\%$), which is a result
628 of the estimated mean function resulting in a ‘singularity’ towards the centre of
the plate; $m(\mathbf{x}) \rightarrow \infty$ as $\rho \rightarrow 0$.

630 Finally, considering models E and F, where the knowledge of the guided
waves is used to modify the prior belief in the model, via the covariance kernel;
632 both of these models perform very well when considering their ability to model
the training data assessed through the LML . When examining their performance
634 on the independent test set, it is seen that the $NMSE$ score is around 9%
compared to the 4% of models A and B; this along with their poor independent
636 predictive log likelihoods PLL_i are indicators that the point-wise predictions
from these models are not as good as models A and B. However, in terms of
638 capturing the complete function space they far exceed all the other modelling
strategies. These approaches appear to best capture the underlying functional

640 behaviour of the guided-wave attenuation. For this reason, these models can
642 be considered to be the most suitable for de-noising or spatially up-sampling the
644 data. These two approaches may also be considered the most robust methodologies
646 for modelling the guided-wave behaviour not currently described by governing
equations, e.g. outlying behaviour. This advantage has been achieved through
incorporation of the physical mechanisms driving the attenuation in the prior
specification of the model, via modification of the covariance kernel.

The results from this dataset and model would support the use of model E for
648 modelling the feature-space of h_m for a guided wave in a CFRP plate. However,
when using this strategy users should consider all models and the system being
650 modelled, as well as the level of physical knowledge that is currently available
for said system.

652 6. Conclusion and further work

As computation capabilities and machine learning techniques are becoming
654 increasingly more accessible, the adoption of such methods to solve engineering
problems is also becoming more prominent. This increasing use of such methods
656 can present many underlying issues with the results, such as unreasonable
assumptions, lack of extrapolation capability and computational costs. However,
658 by implementing physical knowledge to guide learning, more robust models may
be generated which can reduce many of these issues.

A barrier to the progression of using guided waves in an NDE or SHM
660 strategy is the difficulty of modelling the behaviour of these waves in complex
materials. The work presented here shows promising steps towards generating
662 a physics-incorporated, data-driven model for the feature-space of guided-waves
in such materials. Several characteristics of such a strategy, which must be
664 carefully considered to maintain robustness, have been discussed. This strategy
provides a key framework for the development of guided-wave models for complex
666 materials — such as the ones used in this paper — by allowing modelling of
features which define the waves propagating throughout the material. The
668 important distinction of this combined method, in comparison to physics or
data-driven-only methods, is that this allows physics to guide the model, whilst
670 allowing unknown or undescribed physical mechanisms to be incorporated through
the data-driven aspect.

When initially looking at the kernels chosen to represent different learning
674 strategies and levels of constraint, it is not clear which strategy will result in the
optimal model output. By various qualitative assessment values it is possible
676 to see how each model fits in comparison to the validation data, in different
aspects. By leaving the model uninformed, it is possible to get a closer fit to the
678 training and validation data in terms of difference between the predicted and
measured values. However, by guiding the learning process using physics-based
680 implementation of the problem, it is possible to get a higher likelihood model.

The work here has been shown for the case of modelling energy-based features
682 of a guided wave in a CFRP plate structure, and though these specific kernels
cannot be directly applied to some other complex scenarios — such as complex
684 geometry, quasi-periodic materials, or other features such as *time-of-flight* —
the kernels can be modified in the framework presented here to model such
686 systems. The structure is applied through the kernels to embed prior belief of
the shape of the features over the topology being modelled.

688 Further work will be done to embed structural knowledge into the modelling
 process — such as plate thickness, joints and layup information — in order to
 690 improve extrapolation of data over guides of varying shape.

7. Acknowledgements

692 The authors gratefully acknowledge the support of the UK Engineering and
 Physical Sciences Research Council (EPSRC) [grant numbers EP/R004900/1,
 694 EP/R003645/1, EP/S001565/1, EP/J013714/1 and EP/N010884/1].

Appendices

696 A. Bayesian Linear Regression

Traditional linear regression formulates a model using point estimates of a
 set of parameters which “best” fit an available dataset, based on minimising an
 L^2 -norm between the model predictions and the data. Instead, BLR aims to
 establish a probability distribution of possible model parameters. The model
 has the form,

$$y = \mathbf{w}^\top \phi(\mathbf{x}) + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma^2) \quad (\text{A.1})$$

where ϕ is some basis for expansion of a p -dimensional data point \mathbf{x} ; $\mathbf{w} = \{w_1, w_2, \dots, w_p\}$ are the associated weights of the basis expansion, and ε is an additive Gaussian white noise distributed as $\mathcal{N}(0, \sigma^2)$. The weights \mathbf{w} and the variance σ^2 are the unknown. The Bayesian linear regression model approach was chosen since it returns a quantified uncertainty. The task is then to compute the posterior distribution of the parameters $p(\mathbf{w}, \sigma^2 | D)$. This posterior distribution has the following form,

$$p(\mathbf{w}, \sigma^2 | D) = \text{NIG}(\mathbf{w}, \sigma^2 | \mathbf{w}_N, \mathbf{V}_N, a_N, b_N) \quad (\text{A.2})$$

with,

$$\mathbf{w}_N = \mathbf{V}_N(\mathbf{V}_0^{-1}\mathbf{w}_0 + \mathbf{X}^\top \mathbf{y}) \quad (\text{A.3})$$

$$\mathbf{V}_N = (\mathbf{V}_0^{-1} + \mathbf{X}^\top \mathbf{X})^{-1} \quad (\text{A.4})$$

$$a_N = a_0 + n/2 \quad (\text{A.5})$$

$$b_N = b_0 + \frac{1}{2} (\mathbf{w}_0^\top \mathbf{V}_0^{-1} \mathbf{w}_0 + \mathbf{y}^\top \mathbf{y} - \mathbf{w}_N^\top \mathbf{V}_N^{-1} \mathbf{w}_N) \quad (\text{A.6})$$

where \mathbf{V}_0 , \mathbf{w}_0 , a_0 and b_0 are hyperparameters of the prior. It is possible to set
 698 an uninformative prior for σ^2 by applying $a_0 = b_0 = 0$. Also setting $\mathbf{w}_0 = 0$
 and $\mathbf{V}_0 = g(\mathbf{X}^\top \mathbf{X})^{-1}$ for any positive value g ; leads to Zellner’s *g-prior* [31].
 700 By having the prior variance proportional to $(\mathbf{X}^\top \mathbf{X})^{-1}$, it is ensured that the
 posterior is invariant to scaling of the inputs.

702 **B. Gaussian Process**

704 Conceptually, one can think of this process as estimating, rather than one
 “best” fit through the data, a distribution over all the possible functions that
 could explain the data. By virtue of its construction, the marginal and conditional
 706 distributions of any finite subset of data points in the function is Gaussian.
 In other words, any set of data modelled by the Gaussian process can be
 708 represented by a multivariate Gaussian distribution. The benefit of this result
 is that computations are normally available in closed form; for example, the
 710 conditional distribution of some new test points given the already observed
 data can be recovered exactly. The model is also *nonparametric*; the form of
 712 the function which will fit the data does not need to be specified, i.e. it is not
 necessary to choose a basis, such as a polynomial one. Instead, the function
 714 is modelled by representing the covariance in the data through a kernel or
 covariance function. This kernel is used to embed belief about which *family*
 716 *of functions* the data have come from, e.g. a nonlinear or periodic function.

The Gaussian process can be used to model nonlinear regression problems of the form,

$$\mathbf{y} = f(\mathbf{X}) + \varepsilon \quad \varepsilon \sim \mathcal{N}(\mathbf{0}, \sigma_n^2 \mathbb{I}) \quad (\text{B.1})$$

718 where \mathbf{y} is a vector of N observed targets, \mathbf{X} a matrix of N observed inputs in
 D dimensions, and ε a vector of realisations from a zero-mean Gaussian white
 noise process with variance σ_n^2 .

A GP is fully defined by its mean and covariance function,

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')) \quad (\text{B.2})$$

The mean function, $m(\mathbf{x})$ can be any parametric mapping of \mathbf{x} , e.g. a polynomial. The correlation between the targets is captured by the covariance function which expresses the similarity between two input vectors \mathbf{x} and \mathbf{x}' . To predict at a new test point \mathbf{x}_* , or set of test points X_* , predictive equations are used to determine the expected mean function $\mathbb{E}[f_*]$ and expected covariance $\mathbb{V}[f_*]$ [23],

$$f_* \sim \mathcal{N}(\mathbb{E}[f_*], \mathbb{V}[f_*]) \quad (\text{B.3a})$$

$$\mathbb{E}[f_*] = m(\mathbf{x}_*) + k(\mathbf{x}_*, X)(k(X, X) + \sigma_n^2 \mathbb{I})^{-1} \mathbf{y} \quad (\text{B.3b})$$

$$\mathbb{V}[f_*] = k(\mathbf{x}_*, \mathbf{x}_*) - k(\mathbf{x}_*, X)(k(X, X) + \sigma_n^2 \mathbb{I})^{-1} k(X, \mathbf{x}_*) \quad (\text{B.3c})$$

720 If predicting at noisy output locations, i.e. y_* , it is trivial to add the noise
 variance $\sigma_n^2 \mathbb{I}$ to the predictive covariance in Eq. (B.3c). As such the GP returns
 722 the posterior distribution over f_* or y_* as a Gaussian distribution.

724 For practical implementation of a Gaussian process, the reader is recommended
 to follow the guidance of Rasmussen [23]. The primary influence of the user
 when implementing a GP is in the choice of the kernel, which is calculated as
 726 any other kernel; linear pair-wise distances between points to form a covariance
 matrix. Careful consideration of data should also be applied in implementation,
 728 such as data type (scale, sign, etc.), data size and space on which it operates.

Standard practice to determine the hyperparameters of a Gaussian process
 730 is to maximise the *marginal likelihood*, which in practice is done by minimising
 the *log marginal likelihood* (NLML),

$$\hat{\Theta} = \arg \min(-\log p(\mathbf{y}|\Theta)) \quad (\text{B.4})$$

732 where the negative log marginal likelihood of the Gaussian process is given by,

$$-\log p(\mathbf{y}|\mathbf{X}, \Theta) = \frac{1}{2} \log |K_y| + \frac{1}{2} \mathbf{y}^\top K_y^{-1} \mathbf{y} + \frac{N}{2} \log(\sigma_n^2) \quad (\text{B.5})$$

734 When defining $K_y = K(X, X) + \sigma_n^2 \mathbb{I}$, $K(X, X)$ is the pairwise covariance matrix of all of the training inputs and N is the number of data points.

C. Polar GP

To make this modification is not as trivial as it may seem. Remembering that the covariance function is a measure of similarity between two data points it is necessary to define a kernel which encodes this. Specifically, it is necessary to have high covariance between points that are close to each other in angle. For example a point with angle 359° should have a high covariance with 1° if the radii are also close. This will require modifications to the kernel in terms of the distance used to assess how close points are together and also to the covariance function itself. Padonou and Roustant [25] suggest two potential definitions for a distance which fulfils this criteria, full details of setting up a polar coordinates kernel can be found in that work but it is briefly reviewed here. These two distances are: the *chordal* distance $d_1(\theta, \theta') = 2 \sin\left(\frac{\theta - \theta'}{2}\right)$ or the *geodesic* distance $d_2(\theta, \theta') = \arccos(\cos(\theta - \theta'))$. Using these definitions, it is possible to define the covariance in the θ dimension of a $\{\rho, \theta\}$ polar coordinate system. The C^2 -Wendland function is used as the kernel, since this produces a covariance of 0 when $d_2 = \pi$ and is strictly positive when $d_2 > \pi$, both necessary conditions for the polar kernel design. The C^2 -Wendland function is defined as,

$$W_c(t) = \left(1 + \tau \frac{t}{c}\right) \left(1 - \frac{t}{c}\right)_+^\tau, \quad c \in [0, \pi]; \tau \geq 4 \quad (\text{C.1})$$

When applying the Wendland function as the covariance function, the value of c must change depending on the angular distance chosen,

$$k_W = \begin{cases} k_{chord}(\theta, \theta') = W_2(d_1(\theta, \theta')) \\ k_{geo}(\theta, \theta') = W_\pi(d_2(\theta, \theta')) \end{cases} \quad (\text{C.2})$$

736 Here, the value of τ acts as a ‘steepening’ parameter on the angular covariance; this can be seen as the angular analogue to the inverse of the length scale
738 parameter described previously.

To form a full polar covariance function, a different kernel is applied only on
740 the radial dimension of the input. This kernel could be any stationary isotropic kernel; in this work the Matérn 5/2 kernel is used as in [25]. In that case the
742 distance used in the Matérn kernel is the absolute difference between the two radial components $|\rho - \rho'|$. For the angular component, Eq. (C.1) is used with
744 the geodesic distance such that kernel k_{geo} is used.

These choices define the covariance in the model along each of the directions - the radial ρ and the angular θ . To form the total covariance it is necessary to combine these two. It is known that the addition or pointwise multiplication of any two valid covariance functions is itself a valid covariance [23]. In this work an ANOVA combination [32] of the kernels in each dimension is used, as in [25]. This allows variations in each dimension as well as the combination to

contribute to variation in the function. The combined *ANOVA* kernel is defined as,

$$k_2(\mathbf{x}, \mathbf{x}') = \sigma_f^2 (1 + \sigma_{f,r}^2 k_{\text{mat}}(\rho, \rho')) (1 + \sigma_{f,a}^2 k_{\text{W}}(\theta, \theta')) \quad (\text{C.3})$$

where $\sigma_{f,m}$ and $\sigma_{f,a}$ act as weights representing the influence of changes in each dimension on a change in the output.

D. Kernels

Many of the kernels shown in this paper use the distance between data points, r , for their calculation; for this paper, $\mathbf{x} - \mathbf{x}'$ is the Euclidean distance when data is in Cartesian coordinates or applied to ρ in polar coordinates and is defined by $d_2 = \arccos(\cos(\theta - \theta'))$ when applied to θ in polar coordinates.

D.1. Squared Exponential

The squared exponential is a general nonlinear kernel which operates in a D -dimensional real space, where r is the distance between two points \mathbf{x} & \mathbf{x}' .

$$k_{\text{sqe}}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left(-\frac{r^2}{2\ell^2}\right) \quad (\text{D.1})$$

Hyperparameters. ℓ is the length scale, σ_f^2 is the scaling factor.

D.2. Matérn Class

The Matérn class of kernels are general nonlinear kernels which operate in a D -dimensional real space, and are specified by a scaling factor ν ; in this case, $\nu = 5/2$ is used,

$$k_{\text{mat},\nu=5/2}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \left(1 + \frac{\sqrt{5}r}{\ell} + \frac{5r^2}{3\ell^2}\right) \exp\left(-\frac{\sqrt{5}r}{\ell}\right) \quad (\text{D.2})$$

Hyperparameters. ℓ is the length scale, σ_f^2 is the scaling factor.

D.3. Polynomial

The polynomial kernel is an inhomogeneous linear kernel which operates in a D -dimensional real space, and is specified by the coefficient p ,

$$k_{\text{pol}}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 (\mathbf{x} \cdot \mathbf{x}' + \sigma_0^2)^p \quad (\text{D.3})$$

Hyperparameters. $\sigma_0^2 \geq 0$ is the offset term trading off the influence of higher-order versus lower-order terms, σ_f^2 is the scaling factor.

D.4. C^2 -Wendland

The C^2 -Wendland kernel is a stationary kernel which operates on the θ -dimension of a real polar space, and was shown for use in Gaussian process modelling in [25]. $c = 2$ when $d = d_1(\theta, \theta') = 2 \sin((\theta - \theta')/2)$ and $c = \pi$ when $d = d_2(\theta, \theta') = \arccos(\cos(\theta - \theta'))$,

$$k_{\text{W}}(\theta, \theta') = \sigma_f^2 \left(1 + \tau \frac{d}{c}\right) \left(1 - \frac{d}{c}\right)^\tau \quad (\text{D.4})$$

772 *Hyperparameters.* $\tau \geq 4$ is the variance steepening parameter, σ_f^2 is the scaling factor.

D.5. Exponential Decay

774 The exponential decay kernel is a non-general linear kernel which operates in a D -dimensional real space,

$$k_{\text{ed}}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp(-\mathbf{x}l) \cdot \exp(-\mathbf{x}'^\top l) \quad (\text{D.5})$$

776 *Hyperparameters.* l is the length scale, σ_f^2 is the scaling factor.

D.6. Strictly Periodic

778 The strictly-periodic kernel is a non-general linear kernel which operates on the θ -dimension of a real polar space; the nature is enforced by the number of
780 equally distributed symmetry axes, n ,

$$k_{\text{sym}}(\theta, \theta') = \sigma_f^2 (\alpha_1 + \alpha_2 \cos(2nd_2(\theta, \theta'))) \quad (\text{D.6})$$

782 *Hyperparameters.* α_1 & α_2 are the relative weighting of the maximum value and periodicity of the function respectively.

D.7. Generic Polar

784 The generic polar kernel is an ANOVA combination of k_{mat} and k_W which operates on the $\{\theta, \rho\}$ -dimension real polar space,

$$k_2(\mathbf{x}, \mathbf{x}') = \sigma_f^2 (1 + \sigma_{f,r}^2 k_{\text{mat}}(\rho, \rho')) (1 + \sigma_{f,a}^2 k_W(\theta, \theta')) \quad (\text{D.7})$$

786 *Hyperparameters.* $\sigma_{f,r}^2$ & $\sigma_{f,a}^2$ are the relative importance of changes in the radial & angular dimension respectively, σ_f^2 is the overall scaling factor.

788 D.8. Angular Informed

790 The angular informed kernel is an OR combination of k_{sqe} and k_{sym} which operates on the θ -dimension of a real polar space,

$$k_{\text{ang}}(\mathbf{x}, \mathbf{x}') = \sigma_{f,\text{sqe}}^2 k_{\text{sqe}}(\theta, \theta') + \sigma_{f,\text{sym}}^2 k_{\text{sym}}(\theta, \theta') \quad (\text{D.8})$$

Hyperparameters. $\sigma_{f,\text{sqe}}^2$ & $\sigma_{f,\text{sym}}^2$ are the relative scaling factors for each kernel.

792 D.9. Radial Informed

794 The radial informed kernel is an AND combination of k_{pol} , $p = -1/2$ and k_{ed} which operates on the ρ -dimension of a real polar space,

$$k_{\text{rad}}(\mathbf{x}, \mathbf{x}') = \sigma_{f,r}^2 k_{\text{pol}}(\rho, \rho') \cdot k_{\text{ed}}(\rho, \rho') \quad (\text{D.9})$$

Hyperparameters. σ_f^2 is the overall scaling factor.

796 *D.10. Informed Guided Wave, smooth & strictly periodic*

798 The first UGW-informed kernel is an ANOVA combination of k_{ang} and k_{rad} which operates on the $\{\theta, \rho\}$ -dimension of a real polar space. The periodicity in θ is enforced and smooth,

$$k_3(\mathbf{x}, \mathbf{x}') = \sigma_f^2 (1 + \sigma_{f,a}^2 k_{\text{ang}}(\theta, \theta')) (1 + \sigma_{f,r}^2 k_{\text{rad}}(\rho, \rho')) \quad (\text{D.10})$$

800 *Hyperparameters.* $\sigma_{f,a}^2$ & $\sigma_{f,r}^2$ are the relative importance of changes in the angular & radial dimension respectively, σ_f^2 is the overall scaling factor.

802 *D.11. Informed Guided Wave, smooth & strictly periodic*

804 The second UGW-informed kernel is an ANOVA combination of k_{W} and k_{rad} which operates on the $\{\theta, \rho\}$ -dimension of a real polar space. The periodicity in θ is enforced but not necessarily smooth,

$$k_4(\mathbf{x}, \mathbf{x}') = \sigma_f^2 (1 + \sigma_{f,a}^2 k_{\text{W}}(\theta, \theta')) (1 + \sigma_{f,r}^2 k_{\text{rad}}(\rho, \rho')) \quad (\text{D.11})$$

806 *Hyperparameters.* $\sigma_{f,a}^2$ & $\sigma_{f,r}^2$ are the relative importance of changes in the angular & radial dimension respectively, σ_f^2 is the overall scaling factor.

808 E. Performance Metrics

810 It is necessary at this point to develop some metrics by which the models can be assessed. For this purpose, the experimental dataset was split into a training set, \mathbf{x}_t and a test set \mathbf{x}_v , which consisted of 75% and 25% of the total dataset respectively. The performance of each model is reported on both the training and the test data; it is important to consider the test data performance, as this is the best indicator of which models are able to *generalise*, i.e. which will work best on new unseen data.

816 The first metric used here is the *normalised mean squared error* (NMSE) which can be computed for both training set ($NMSE_{tr}$) and test set ($NMSE_t$). For descriptive purposes, the NMSE indicates how well the estimate of the output fits with the observed values. The NMSE will return a score of zero when the predicted values are identical to those measured (this is impossible in the presence of any noise). A score of 100 is equivalent to simply taking the mean of the observed data as the prediction at every instance. The calculation for the NMSE is,

$$NMSE = \frac{100}{n\sigma_y^2} (\mathbf{y} - \mathbf{y}^*)^\top (\mathbf{y} - \mathbf{y}^*) \quad (\text{E.1})$$

824 where \mathbf{y} is the vector of observed outputs and \mathbf{y}^* the predicted outputs (in this case the mean of the predictive distribution from a Gaussian process). n is the number of observations in \mathbf{y} and σ_y^2 the variance of those observations.

826 The second metric will be to compare the predictive likelihoods of the model. This metric can be a more informative way of assessing the models as it takes into account the uncertainty in the prediction as well as the quality of the mean fit. The predictive likelihood is given as $p(\mathbf{y}^* | \mathbf{x}^*, \mathbf{y}, \mathbf{x})$; this will change

dependent upon the model being assessed but since all models in this work have a tractable Gaussian posterior it is given here by,

$$p(\mathbf{y}^*|\mathbf{x}^*, \mathbf{y}, \mathbf{x}) = \mathcal{N}(\mathbb{E}[\mathbf{y}^*], \mathbb{V}[\mathbf{y}^*]) \quad (\text{E.2})$$

This work will use this quantity in two ways; the first will consider each prediction to be independent, by not including the cross covariance terms in $\mathbb{V}[\mathbf{y}^*]$ the predictive variance matrix. The full covariance of the prediction will also be considered from the Gaussian process, as this can give better insight into how well the function has been modelled.

To ensure computational stability these likelihood estimates are both calculated in the log space. The first quantity will be referred to as the independent predictive log likelihood PLL_i , and is defined by,

$$PLL_i = \sum_i^N \log \mathcal{N}(\mathbf{y}_i|\mathbb{E}[\mathbf{y}_i], \mathbb{V}[\mathbf{y}_i], \Theta) \quad (\text{E.3})$$

for N data points. This is the product over the predictive likelihoods for every point, i.e. the joint likelihood if they were uncorrelated. The second will be considered the co-dependent predictive log likelihood PLL_c , defined by,

$$PLL_c = \log p(\mathbf{y}|\mathbb{E}[\mathbf{y}], \mathbb{V}[\mathbf{y}], \Theta) \quad (\text{E.4})$$

where PLL_c is computed as the likelihood of the full multivariate Gaussian over the predictive points, including the predicted covariance between those points.

References

- [1] N. Dowling, Mechanical Behavior of Materials: Engineering Methods for Deformation, Fracture, and Fatigue, Pearson, 2012.
- [2] P. Schubel, R. Crossley, Wind turbine blade design, *Energies* 5 (2012) 3425–3449.
- [3] M. Wymore, J. V. Dam, H. Ceylan, D. Qiao, A survey of health monitoring systems for wind turbines, *Renewable and Sustainable Energy Reviews* 52 (2015) 976–990.
- [4] C. R. Farrar, K. Worden, Structural Health Monitoring.: A Machine Learning Perspective, John Wiley & Sons, 2012.
- [5] H. Mao, S. Mahadevan, Fatigue damage modelling of composite materials, *Composite Structures* 58 (2002) 405–410.
- [6] A. Orifici, I. Herszberg, R. Thomson, Review of methodologies for composite material modelling incorporating failure, *Composite structures* 86 (2008) 194–210.
- [7] I. Viktorov, Rayleigh and Lamb Waves: Physical Theory and Applications, Plenum Press, 1967.
- [8] E. Adler, Matrix methods applied to acoustic waves in multilayers, *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 37 (1990) 485–490.

- 854 [9] G. Mook, C. Willberg, U. Gabbert, J. Pohl, Lamb wave mode conversion
856 in CFRP plates, 11th European Conference on Non-Destructive Testing
(2014).
- [10] C. Willberg, S. Koch, G. Mook, J. Pohl, U. Gabbert, Continuous mode
858 conversion of Lamb waves in CFRP plates, *Smart Materials and Structures*
21 (2012) 1–9.
- 860 [11] L. Wang, S. I. Rokhlin, Ultrasonic wave interaction with multidirectional
composites: Modeling and experiment, *The Journal of the Acoustical*
862 *Society of America* 114 (2003) 2582–2595.
- [12] O. Simeone, A brief introduction to machine learning for engineers,
864 *Foundations and Trends® in Signal Processing* 12 (2018) 200–431.
- [13] J. Rose, Ultrasonic guided waves in structural health monitoring, *Key*
866 *Engineering Materials* 270-273 (2004) 14–21.
- [14] K. Worden, Rayleigh and Lamb Waves - Basic Principles, *Strain* 37 (2001)
868 167–172.
- [15] A. Pollock, Classical wave theory in practical AE testing, *Progress in*
870 *Acoustic Emission III: Proceedings of the 8th International Acoustic*
Emission Symposium (1986).
- 872 [16] C. Ramadas, K. Balasubramaniam, A. Hood, M. Joshi, C. V.
Krishnamurthy, Modelling of attenuation of Lamb waves using Rayleigh
874 damping: Numerical and experimental studies, *Composite Structures* 93
(2011) 2020 – 2025.
- 876 [17] K. J. Schubert, A. S. Herrmann, On attenuation and measurement of Lamb
waves in viscoelastic composites, *Composite Structures* 94 (2011) 177–185.
- 878 [18] M. J. S. Lowe, Matrix techniques for modeling ultrasonic waves in
multilayered media, *IEEE Transactions on Ultrasonics, Ferroelectrics, and*
880 *Frequency Control* 42 (1995) 525–542.
- [19] D. N. Alleyne, P. Cawley, The interaction of Lamb waves with defects,
882 *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*
39 (1992) 381–397.
- 884 [20] N. Guo, P. Cawley, The interaction of Lamb waves with delaminations in
composite laminates, *The Journal of the Acoustical Society of America* 94
886 (1993) 2240–2246.
- [21] S. da Silva, Data-driven model identification of guided wave propagation
888 in composite structures, *Journal of the Brazilian Society of Mechanical*
Sciences and Engineering 40 (2018) 543.
- 890 [22] K. Murphy, *Machine Learning: A Probabilistic Perspective*, The MIT
Press, 2012.
- 892 [23] C. Rasmussen, C. Williams, *Gaussian Processes for Machine Learning*
(Adaptive Computation and Machine Learning), The MIT Press, 2005.

- 894 [24] A. O'Hagan, Curve fitting and optimal design for prediction, *Journal of the Royal Statistical Society: Series B (Methodological)* 40 (1978) 1–24.
- 896 [25] E. Padonou, O. Roustant, Polar Gaussian processes for predicting on circular domains (Feb. 2015).
- 898 [26] D. Duvenaud, Automatic model construction with Gaussian processes, Ph.D. thesis, University of Cambridge (2014).
- 900 [27] C. E. Rasmussen, Z. Ghahramani, Occam's razor, *Advances in Neural Information Processing Systems* (2001).
- 902 [28] S. Yang, M. Wang, et al., A quantum particle swarm optimization, *Proceedings of the 2004 Congress on Evolutionary Computation* (2004).
- 904 [29] T. J. Rogers, G. Manson, K. Worden, E. J. Cross, On the choice of optimisation scheme for Gaussian process hyperparameters in SHM problems, *Structural Health Monitoring 2017: Real-Time Material State Awareness and Data-Driven Safety Assurance - Proceedings of the 11th International Workshop on Structural Health Monitoring* (2017).
- 906
- 908 [30] K. Kersting, C. Plagemann, P. Pfaff, W. Burgard, Most likely heteroscedastic Gaussian process regression, *Proceedings of the 24th International Conference on Machine Learning* (2007).
- 910
- 912 [31] A. Zellner, On assessing prior distributions and bayesian regression analysis with g prior distributions, in: P. Goel, A. Zellner (Eds.), *Bayesian Inference and Decision Techniques, Studies of Bayesian and Econometrics and Statistics*, Elsevier Science Ltd., 1986, pp. 233–243.
- 914
- 916 [32] G. Wahba, *Spline Models for Observational Data*, Society for Industrial and Applied Mathematics, 1990.